BROWNIAN SMOOTHING VIA THE DRIVING NOISE

Anonymous authors
Paper under double-blind review

ABSTRACT

Amortized posterior sampling offers a promising route to accurate and efficient smoothing. Learned smoothers can model complex trajectory posteriors, while drawing a sample still often requires full-trajectory modeling or a large population of particles. In this paper, we introduce Brownian Smoothing, which generates complete posterior trajectories in one rollout by conditioning the initial state and shifting the simulator's Gaussian driving noise. To maintain accuracy under one-rollout sampling, Brownian Smoothing is trained by conditional maximum likelihood on model simulations. And we prove that this objective is a forward-KL projection and that the optimal drift correction converges to the one induced by Doob's h-transform under grid refinement. Empirically, Brownian Smoothing closely matches exact posterior statistics and achieves the best marginal and joint scores on nonlinear, high-dimensional benchmarks. It also substantially reduces sampling time and memory compared with global flow and score-based baselines.

1 INTRODUCTION

Smoothing combines dynamical models with noisy observations to reconstruct how a system evolved over time (Briers et al., 2010; Särkkä, 2013). It uses the full observation record and plays an important role when measurements provide only a partial view of the system. For example, in tracking and navigation, smoothing helps recover an object's past motion from sensor measurements (Särkkä, 2013). In geophysics, it is used to reconstruct physical fields from sparse observations (Browne, 2016; Rozet & Louppe, 2023). In neuroscience, it helps infer latent neural dynamics from recorded activity and it can quantify uncertainty in these latent states as well (Lawson et al., 2023).

Modern smoothing took shape in the 1960s, alongside the progress of Kalman filtering (Särkkä, 2013). Rauch, Tung, and Striebel showed how a backward pass could refine an entire sequence of filtered estimates, giving an exact solution for linear Gaussian systems (Rauch et al., 1965). As applications expanded to nonlinear dynamics and non-Gaussian noise, researchers turned to simulation to move beyond these closed-form solutions. In the 1990s, Kitagawa developed a Monte Carlo filter and smoother that represented distributions through simulated samples rather than a Gaussian approximation (Kitagawa, 1996). Godsill, Doucet, and West later developed a particle-based forward–backward procedure for sampling complete state histories and tested it on noisy speech signals (Godsill et al., 2004). But such developments brought more complex time series within reach of smoothing methods and this broader scope also brought a high computational cost, because particle methods may need large populations to maintain accuracy, while path-space Markov chain Monte Carlo methods may need many updates to explore the posterior (Chopin & Papaspiliopoulos, 2020; Roberts & Stramer, 2001; Papaspiliopoulos et al., 2013). In both cases, much of the computation is tied to a single observation record and must be repeated whenever a new observation record arrives.

Learning from simulated data offers a way to share the work of inference across observation records. Amortized methods train a conditional model on many simulated records and reuse it for new observations (Papamakarios & Murray, 2016; Cranmer et al., 2020; Radev et al., 2022). The cost of drawing each trajectory remains a separate challenge. For example, a global conditional transport models all time steps together, so its input dimension grows with the horizon and state size (Wildberger et al., 2023). Score-based smoothers also work with complete trajectories and update them repeatedly to obtain a posterior sample (Rozet & Louppe, 2023). Learned particle smoothers use sequential sampling with a population of particles, weighting, and resampling (Lawson et al., 2022; 2023). These sampling costs motivate an amortized smoother that generates accurate trajectories in a single simulator rollout. Appendix B provides a more detailed comparison of related approaches.

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.

We introduce Brownian Smoothing, which generates a complete posterior trajectory in one simulator rollout. Sampling requires no global trajectory model, repeated full-trajectory updates, or particle population. This avoids the extra work of repeatedly processing whole trajectories or maintaining many particles to obtain a sample. The method learns an observation-conditioned initial distribution and shifts the simulator's Gaussian driving noise at each step. The supplied dynamics map these inputs to a complete, equally weighted trajectory, with the drift and diffusion functions kept fixed. The learned initial law and noise corrections can then be reused across unseen observation records.

Training Brownian Smoothing is direct. Each model simulation records the initial state and Gaussian driving noise, which directly supervise both components. Conditional maximum likelihood requires neither posterior samples nor the observation density. The learned initial law and noise shifts then generate one equally weighted trajectory per rollout. This objective is a forward KL projection; under a curvature condition, its Gaussian residual is $O(|\Delta|)$. As the grid is refined, the optimal discrete correction converges to the continuous-time drift correction induced by Doob's h-transform.

Empirically, Brownian Smoothing produces accurate posterior trajectories across exact, nonlinear, and high-dimensional systems. It matches exact Gaussian statistics and approaches particle smoothing with 4,096 particles on Double Well. On Lorenz–96 and Kolmogorov flow, it achieves the lowest RMSE, marginal CRPS, and joint trajectory score among the evaluated methods. Fixing X_0 shows that the learned correction to the simulator's Gaussian driving noise captures observation information that cannot be absorbed into the initial state. With one simulator rollout per sample, Brownian Smoothing is also substantially faster and uses less memory than global flow and score baselines.

2 Problem Statement

We study a stochastic system whose latent state $X_t \in \mathbb{R}^d$ evolves on $[0, T]$ according to

$$dX_t = b(t, X_t)\, dt + \sigma(t, X_t)\, dW_t, \qquad X_0 \sim p_0, \tag{1}$$

where W is r-dimensional Brownian motion and $\sigma(t, x) \in \mathbb{R}^{d\times r}$ can be rectangular or rank deficient. We assume a unique, nonexplosive solution (Bain & Crisan, 2009; Särkkä, 2013). The state is observed through noisy, partial measurements $Y_m \sim g_m(\cdot \mid X_{t_m})$ at $0 < t_1 < \cdots < t_M \le T$, independently given the trajectory. The dynamics and observation model are known; their dependence on the supplied parameter λ is suppressed throughout.

From these measurements, we seek to reconstruct how the latent state evolved and quantify uncertainty in that reconstruction. Writing $Y = ((t_m, Y_m))_{m=1}^M$ for the full record and y for its realization, this amounts to inferring the joint smoothing posterior $\mathcal{L}(X_{[0,T]} \mid Y = y)$. This posterior uses all available observations to describe both uncertainty in each state and dependence across time.

We pursue this goal numerically by discretizing the dynamics on an Euler grid $0 = t_0 < \cdots < t_K = T$ containing all observation times. With $\Delta_k = t_{k+1} - t_k$, the initial state $X_0 \sim p_0$ and independent Gaussian inputs $\varepsilon_k \sim \mathcal{N}(0, I_r)$ generate

$$X_{k+1} = X_k + b(t_k, X_k)\Delta_k + \sigma(t_k, X_k)\sqrt{\Delta_k}\,\varepsilon_k. \tag{2}$$

The resulting joint posterior $\mathcal{L}(X_{0:K} \mid Y = y)$ is the target of our method, providing state estimates and uncertainty on the chosen grid.

3 Brownian Smoothing

The Euler recursion suggests a direct way to represent this posterior: once the initial state and Gaussian inputs are fixed, the entire trajectory is determined. Let $\Phi(x_0, \varepsilon_{0:K-1})$ denote this simulator map, where $\varepsilon_{0:K-1} = (\varepsilon_0, \ldots, \varepsilon_{K-1})$ and the input prior is $\mu = p_0 \otimes \mathcal{N}(0, I_r)^{\otimes K}$. Applying the same map to the posterior inputs gives

$$\mathcal{L}(X_{0:K} \mid Y = y) = \Phi_{\#}\mathcal{L}(X_0, \varepsilon_{0:K-1} \mid Y = y), \tag{3}$$

where $\Phi_{\#}$ denotes pushforward. Brownian Smoothing builds on this identity by learning how observations change the initial state and the Gaussian inputs while keeping the simulator fixed. Each joint simulation records $(X_0, \varepsilon_{0:K-1}, Y)$, so these inputs are available for training even when Φ has no inverse. The learned conditional model then supplies the inputs for one-rollout trajectory generation.

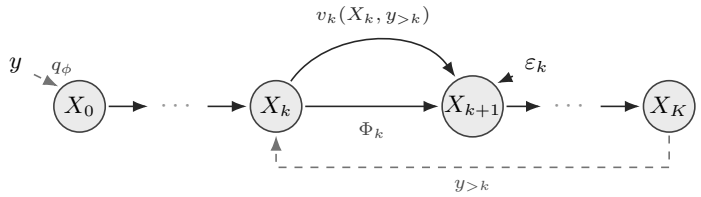


Figure 1: Brownian Smoothing sampler in one rollout. The record y conditions the initial state X_0. At step k, the current state X_k and future observations $y_{>k}$ determine a Gaussian mean shift, while fresh noise ξ_k and the supplied Euler dynamics generate X_{k+1}.

3.1 Posterior sampling in one rollout

The input model follows the simulator's sequential structure. The complete record y conditions the initial state through $q_\phi(x_0 \mid y)$. At step k, the current state X_k carries the effect of all previous inputs. Write $Y_{\le k}$ and $Y_{>k}$ for observations no later than and strictly after t_k, respectively. Given X_k, the Markov property leaves only future observations $y_{>k}$ to determine the next noise correction. Let $v_k(x, y_{>k}) \in \mathbb{R}^r$ denote the learned control. In order to generate a trajectory, sample the posterior simulator inputs as

$$\varepsilon_k = \xi_k + \sqrt{\Delta_k}\, v_k(X_k, y_{>k}), \quad k = 0, \dots, K-1. \tag{4}$$

where $X_0 \sim q_\phi(\cdot \mid y)$ and $\xi_k \sim \mathcal{N}(0, I_r)$. At each step, the shifted input ε_k is passed through the supplied Euler update in equation 2. The shift contributes the posterior drift correction $\sigma(t_k, X_k)v_k(X_k, y_{>k})$, while the supplied functions b and σ remain fixed. Repeating this update returns a complete, equally weighted trajectory in one simulator rollout. Thus, no likelihood evaluation, weighting, or resampling is required in Brownian Smoothing.

Let φ_r denote the standard Gaussian density on $\mathbb{R}^r$. The sampler then defines this density over the simulator inputs:

$$q(x_0, \varepsilon_{0:K-1} \mid y) = q_\phi(x_0 \mid y) \prod_{k=0}^{K-1} \varphi_r\Big(\varepsilon_k - \sqrt{\Delta_k}\, v_k(x_k, y_{>k})\Big), \tag{5}$$

where x_k is generated by the supplied Euler dynamics. The complete record conditions X_0, and each later factor learns a Gaussian mean shift in $\mathbb{R}^r$ from the current state and future observations. The factors retain identity covariance; Section 4 analyzes the approximation.

The sampler only needs q_ϕ and v_k, which joint model simulations train directly.

3.2 Learning from model simulations

Both q_ϕ and v_k are learned from model simulations recording $(X_0, \varepsilon_{0:K-1}, Y)$, exactly the variables modeled in equation 5. Training therefore requires no posterior samples.

We fit q by conditional maximum likelihood. Dropping terms that do not depend on the learned components gives the loss for one simulation:

$$\ell(X_0, \varepsilon_{0:K-1}, Y) = -\log q_\phi(X_0 \mid Y) + \sum_{k=0}^{K-1} \left[\frac{\Delta_k}{2} \|v_k\|^2 - \sqrt{\Delta_k}\, v_k^\top \varepsilon_k \right], \tag{6}$$

where $v_k = v_k(X_k, Y_{>k})$. Neither the posterior nor the observation density must be evaluated.

Minimizing equation 6 learns the posterior of X_0 and the Gaussian shifts implied by future observations. Completing the square gives the optimal control:

$$u_k(x, y_{>k}) = \frac{1}{\sqrt{\Delta_k}} \mathbb{E}[\varepsilon_k | X_k = x,\, Y_{>k} = y_{>k}]\,. \tag{7}$$

Thus, $\sqrt{\Delta_k} u_k(x, y_{>k})$ is the posterior mean shift of input k. Together, q_ϕ and the controls learn where the posterior path starts and how later observations correct the Gaussian driving noise.

In practice, we optimize equation 6 instead of directly regressing onto $\varepsilon_k/\sqrt{\Delta_k}$, whose variance grows as Δ_k^{-1}. For long trajectories, time indices can be sampled and weighted by the inverse sampling probability to obtain an unbiased loss estimate. Section 4 shows that the same population objective is a forward KL projection and characterizes the error introduced by the Gaussian factors.

3.3 Complete training and sampling procedure

Algorithm 1 separates amortized training from posterior sampling and gives the complete procedure.

Algorithm 1 Brownian Smoothing

▷ **Amortized Training**

Input: supplied model (p_0, b, σ, g), Euler grid $\{t_k\}_{k=0}^{K}$, initial density q_ϕ, control v_θ
for training step **do**
  Simulate a mini-batch $(X_0, \varepsilon_{0:K-1}, X_{1:K}, Y)$ from the supplied model
  Encode each observation record Y in reverse order
  Update q_ϕ and v_θ by minimizing the batch mean of equation 6
end for

▷ **Posterior Sampling**

Input: trained q_ϕ and v_θ, new observation record y, number of paths B
Encode y once
Draw $X_0^{(i)} \sim q_\phi(\cdot \mid y)$ for $i = 1, \ldots, B$
for Euler step $k = 0, \ldots, K-1$ **do**
  Draw fresh $\xi_k^{(i)} \sim \mathcal{N}(0, I_r)$ for every path
  Shift the Gaussian input: $\varepsilon_k^{(i)} = \xi_k^{(i)} + \sqrt{\Delta_k}\, v_\theta(X_k^{(i)}, y_{>k})$
  Advance every state with the supplied Euler dynamics in equation 2
end for
return B complete posterior trajectories with equal weights

Generating B posterior trajectories requires one record encoding, K batched control evaluations, and K batched simulator updates. All paths share the encoded summaries and advance together in one rollout. Global transports model the full trajectory, and score-based smoothers repeatedly update it; Brownian Smoothing applies each correction locally within the supplied simulator. Particle smoothers maintain weighted populations, while our paths are generated independently given the record, with no weighting or resampling. Thus, B sets the number of output samples rather than the size of an interacting particle population. This procedure reuses the trained model across records and avoids record-specific optimization and repeated full-trajectory processing at sampling time.

4 Theoretical Analysis

Brownian Smoothing samples in one rollout by conditioning the initial state and shifting only the mean of the simulator's Gaussian driving noise at each step. This restriction raises a direct question: which parts of the smoothing posterior can be recovered by mean shifts alone? We answer it in two steps. We first analyze the objective on a finite grid and separate errors in the learned initial law and controls from the Gaussian approximation error that training cannot remove. We then study grid refinement and identify the limiting control with Doob's h-transform.

4.1 Finite grid objective and approximation error

Fix a record y and an Euler grid. Let p and q denote the exact and learned conditional laws of the simulator inputs $(X_0, \varepsilon_{0:K-1})$. For each step, let $\pi_k(\cdot \mid x)$ be the exact conditional law of ε_k given $X_k = x$ and $Y_{>k} = y_{>k}$, and let $m_k(x)$ be its mean. The Markov property makes X_k sufficient for the previous inputs; Lemma 1 states this precisely. The population target in equation 7 gives $m_k(x) = \sqrt{\Delta_k}\, u_k(x)$, where we suppress the fixed argument for future observations.

Brownian Smoothing approximates $\pi_k(\cdot \mid x)$ by a Gaussian with mean $\sqrt{\Delta_k}\, v_k(x)$ and identity covariance. Even after matching the mean, a residual remains when the exact conditional law has a

different covariance or is non-Gaussian. Define this residual by

$$R(y) = \sum_{k=0}^{K-1} \mathbb{E}_p D_{\mathrm{KL}}(\pi_k(\cdot \mid X_k) \| \mathcal{N}(m_k(X_k), I_r)) .$$

We then separate this Gaussian approximation error from the two errors that training can reduce.

Theorem 1 (Forward KL decomposition on an Euler grid)**.** *For almost every record, the conditional second moments of the Gaussian increments are finite and, for every measurable initial density and control, the following identity holds in* $[0, \infty]$*:*

$$D_{\mathrm{KL}}(p\|q) = D_{\mathrm{KL}}(p(X_0)\|q(X_0)) + \frac{1}{2}\sum_{k=0}^{K-1} \Delta_k \mathbb{E}_p \left\|v_k(X_k) - u_k(X_k)\right\|^2 + R. \tag{8}$$

The decomposition separates two trainable errors from one irreducible residual. The first term is minimized when the learned initial law matches the posterior law of X_0. The second is minimized when $v_k = u_k$, so each control matches the posterior mean shift. These minimizers are unique up to posterior null sets. Whenever the posterior entropy and model cross entropy are finite, averaging equation 8 shows that conditional maximum likelihood is a forward KL projection. Appendix C.3 gives the proof and treats the extended real cases separately.

The residual R cannot be reduced by changing the initial law or the controls. Its size is instead controlled by the Euler step width.

Proposition 1. *Under the regularity conditions stated in Appendix C.3, suppose the Hessian of the log future likelihood along Brownian directions is uniformly bounded by* M*. Then*

$$\mathbb{E}_Y R(Y) \le \frac{rM^2}{2} \sum_{k=0}^{K-1} \Delta_k^2 \le \frac{rM^2 T}{2} |\Delta|. \tag{9}$$

At the population optimum, only R remains. Proposition 1 shows that this grid restriction vanishes in averaged forward KL under refinement. The rate is sharp for a scalar linear–Gaussian smoother; Appendix C.4 gives the example. The constant M may still depend on dimension, observation strength, horizon, and mixing. We next identify the continuous limit of the optimal discrete controls.

4.2 Continuous time limit

To identify this limit, fix a record y and let $h(t, x)$ be the likelihood of the observations after time t, conditional on $X_t = x$. Thus, h measures how strongly future observations favor each continuation of the current state. Appendix C.5.1 constructs the corresponding density process. With C denoting the normalizer, the continuous posterior has initial density and target control

$$p(x \mid y) = \frac{p_0(x)h(0, x)}{C}, \qquad u(t, x) = \sigma(t, x)^\top \nabla_x \log h(t, x). \tag{10}$$

The function h identifies the population target but is not evaluated by the algorithm. Brownian Smoothing instead learns both quantities in equation 10 from model simulations.

Let $v(t, x) = v_\theta(t, x, y_{>t})$, and let $q(\cdot \mid y)$ be the joint law obtained by drawing $X_0 \sim q_\phi(\cdot \mid y)$, solving

$$dX_t = [b(t, X_t) + \sigma(t, X_t)v(t, X_t)]\, dt + \sigma(t, X_t)\, dB_t, \tag{11}$$

and setting $W_t = B_t + \int_0^t v(s, X_s)\, ds$. Thus (X_0, W) records the initial state and driving noise of the same controlled path; $p(X_0)$ and $q(X_0)$ denote the corresponding initial marginals.

Proposition 2. *Under Assumption 1, replacing* v *by* u *in equation 11 gives the continuous smoothing law with the initial density in equation 10. Moreover,*

$$D_{\mathrm{KL}}(p(\cdot \mid y)\|q(\cdot \mid y)) = D_{\mathrm{KL}}(p(X_0)\|q(X_0)) + \frac{1}{2}\mathbb{E}_p \int_0^T \|v(t, X_t) - u(t, X_t)\|^2 dt. \tag{12}$$

The proof is given in Appendix C.7.

The continuous identity has the same interpretation: its terms measure errors in the initial state and the correction to the driving noise. Proposition 3 connects the discrete and continuous analyses.

Proposition 3. *Consider grids with $|\Delta| \to 0$ and indices such that $t_k \to t$. Under the consistency and uniform integrability conditions in Appendix C.3, the optimal discrete control satisfies*

$$u_k(x) \longrightarrow \sigma(t,x)^\top \nabla_x \log h(t,x) = u(t,x). \tag{13}$$

Proposition 3 identifies the population target with the correction induced by Doob's h-transform. The training objective targets this correction without evaluating h, solving its backward equation, or computing its gradient. The control remains in a driving space of dimension r and requires no inversion of $\sigma\sigma^\top$, including when σ is rectangular or rank deficient. These are population statements; they do not account for finite data, limited network capacity, or optimization error.

5 Numerical Studies

The preceding sections established the factorization, likelihood, and interpretation in continuous time. The experiments test its failure modes: linear–Gaussian checks separate discretization from learning, nonlinear systems test multimodality and scale, an exact wave posterior tests dependence in high dimensions, and interventions isolate the control contribution.

5.1 Analytic checks against exact smoothers

Before turning to nonlinear systems, we use three linear–Gaussian models with exact Rauch–Tung–Striebel smoothers: a scalar Ornstein–Uhlenbeck process, an integrated Langevin process with noise only in velocity, and a partially observed periodic field. They test regular dynamics, rank deficient diffusion, and spatial forcing of low rank. Brownian Smoothing trains on 10^5 prior predictive paths and returns 64 unweighted joint trajectories; Appendices A.3 and A.4 give the protocol. Figure 2 shows that field covariance remains close to the reference based on 64 draws, while posterior mean and marginal scale retain an amortization gap caused by finite data, capacity, and optimization.

The theoretical target is defined in continuous time, and training uses a finite Euler grid. Appendix A.5 therefore compares Euler smoothing with smoothing under the exact OU transition on the same 512 records: mean, marginal standard deviation, and mean control shift errors all decrease by about a factor of eight from $K = 32$ to $K = 256$. These checks establish agreement between the learned target and exact Gaussian smoothers across scalar, hypoelliptic, and spatial dynamics. We next test the factorization on nonlinear, non-Gaussian posteriors without scalable exact references.

5.2 Nonlinear benchmarks

We test three nonlinear regimes. Double Well asks whether local Gaussian shifts can produce multimodal paths with uncertain well occupancy and barrier crossing. Lorenz–96 tests the sequential factorization in a chaotic ring at $D = 40$ and $D = 320$, while a 64^2 Kolmogorov flow tests a large state driven by 32 Brownian modes. Every learned method receives 10^5 prior predictive paths per seed and returns 64 joint trajectories on shared records. Since no common exact posterior is available, the comparison uses trajectory metrics; Appendices A.1 and A.6 specify the protocol.

The baselines expose distinct approximation choices. Matched FMPE receives the same simulator inputs, observation encoder, and training pairs, then models their posterior with one global conditional transport (Lipman et al., 2023; Wildberger et al., 2023). SDA applies observation guidance to a learned trajectory score (Rozet & Louppe, 2023); LETKS extends localized ensemble-transform updates backward through time (Browne, 2016); and the particle smoother uses likelihood weights and resampling (Gordon et al., 1993; Godsill et al., 2004). Shared simulators, test records, and output counts align the task, while matched FMPE directly compares global posterior models.

Marginal metrics do not determine a trajectory law. The joint score retains time, directly for Double Well and through fixed projections in the state space for the systems with high dimension. The relative Euler residual checks innovation scale. Since its target value confirms only the diffusion scale, posterior conclusions combine it with the joint score, RMSE, and CRPS.

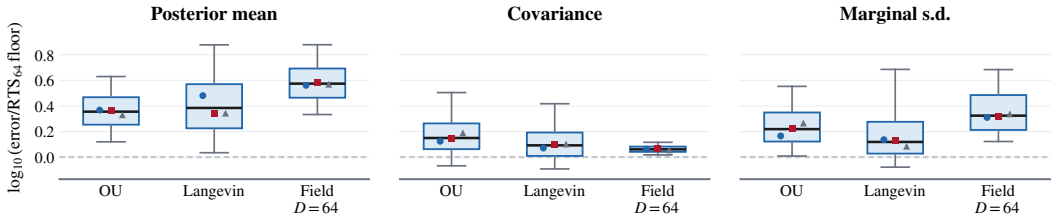


Figure 2: Errors for each record over seeds 0–2, with quartiles, 5th and 95th percentiles, and seed medians. Lower is better. The vertical axis shows $\log_{10}$ error normalized by Monte Carlo error from 64 exact RTS draws; zero is the finite sampling floor.

Table 1: **Nonlinear smoothing benchmarks.** Mean ± sample standard deviation over seeds 0–4. Lower is better except for the relative Euler residual, whose target is one. Joint trajectory quality uses the energy score for Double Well and the normalized sliced energy score elsewhere. Blue marks the best score and the residual closest to one within each setting. Runtime and VRAM measure 64 returned paths for one record with evaluation batch size 1.

Setting	Method	Marginal Quality		Trajectory Evidence		Efficiency	
		RMSE ↓	CRPS ↓	Joint traj. score ↓	Rel. Euler residual (≈ 1)	Runtime per record (s) ↓	VRAM (GB) ↓
Double Well	**Brownian Smoothing**	0.1532 ± 0.0084	0.0823 ± 0.0041	1.6859 ± 0.0725	1.0371 ± 0.0043	1.24	0.013
	FMPE (matched)	0.3856 ± 0.0412	0.1891 ± 0.0125	3.3743 ± 0.3150	1.0166 ± 0.0482	2.00	0.065
	SDA	0.3902 ± 0.0195	0.2096 ± 0.0284	7.4845 ± 0.6310	33.07 ± 14.28	2.61	0.159
	Bootstrap PF	0.1484 ± 0.0012	0.0794 ± 0.0008	1.6206 ± 0.0065	0.9998 ± 0.0015	1.13	0.032
Lorenz–96 $D = 40$, O50	**Brownian Smoothing**	2.1389 ± 0.0845	0.9825 ± 0.0412	1.4523 ± 0.0630	1.0001 ± 0.0028	1.30	0.042
	FMPE (matched)	2.5253 ± 0.1140	1.2401 ± 0.0685	1.8036 ± 0.0920	0.9993 ± 0.0315	18.92	0.278
	SDA	2.6101 ± 0.1650	1.5163 ± 0.0940	2.4403 ± 0.5280	410.3 ± 185.6	21.79	0.789
	LETKS	2.5274 ± 0.0480	1.1734 ± 0.0235	1.7938 ± 0.0145	18.96 ± 4.32	1.24	0.136
	Bootstrap PF	5.0929 ± 0.3210	3.9921 ± 0.2850	5.0164 ± 0.0890	1.0041 ± 0.0035	1.25	0.096
Lorenz–96 $D = 320$, O50	**Brownian Smoothing**	1.7912 ± 0.0620	0.7783 ± 0.0315	1.2537 ± 0.0480	1.0000 ± 0.0021	1.69	0.176
	FMPE (matched)	2.0627 ± 0.0890	0.9425 ± 0.0490	1.4505 ± 0.0710	1.0015 ± 0.0245	127.25	1.165
	SDA	3.3760 ± 0.1420	1.9129 ± 0.1150	2.4859 ± 0.1860	162.3 ± 68.5	74.68	3.306
	LETKS	2.5308 ± 0.0510	1.1761 ± 0.0280	1.8097 ± 0.0265	19.59 ± 5.12	1.49	0.570
	Bootstrap PF	5.3234 ± 0.2640	4.2320 ± 0.1980	5.2699 ± 0.1540	1.0009 ± 0.0031	1.50	0.402
Kolmogorov 64^2	**Brownian Smoothing**	0.4785 ± 0.0152	0.2671 ± 0.0094	0.3468 ± 0.0118	1.0007 ± 0.0019	5.18	1.179
	FMPE (matched)	0.4996 ± 0.0264	0.3031 ± 0.0142	0.3567 ± 0.0185	1.0438 ± 0.0385	23.18	3.889
	SDA	0.5998 ± 0.0345	0.2774 ± 0.0195	1.0834 ± 0.1420	423.6 ± 164.8	457.42	4.390
	LETKS	0.5990 ± 0.0185	0.2765 ± 0.0112	0.3561 ± 0.0160	43.69 ± 9.82	5.28	2.131
	Bootstrap PF	1.0318 ± 0.0960	0.8226 ± 0.0780	1.0137 ± 0.0540	1.0152 ± 0.0042	5.59	1.121

Table 1 gives the primary comparison. On Double Well, Brownian Smoothing is within 4% of the particle smoother in RMSE, CRPS, and joint score and improves on the learned references. It also obtains the lowest joint score on both Lorenz–96 settings and Kolmogorov flow, extending its marginal advantage to complete trajectories. Since matched FMPE uses the same inputs and training pairs, this comparison isolates the sequential factorization. Brownian Smoothing has a relative Euler residual between 1.000 and 1.037 throughout. FMPE and the particle smoother also decode paths through the simulator and remain near one, whereas updates in state space from LETKS and SDA can depart sharply from the stipulated diffusion scale.

We test temporal dependence directly by independently permuting the 64 path labels at every time while preserving every empirical time marginal; Appendix A.7 gives the full protocol. Marginal metrics change by at most 2.4×10^{-7}, but the mean joint score worsens from 1.5148 to 1.5272 on Lorenz–96 at $D = 40$ and from 0.34508 to 0.34569 on Kolmogorov flow. The same intervention raises lag-one correlation error from 0.00298 to 0.45295 on Lorenz–96 and temporal autocorrelation error from 5.3×10^{-5} to 0.82602 on Kolmogorov flow. The original paths perform better at all three checkpoints in both systems, showing that the joint score and correlation diagnostics capture temporal information absent from the unchanged marginal distributions at individual times.

The runtime results separate the cost of the posterior model from simulator cost. Matched FMPE makes 128 evaluations of the complete $d + Kr$ input vector, whereas SDA makes 128 evaluations of the complete $(K + 1)d$ state trajectory. In Lorenz–96 at $D = 320$, $r = d = 320$ and $K = 256$,

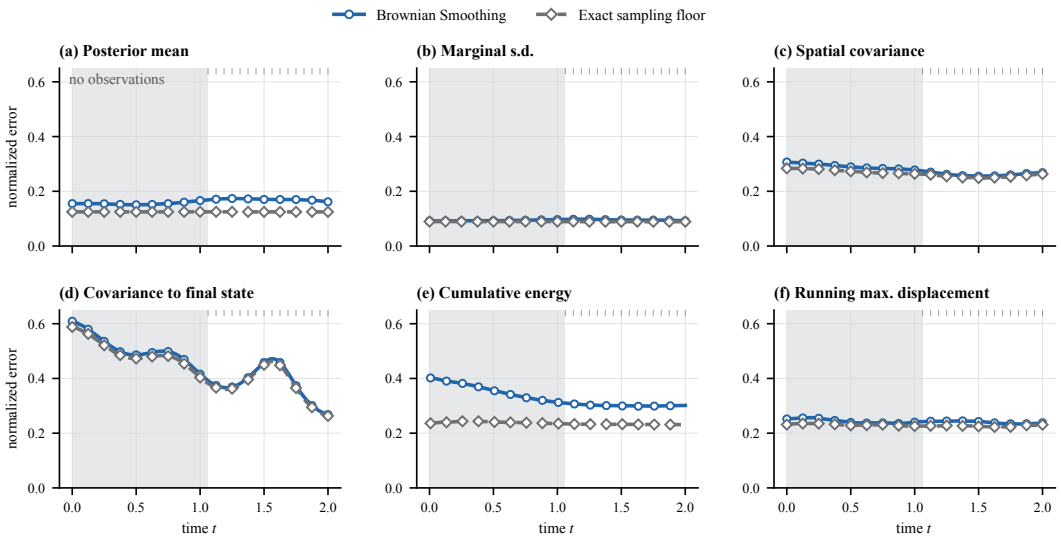


Figure 3: **Exact joint posterior validation over time.** Errors from 64 paths per record, averaged over 64 records and three checkpoints. Panels share one scale; gray marks the interval without observations, and top ticks mark the 16 observation times. Independent exact paths define the floor for 64 samples. Panel (d) reports $\mathrm{Cov}(X_t, X_T)$; panels (e)–(f) report cumulative functionals.

so both global objects contain 82,240 scalars. Brownian Smoothing instead predicts one local shift per step in one Euler rollout, taking 1.69 seconds rather than 127.25 for FMPE or 74.68 for SDA.

Kolmogorov flow reverses this dimensional balance: $d = 4096$, $r = 32$, and $K = 128$, so FMPE handles 8,192 inputs but SDA a trajectory containing 528,384 state values, taking 23.18 and 457.42 seconds. Brownian Smoothing predicts a correction of dimension 32, but its rollout still advances the full pseudospectral state and takes 5.18 seconds. The method removes repeated global transformations; the remaining runtime is the cost of one simulator rollout.

These measurements explain the runtime gains. Exact posterior error is unavailable for the nonlinear systems, so we next turn to a model in high dimensions with a tractable joint posterior.

5.3 Exact wave validation in high dimensions

Figure 3 compares Brownian Smoothing with an exact $D = 512$ Gaussian wave posterior over 131,584 trajectory coordinates. Averaged over time, errors in marginal scale, spatial covariance, covariance with the terminal state, and running maximum displacement are only 3–6% above the floor from exact samples. Posterior mean and cumulative energy remain 30% and 39% above the floor, exposing current approximation and training limits. Relative time evaluates one checkpoint over $T \in [2, 4]$, where moment quality is stable; the unseen $T = 5$ marks the extrapolation boundary.

This exact reference serves a different role from the controlled comparison in Table 1. Matched FMPE is omitted because 128 evaluations of a global transport over the vector $(X_0, \varepsilon_{0:255})$, which has 66,048 dimensions, exceeded our compute budget. Figure 3 therefore measures the absolute error of Brownian Smoothing, and Table 1 supplies the comparison with other methods. With distributional accuracy established, two questions remain: whether the samples recover nonlinear events along a path and whether the controls carry information after the initial state is known.

5.4 Exact path quantities and control contribution

These audits test path events and control. Double Well measures well occupancy, barrier crossing, and hitting time; wave fixes X_0 to isolate the control. Table 2 uses 512 Double Well records across five seeds and 32 wave records across three checkpoints; Appendices A.9 and A.8 give details.

On Double Well, Brownian Smoothing ranks second in all three event errors with 64 paths from one rollout. The particle smoother leads using 4,096 particles weighted by likelihood for each record. Amortized correction therefore approaches much costlier inference and extends the marginal and joint results in Table 1 to trajectory events.

The event comparison establishes accuracy of the complete trajectories, after which the wave intervention isolates the mechanism that produces their conditional evolution. We supply the simulated true X_0 and evaluate learned and zero control against $p(X_{1:K} \mid X_0, y)$. The learned control reduces posterior mean error from 1.1375 to 0.1646, temporal covariance error from 0.3687 to 0.1754, and energy error from 0.6595 to 0.3351. It also reduces every remaining error in Table 6, and direct comparison with the exact control gives cosine alignment 0.9804 ± 0.0019. The two audits thus connect accurate path quantities to the control identified by the likelihood analysis. This connection determines both the advantages and the limits discussed next.

Table 2: **Exact path and control audits** (mean $\pm$ s.d.; lower is better; 64 paths per method).

Double Well event probabilities

Method	**Exact posterior error**		
	Well occupancy	**Crossing prob.**	**Hitting time W_1**
Brownian Smoothing	0.0061 ± 0.0006	0.0254 ± 0.0080	0.0237 ± 0.0072
FMPE (matched)	0.1344 ± 0.0060	0.1891 ± 0.0149	0.1443 ± 0.0054
SDA	0.2649 ± 0.0091	0.74150 ± 0.00004	0.7505 ± 0.0024
Particle smoother	0.0034 ± 0.0000	0.0126 ± 0.0008	0.0112 ± 0.0006

Wave control intervention with true X_0

Control	**Learned control quality**		
	Posterior mean	**Temporal cov.**	**Energy**
Learned	0.1646 ± 0.0028	0.1754 ± 0.0019	0.3351 ± 0.0447
Zero	1.1375 ± 0.0018	0.3687 ± 0.0004	0.6595 ± 0.0139

6 Discussion and Limitations

We have presented an amortized smoother that generates complete posterior trajectories through one simulator rollout. Across exact Gaussian models, nonlinear systems, and dynamics with high dimension, Brownian Smoothing recovers marginal and temporal posterior structure while using the supplied transition map. It approaches a particle smoother with 4,096 particles on Double Well and obtains the best marginal and joint scores on Lorenz–96 and Kolmogorov flow. It also generates 64 trajectories substantially faster and with less memory than the global flow and score methods. The event study shows that the samples recover path quantities, and the intervention shows that the learned controls use observation information beyond the initial state.

These results are made possible by locating posterior inference in the simulator inputs. Conditioning X_0 and shifting each Gaussian increment keeps the learned correction in the driving dimension, while the supplied dynamics propagate it through the complete state. The construction therefore applies when the diffusion coefficient is rectangular or rank deficient, without inversion of state covariance. It also replaces repeated global transformations, iterative guidance, weighting, and resampling with local corrections applied during a single rollout. Conditional maximum likelihood learns these corrections directly from model simulations. The theory identifies its population objective as a forward KL projection and connects the optimal correction to Doob's h-transform.

The current formulation also points to several directions for extension. On a finite grid, Gaussian mean shifts leave covariance changes and other non-Gaussian structure in the residual, and finite training adds an amortization gap. The theory assumes a correct simulator and conditions represented during training; accordingly, the wave sampler degrades at an unseen longer horizon and outside the training range of observation noise. Richer corrections in the driving space, joint inference of unknown parameters, and explicit treatment of model mismatch could extend the method to broader posterior laws. Together, these extensions could make Brownian Smoothing reusable across simulators, horizons, and observation regimes. Thus, these will all be our future work directions.

REPRODUCIBILITY STATEMENT

The main text and Appendix A specify the systems, parameter distributions, Euler grids, observation schemes, training budgets, test cache sizes, baselines, metrics. All reported learned models use terminal checkpoints without selection by validation metrics. The supplementary material includes full Brownian Smoothing implementation codes and experiment scripts.

ETHICS STATEMENT

This study uses synthetic stochastic systems and contains no human, personal, or sensitive data. In consequential applications, posterior trajectories should be interpreted conditional on the stipulated dynamical and observation models and should be accompanied by independent model validation.

AI USE STATEMENT

Per the ICLR 2027 author guidelines, we disclose that large language models (LLMs) were used in preparing this work. (i) *Related-work support*: We used LLMs to help identify potentially relevant literature; all cited works and bibliographic information were checked by the authors. (ii) *Engineering assistance*: We used LLMs to assist with coding tasks. The resulting code and outputs were reviewed and validated by the authors. (iii) *Mathematical verification*: We used LLMs to construct numerical and symbolic sanity checks of our own derivations, including checks of signs, limiting constants, and finite-grid identities. (iv) *Language polishing*: We used LLMs to improve sentence organization in appendix paragraphs describing additional experiments and for limited sentence organization in the main text. The authors reviewed all such edits. No LLMs were used to fabricate scientific data or reported results. We take full responsibility for the content of this paper.

REFERENCES

Juha Ala-Luhtala, Simo Särkkä, and Robert Piché. Gaussian filtering and variational approximations for bayesian smoothing in continuous-discrete stochastic dynamic systems. *Signal Processing*, 111:124–136, 2015. doi: 10.1016/j.sigpro.2014.12.013. URL https://doi.org/10.1016/j.sigpro.2014.12.013.

Cédric Archambeau, Manfred Opper, Yuan Shen, Dan Cornford, and John Shawe-Taylor. Variational inference for diffusion processes. In John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis (eds.), *Advances in Neural Information Processing Systems 20*, pp. 17–24. Curran Associates, Inc., 2007. URL https://proceedings.neurips.cc/paper/2007/hash/818f4654ed39a1c147d1e51a00ffb4cb-Abstract.html.

Alan Bain and Dan Crisan. *Fundamentals of Stochastic Filtering*. Springer, New York, 2009. doi: 10.1007/978-0-387-76896-0. URL https://doi.org/10.1007/978-0-387-76896-0.

Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth O. Roberts. Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, 12(6):1077–1098, 2006. doi: 10.3150/bj/1165269151. URL https://doi.org/10.3150/bj/1165269151.

Mark Briers, Arnaud Doucet, and Simon Maskell. Smoothing algorithms for state-space models. *Annals of the Institute of Statistical Mathematics*, 62(1):61–89, 2010. doi: 10.1007/s10463-009-0236-2. URL https://doi.org/10.1007/s10463-009-0236-2.

Philip A. Browne. A comparison of the equivalent weights particle filter and the local ensemble transform Kalman filter in application to the barotropic vorticity equation. *Tellus A: Dynamic Meteorology and Oceanography*, 68(1):30466, 2016. doi: 10.3402/tellusa.v68.30466. URL https://doi.org/10.3402/tellusa.v68.30466.

Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer, New York, 2005. doi: 10.1007/0-387-28982-8. URL https://doi.org/10.1007/0-387-28982-8.

C. K. Carter and R. Kohn. On gibbs sampling for state space models. *Biometrika*, 81(3):541–553, 1994. doi: 10.1093/biomet/81.3.541. URL https://doi.org/10.1093/biomet/81.3.541.

Nicolas Chopin and Omiros Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*. Springer, Cham, 2020. doi: 10.1007/978-3-030-47845-2. URL https://doi.org/10.1007/978-3-030-47845-2.

Quentin Clairon and Adeline Samson. Optimal control for estimation in partially observed elliptic and hypoelliptic linear stochastic differential equations. *Statistical Inference for Stochastic Processes*, 23(1):105–127, 2020. doi: 10.1007/s11203-019-09199-9. URL https://doi.org/10.1007/s11203-019-09199-9.

Marc Corstanje, Frank van der Meulen, and Moritz Schauer. Conditioning continuous-time markov processes by guiding. *Stochastics*, 95(6):963–996, 2023. doi: 10.1080/17442508.2022.2150081. URL https://doi.org/10.1080/17442508.2022.2150081.

Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020. doi: 10.1073/pnas.1912789117.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34*, pp. 17695–17709, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/940392f5f32a7ade1cc201767cf83e31-Abstract.html.

Bernard Delyon and Ying Hu. Simulation of conditioned diffusion and application to parameter estimation. *Stochastic Processes and their Applications*, 116(11):1660–1675, 2006. doi: 10.1016/j.spa.2006.04.004. URL https://doi.org/10.1016/j.spa.2006.04.004.

Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000. doi: 10.1023/A:1008935410038. URL https://doi.org/10.1023/A:1008935410038.

Arnaud Doucet, Nando de Freitas, and Neil Gordon (eds.). *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001. doi: 10.1007/978-1-4757-3437-9. URL https://doi.org/10.1007/978-1-4757-3437-9.

Garland B. Durham and A. Ronald Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business and Economic Statistics*, 20(3):297–338, 2002. doi: 10.1198/073500102288618397. URL https://doi.org/10.1198/073500102288618397.

Oskar Eklund, Annika Lang, and Moritz Schauer. Guided smoothing and control for diffusion processes. *Stochastic Processes and their Applications*, 192:104806, 2026. doi: 10.1016/j.spa.2025.104806. URL https://doi.org/10.1016/j.spa.2025.104806.

Simon J. Godsill, Arnaud Doucet, and Mike West. Monte carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, 99(465):156–168, 2004. doi: 10.1198/016214504000000151. URL https://doi.org/10.1198/016214504000000151.

Andrew Golightly and Darren J. Wilkinson. Bayesian inference for nonlinear multivariate diffusion models observed with error. *Computational Statistics and Data Analysis*, 52(3):1674–1693, 2008. doi: 10.1016/j.csda.2007.05.019. URL https://doi.org/10.1016/j.csda.2007.05.019.

N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F Radar and Signal Processing*, 140(2):107–113, 1993. doi: 10.1049/ip-f-2.1993.0015.

Jeremy Heng, Adrian N. Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential monte carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020. doi: 10.1214/19-AOS1914. URL https://doi.org/10.1214/19-AOS1914.

Genshiro Kitagawa. Monte carlo filter and smoother for non-gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996. doi: 10.1080/10618600.1996.10474692. URL https://doi.org/10.1080/10618600.1996.10474692.

Dieterich Lawson, Allan Raventós, Andrew Warrington, and Scott Linderman. Sixo: Smoothing inference with twisted objectives. In *Advances in Neural Information Processing Systems 35*, pp. 38844–38858, 2022. doi: 10.52202/068431-2815. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/fddc79681b2df2734c01444f9bc2a17e-Abstract-Conference.html.

Dieterich Lawson, Michael Y. Li, and Scott Linderman. Nas-x: Neural adaptive smoothing via twisting. In *Advances in Neural Information Processing Systems 36*, pp. 8602–8633, 2023. doi: 10.52202/075280-0376. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b3d005a2cb0e71e698e0b13ac657473-Abstract-Conference.html.

Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3870–3882, 2020. URL https://proceedings.mlr.press/v108/li20i.html.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Chris J. Maddison, Dieterich Lawson, George Tucker, Nicolas Heess, Mohammad Norouzi, Andriy Mnih, Arnaud Doucet, and Yee Whye Teh. Filtering variational objectives. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6573–6583, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/fa84632d742f2729dc32ce8cb5d49733-Abstract.html.

Christian A. Naesseth, Scott W. Linderman, Rajesh Ranganath, and David M. Blei. Variational sequential monte carlo. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 968–977, 2018. URL https://proceedings.mlr.press/v84/naesseth18a.html.

Manfred Opper and Cédric Archambeau. The variational gaussian approximation revisited. *Neural Computation*, 21(3):786–792, 2009. doi: 10.1162/neco.2008.08-07-592. URL https://doi.org/10.1162/neco.2008.08-07-592.

George Papamakarios and Iain Murray. Fast ϵ-free inference of simulation models with Bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*, volume 29, pp. 1028–1036, 2016.

Omiros Papaspiliopoulos, Gareth O. Roberts, and Osnat Stramer. Data augmentation for diffusions. *Journal of Computational and Graphical Statistics*, 22(3):665–688, 2013. doi: 10.1080/10618600.2013.783484. URL https://doi.org/10.1080/10618600.2013.783484.

Stefan T. Radev, Ulf K. Mertens, Andreas Voss, Lynton Ardizzone, and Ullrich Köthe. BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(4):1452–1466, 2022. doi: 10.1109/TNNLS.2020.3042395.

H. E. Rauch, F. Tung, and C. T. Striebel. Maximum likelihood estimates of linear dynamic systems. *AIAA Journal*, 3(8):1445–1450, 1965. doi: 10.2514/3.3166. URL https://doi.org/10.2514/3.3166.

Gareth O. Roberts and Osnat Stramer. On inference for partially observed nonlinear diffusion models using the metropolis-hastings algorithm. *Biometrika*, 88(3):603–621, 2001. doi: 10.1093/biomet/88.3.603. URL https://doi.org/10.1093/biomet/88.3.603.

François Rozet and Gilles Louppe. Score-based data assimilation. In *Advances in Neural Information Processing Systems*, volume 36, pp. 40521–40541, 2023. doi: 10.52202/075280-1763.

Hans-Christian Ruiz and Hilbert J. Kappen. Particle smoothing for hidden diffusion processes: Adaptive path integral smoother. *IEEE Transactions on Signal Processing*, 65(12):3191–3203, 2017. doi: 10.1109/TSP.2017.2686340. URL https://doi.org/10.1109/TSP.2017.2686340.

Tom Ryder, Andrew Golightly, A. Stephen McGough, and Dennis Prangle. Black-box variational inference for stochastic differential equations. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4423–4432, 2018. URL https://proceedings.mlr.press/v80/ryder18a.html.

Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, Cambridge, 2013. doi: 10.1017/CBO9781139344203. URL https://doi.org/10.1017/CBO9781139344203.

Simo Särkkä and Juha Sarmavuori. Gaussian filtering and smoothing for continuous-discrete dynamic systems. *Signal Processing*, 93(2):500–510, 2013. doi: 10.1016/j.sigpro.2012.09.002. URL https://doi.org/10.1016/j.sigpro.2012.09.002.

Moritz Schauer, Frank van der Meulen, and Harry van Zanten. Guided proposals for simulating multi-dimensional diffusion bridges. *Bernoulli*, 23(4A):2917–2950, 2017. doi: 10.3150/16-BEJ833. URL https://doi.org/10.3150/16-BEJ833.

Frank van der Meulen and Moritz Schauer. Bayesian estimation of incompletely observed diffusions. *Stochastics*, 90(5):641–662, 2018. doi: 10.1080/17442508.2017.1381097. URL https://doi.org/10.1080/17442508.2017.1381097.

Prakhar Verma, Vincent Adam, and Arno Solin. Variational gaussian process diffusion processes. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1909–1917, 2024. URL https://proceedings.mlr.press/v238/verma24a.html.

Yu Wang and Arnab Ganguly. Variational smoothing and inference for SDEs from sparse data with dynamic neural flows, 2026. URL https://arxiv.org/abs/2605.05606.

Jonas Wildberger, Maximilian Dax, Simon Buchholz, Stephen Green, Jakob H. Macke, and Bernhard Schölkopf. Flow matching for scalable simulation-based inference. In *Advances in Neural Information Processing Systems*, volume 36, pp. 16837–16864, 2023. doi: 10.52202/075280-0737. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/3663ae53ec078860bb0b9c6606e092a0-Abstract-Conference.html.

Christian Wildner and Heinz Koeppl. Moment-based variational inference for stochastic differential equations. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 1918–1926, 2021. URL https://proceedings.mlr.press/v130/wildner21a.html.

Gefan Yang, Frank van der Meulen, and Stefan Sommer. Neural guided diffusion bridges. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 71210–71230, 2025. URL https://proceedings.mlr.press/v267/yang25af.html.

A EXPERIMENTAL DETAILS

A.1 MODEL ARCHITECTURES

Each encoder follows the geometry of its simulator, allowing sparse observations to propagate on the same ring or periodic field as the dynamics. Brownian Smoothing and matched FMPE use the same geometry-specific encoder family with independent parameters, so their comparison contrasts posterior factorization within a common encoding design. Table 3 reports the learned components and shared configuration; exact widths and parameter counts are given in and immediately below the table.

Table 3: **Encoder, posterior backbone, and optimization configuration.** Ring and field blocks use periodic padding. Models use independent weights.

	Brownian Smoothing	**FMPE (matched)**
OBSERVATION ENCODING		
Double Well	Two-layer reverse-causal Transformer, 8 heads	Same design, independent parameters
Lorenz–96	Sparse ring lift, dilated ring blocks $1, 2, 4, 8$, reverse GRU	Same design, independent parameters
Kolmogorov flow	Sparse periodic-field lift, dilated field blocks $1, 2, 4, 8$, reverse GRU	Same design, independent parameters
POSTERIOR BACKBONE		
Hidden width; time features	128; 32	128; 32 (Double Well and Lorenz–96), 232; 32 (Kolmogorov)
Double Well	Four-transform, 8-bin conditional spline flow for X_0; 3 residual control blocks	Vector field over the complete input vector $(X_0, \varepsilon_{0:K-1})$, with 9 temporal residual blocks
Lorenz–96	Six-layer circular spline flow for X_0; 4 dilated ring control blocks	6 ring blocks followed by 4 temporal blocks
Kolmogorov flow	Eight-layer periodic spline flow for X_0; 4 field control blocks with 32-mode projection	8 field blocks, 4 temporal blocks, and 2 global mode blocks
TRAINING & INFERENCE PROTOCOL		
Simulator paths	10^5 per seed	
Batch size	64 (Double Well); 16 / 2 (Lorenz–96, $D \in \{40, 320\}$); 1 (Kolmogorov, 8 accum.)	
Optimizer	AdamW (learning rate 3×10^{-4}, weight decay 10^{-4}, grad-norm clip 1.0)	
Precision	BF16 networks with FP32 loss reductions	
Returned paths	64 unweighted joint trajectories	
Inference cost	One Euler rollout (1 control eval/step)	Midpoint ODE integration (128 evals)

The Brownian Smoothing checkpoints contain 967,396 trainable parameters for Double Well, 3,413,144 for either Lorenz–96 dimension, and 4,247,873 for Kolmogorov flow. The corresponding matched FMPE counts are 1,349,000, 1,739,905, and 4,537,225. SDA learns an order-two prior trajectory score on overlapping five-state windows and applies Gaussian observation guidance at inference. Its width-128 score network uses 32 time features, four residual MLP blocks for Double Well, and four periodic blocks with dilations (1,2,4,8) for the ring and field systems. It has 535,048 trainable parameters for Double Well and 402,440 for both Lorenz–96 and Kolmogorov flow.

The wave validation uses a two-channel periodic architecture of the same width. The encoder lifts sparse displacement values and masks to the 256-site ring, applies four circular residual blocks with dilations (1,2,4,8), and propagates future information with a reverse GRU. It receives the gaps between observations and the lead from the initial time to the first observation, while the control receives the time remaining until the next observation. A six-layer checkerboard rational-quadratic spline models the standardized initial displacement and velocity. A four-block ring network predicts the 256-dimensional drift correction.

A.2 COMMON TRAINING AND EVALUATION PROTOCOL

Each learned method in Table 1 is trained independently for seeds 0–4 on one NVIDIA RTX 3090 GPU. Training uses online prior-predictive simulation, AdamW with learning rate 3×10^{-4}, weight decay 10^{-4}, gradient-norm threshold 1, BF16 network evaluation, and FP32 losses and reductions. A fixed simulator budget determines the number of updates, and evaluation uses the terminal checkpoint. Double Well consumes 100,032 paths to preserve its batch shape; every other main benchmark consumes exactly 100,000 paths. Table 4 reports the update counts and available wall-clock measurements.

The wave validation follows the same optimizer and precision protocol for seeds 0–2 on one RTX 3090 per run. Each microbatch uses 256 Euler steps and an independent $T \sim \mathrm{Unif}(2, 4)$; two microbatches are accumulated per update. Each terminal checkpoint is trained on exactly 100,000 paths and then used for every reported wave evaluation.

Table 4: **Training accounting and single-run wall clock on one RTX 3090.** Brownian Smoothing, matched FMPE, and SDA use the same path budget and batch construction within each Table 1 setting. The wave range covers its three checkpoints.

					Training wall time		
Setting	**Simulator paths**	**Microbatch**	**Accum.**	**Updates**	**Brownian**	**FMPE**	**SDA**
Double Well	100,032	64	1	1,563	1.20 min	1.04 min	0.82 min
Lorenz–96, $D = 40$	100,000	16	1	6,250	4.92 min	9.66 min	3.88 min
Lorenz–96, $D = 320$	100,000	2	1	50,000	39.38 min	77.28 min	31.02 min
Kolmogorov, 64^2	100,000	1	8	12,500	5.76 h	4.22 h	5.32 h
Wave validation, $D = 512$	100,000	4	2	12,500	42.21–43.19 min	—	—

One architecture and optimization configuration is specified in advance for Brownian Smoothing and matched FMPE in each setting. SDA uses the same fixed training protocol, with a separate calibration of its inference guidance scale. The vector experiments use the default scale 0.03. For rings and fields, we evaluate $\{0.03, 0.1, 0.3, 1, 3, 10\}$ on four and one independently generated pilot records, respectively, using cache seed 900001, seed-0 checkpoints trained on 10,000 paths, and 16 predictor evaluations. Among settings with finite outputs, pilot marginal CRPS selects 3.0 for rings and 0.3 for fields. These scales are then used for all five final seeds, 128 predictor evaluations, and the disjoint test cache with seed 100,000. Final SDA sampling uses predictor steps only.

The relative-time wave architecture and the training range $T \in [2, 4]$ are specified before three-seed training. Table 8 evaluates $T \in \{2, 3, 4\}$ within this range and $T = 5$ as extrapolation. Wall-clock measurements describe computational cost and are separate from the quality metrics.

A.3 BENCHMARK SYSTEMS

We list the linear–Gaussian systems for exact checks, followed by the nonlinear benchmarks in Section 5.2. All Brownian motions and observation noises are mutually independent. Parameters are drawn uniformly per record and supplied as known conditioning variables.

Scalar Ornstein–Uhlenbeck process. The scalar state follows

$$\begin{aligned} dX_t &= -aX_t\,dt + \gamma\,dW_t, \quad X_0 \sim \mathcal{N}(0, 1), \\ Y_m &= X_{t_m} + \tau\xi_m, \qquad\quad \xi_m \sim \mathcal{N}(0, 1), \end{aligned} \tag{14}$$

on $t \in [0, 4]$, with $a \sim \mathrm{Unif}(0.3, 1.2)$, $\gamma \sim \mathrm{Unif}(0.15, 0.6)$, and $\tau \sim \mathrm{Unif}(0.05, 0.3)$.

Integrated Langevin process. With position and velocity $X_t = (Q_t, V_t)$, the rank-one system is

$$\begin{aligned} dQ_t &= V_t\,dt, \qquad\quad Q_0 \sim \mathcal{N}(0, 1), \\ dV_t &= \gamma\,dW_t, \qquad\quad V_0 \sim \mathcal{N}(0, 0.5^2), \\ Y_m &= Q_{t_m} + \tau\xi_m, \quad \xi_m \sim \mathcal{N}(0, 1), \end{aligned} \tag{15}$$

where Q_0, V_0, and W are independent, $\gamma \sim \mathrm{Unif}(0.15, 0.5)$, and $\tau \sim \mathrm{Unif}(0.05, 0.3)$. Thus velocity is unobserved and position receives no direct diffusion; information and randomness reach these coordinates only through the dynamics.

Periodic linear advection–diffusion field. Let u_t be represented on the 8×8 periodic grid of $\mathbb{T}^2 = [0, 2\pi)^2$. The field satisfies

$$\begin{aligned} \mathrm{d}u_t(x) &= \left[-c\left(\partial_{x_1} + \tfrac{1}{2}\partial_{x_2}\right) u_t(x) + \nu \Delta u_t(x) - 0.2 u_t(x)\right] \mathrm{d}t + \gamma \sum_{r=1}^{16} \psi_r(x)\, \mathrm{d}W_t^r, \\ Y_m &= P_{\mathcal{S}} u_{t_m} + \tau \xi_m, \qquad \xi_m \sim \mathcal{N}(0, I_{16}), \end{aligned} \tag{16}$$

with initial state $u_0 \sim \mathcal{N}(0, 0.75^2 I_{64})$, where $\mathcal{S} = \{0, 2, 4, 6\}^2$ is the observed 4×4 subgrid. The real, discretely orthonormal forcing functions $\{\psi_r\}_{r=1}^{16}$ are the cosine–sine pairs associated with

$$\mathcal{K} = \{(0,1), (1,0), (1,-1), (1,1), (0,2), (2,0), (1,-2), (1,2)\}.$$

The parameter ranges are $c \sim \mathrm{Unif}(0.2, 0.8)$, $\nu \sim \mathrm{Unif}(0.02, 0.08)$, $\gamma \sim \mathrm{Unif}(0.08, 0.25)$, and $\tau \sim \mathrm{Unif}(0.05, 0.2)$. Spatial derivatives in equation 16 are evaluated spectrally with a unitary discrete Fourier transform.

Partially observed stochastic damped wave. On a periodic ring of $N = 256$ sites, let L be the second-difference Laplacian and write the state as displacement and velocity $X_t = (q_t, v_t) \in \mathbb{R}^{512}$. The exact wave validation uses

$$\begin{aligned} dq_t &= v_t\, dt, \\ dv_t &= \big(\kappa L q_t - q_t - \gamma v_t\big) dt + \sigma\, dW_t, \\ Y_m &= S_4 q_{t_m} + \tau \xi_m, \qquad \xi_m \sim \mathcal{N}(0, I_{64}), \end{aligned} \tag{17}$$

where $W_t \in \mathbb{R}^{256}$ and S_4 retains every fourth displacement coordinate; velocity is never observed. Parameters are independently drawn as $\kappa \sim \mathrm{Unif}(3, 5)$, $\gamma \sim \mathrm{Unif}(0.4, 0.6)$, $\sigma \sim \mathrm{Unif}(0.5, 0.7)$, and $\tau \sim \mathrm{Unif}(0.12, 0.18)$. The parameter-conditioned stationary initial law is Gaussian. For Fourier frequency j, its displacement and velocity variances are

$$\mathrm{Var}(\widehat{q}_{0,j}) = \frac{\sigma^2}{2\gamma\{1 + 4\kappa \sin^2(\pi j / N)\}}, \qquad \mathrm{Var}(\widehat{v}_{0,j}) = \frac{\sigma^2}{2\gamma}, \qquad \mathrm{Cov}(\widehat{q}_{0,j}, \widehat{v}_{0,j}) = 0.$$

Double Well. The bistable scalar diffusion and its observation model are

$$\begin{aligned} dX_t &= \alpha\big(X_t - X_t^3\big) dt + \gamma\, dW_t, \quad X_0 \sim \mathcal{N}(0, 1.25^2), \\ Y_m &= X_{t_m} + \tau \xi_m, \qquad\qquad\quad \xi_m \sim \mathcal{N}(0, 1), \end{aligned} \tag{18}$$

on $t \in [0, 4]$, with $\alpha \sim \mathrm{Unif}(0.6, 1.4)$, $\gamma \sim \mathrm{Unif}(0.25, 0.65)$, and $\tau \sim \mathrm{Unif}(0.05, 0.3)$.

Lorenz–96. For $i = 1, \ldots, D$, with indices interpreted periodically, the ring dynamics are

$$\begin{aligned} dX_t^{(i)} &= \Big[\big(X_t^{(i+1)} - X_t^{(i-2)}\big) X_t^{(i-1)} - X_t^{(i)} + F\Big]\, dt + \gamma\, dW_t^{(i)}, \\ Y_m &= P_2 X_{t_m} + \tau \xi_m, \qquad \xi_m \sim \mathcal{N}(0, I_{D/2}), \end{aligned} \tag{19}$$

with initial state $X_0 \sim \mathcal{N}(2.3\mathbf{1}_D, 3.7^2 I_D)$, where P_2 selects every second ring coordinate. The nonlinear benchmarks use $D \in \{40, 320\}$, $t \in [0, 1.6]$, and $F \sim \mathrm{Unif}(6, 10)$, $\gamma \sim \mathrm{Unif}(0.05, 0.2)$, and $\tau \sim \mathrm{Unif}(0.5, 1.5)$.

Stochastic Kolmogorov flow. Let ω_t denote vorticity on $\mathbb{T}^2 = [0, 2\pi)^2$ and let $u_t = \nabla^\perp \Delta^{-1} \omega_t$. The field evolves as

$$\begin{aligned} d\omega_t &= [-u_t \cdot \nabla \omega_t + \nu \Delta \omega_t - 0.1 \omega_t + a \cos(4x_2)]\, dt + \gamma \sum_{r=1}^{32} \psi_r\, dW_t^r, \\ Y_m &= P_{\mathcal{G}} \omega_{t_m} + \tau \xi_m, \qquad \xi_m \sim \mathcal{N}(0, I_{64}), \end{aligned} \tag{20}$$

where $\mathcal{G} = \{0, 8, \ldots, 56\}^2$ is an 8×8 observation subgrid of the 64×64 state grid. The ψ_r are 16 sine–cosine pairs associated with the lowest nonzero Fourier wavevectors and are discretely orthonormal. The full-rank initial law is obtained by applying the normalized spectral filter $(4 + |k|^2)^{-1}$ to spatial white noise using a unitary Fourier transform. We use $t \in [0, 1]$, with $\nu \sim \mathrm{Unif}(8 \times 10^{-4}, 1.2 \times 10^{-3})$, $a \sim \mathrm{Unif}(0.8, 1.2)$, $\gamma \sim \mathrm{Unif}(0.03, 0.08)$, and $\tau \sim \mathrm{Unif}(0.05, 0.15)$. The nonlinear term is evaluated pseudo-spectrally with $2/3$ dealiasing.

A.4 Linear–Gaussian validation

Discretization and observations. All paths and exact references use the same Euler map. OU and integrated Langevin use $K = 128$ steps on $[0, 4]$, with $\Delta = 1/32$. OU has 16 event times sampled without replacement from the positive grid times; integrated Langevin has regular events at $t_m = m/4$, $m = 1, \ldots, 16$. The field uses $K = 64$ steps on $[0, 1]$ and events at $t_m = m/16$. Its partial observation is specified by $P_{\mathcal{S}}$. The fixed test cache contains 512 records for each scalar system and 256 field records, generated with seed 100,000.

Training and posterior sampling. The analytic checks follow Appendix A.2 and use a nominal budget of 100,000 prior-predictive records. Batch sizes of 64, 8, and 8 give 100,032 OU records and 100,000 records for each multivariate system. Each run uses one RTX 3090 and returns 64 equally weighted trajectories from its terminal checkpoint.

All three models use context width 128. OU combines a four-transform, eight-bin conditional spline flow for X_0, a two-layer eight-head reverse-time encoder, and a three-block residual control. The multivariate systems use a fixed-window encoder, a conditional full-covariance Gaussian $q(X_0 \mid y, \lambda)$ parameterized by a precision Cholesky factor, and a three-block affine control. Disjoint projections and separately clipped parameter groups balance the initial-state target against the K control targets.

Exact reference and reported errors. For every record, an FP64 Kalman filter and Rauch–Tung–Striebel recursion use the transition matrices, process covariances, observation operator, and observation mask of the Euler simulator. Let m_k, C_k be the exact smoothing moments and $\widehat{m}_k, \widehat{C}_k$ the empirical moments of the 64 Brownian Smoothing trajectories, with $\widehat{C}_k$ using divisor 64. Define $s_{k,d} = (C_{k,dd})^{1/2}$ and $\widehat{s}_{k,d} = (\widehat{C}_{k,dd})^{1/2}$. The reported errors are

$$
\begin{aligned}
E_m^2 &= \frac{1}{(K+1)D} \sum_{k=0}^{K} (\widehat{m}_k - m_k)^\top C_k^{-1} (\widehat{m}_k - m_k), \\
E_C^2 &= \frac{\sum_{k=0}^{K} \|\widehat{C}_k - C_k\|_F^2}{\sum_{k=0}^{K} \|C_k\|_F^2}, \\
E_s^2 &= \frac{1}{(K+1)D} \sum_{k=0}^{K} \sum_{d=1}^{D} \left(\frac{\widehat{s}_{k,d}}{s_{k,d}} - 1 \right)^2 .
\end{aligned}
\tag{21}
$$

Figure 2 plots $E/E_{\mathrm{RTS},64}$ on a logarithmic axis. The denominator uses the same empirical estimators on 64 exact RTS draws. We average eight independent RTS repetitions for the multivariate systems and evaluate the scalar OU floors analytically. The resulting (E_m, E_C, E_s) floors are $(0.1250, 0.1761, 0.0890)$, $(0.1250, 0.1890, 0.0888)$, and $(0.1250, 0.7300, 0.0889)$, respectively. The field covariance floor is larger because 64 centered samples in $D = 64$ produce a covariance of rank at most 63.

Coverage uses empirical weighted order statistics. For 64 equally weighted draws from a calibrated continuous posterior, the nominal central 90% interval has finite-sample expectation $57/65$.

A.5 OU EULER REFINEMENT

The Euler-grid objective is a numerical approximation to the continuous-time construction in Proposition 2. We isolate that approximation in a scalar OU model with

$$dX_t = -0.7X_t\,dt + 0.4\,dW_t, \qquad X_0 \sim \mathcal{N}(0,1),$$

and observation noise standard deviation 0.15. We generate 512 fixed records from the exact transition, with 16 equally spaced noisy observations on $[0,4]$. At each $K \in \{32,64,128,256\}$, the exact OU transition and the Euler transition are separately passed through an FP64 Kalman filter and RTS recursion on the same observations. The exact-transition recursion is the continuous-time posterior restricted to the chosen grid. We report relative RMS errors for the smoothing mean, marginal standard deviation, and the conditional mean shift of the standardized Gaussian increment divided by $\sqrt{\Delta}$.

Figure 4 shows first-order refinement. From $K = 32$ to $K = 256$, the three errors decrease from 0.02697, 0.02631, and 0.02573 to 0.00305, 0.00320, and 0.00315, respectively. This experiment isolates discretization error: the Euler target approaches the continuous OU smoothing law and its driving-noise control as the grid is refined. Learned approximation error is evaluated separately in Figure 2.

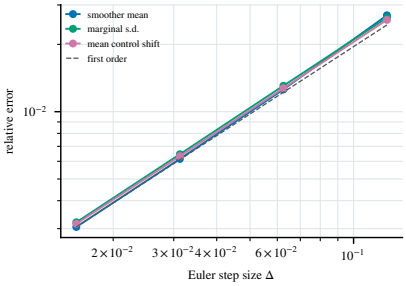


Figure 4: **Step refinement from continuous OU smoothing to the finite Euler target.** Each point compares the Euler RTS smoother with the exact-transition RTS reference on the same 512 observation records. Errors are relative RMS differences in the smoothing mean, marginal standard deviation, and conditional mean shift of the standardized Gaussian increment divided by $\sqrt{\Delta}$. The dashed guide has first-order slope.

A.6 NONLINEAR BENCHMARKS

Simulation and observations. Double Well is discretized by $K = 256$ Euler steps on $[0,4]$, and each record contains 24 event times sampled uniformly without replacement from the positive grid times. Lorenz–96 uses $K = 256$ steps on $[0,1.6]$ and 32 similarly sampled events; at each event, P_2 exposes $D/2$ alternating ring coordinates. Kolmogorov flow uses $K = 128$ steps on $[0,1]$, with observations at $t_m = m/32$, $m = 1,\ldots,32$, on the fixed 8×8 grid $\mathcal{G}$. Each scheduled scalar measurement is independently missing with probability 0.1, subject to retaining at least one component per event. The raw system parameters are fixed within a record, known at inference, and normalized only at the neural-network input.

Training and inference budgets. The learned methods follow Appendix A.2: each receives 10^5 prior-predictive paths for every seed in $\{0,1,2,3,4\}$. Table 4 gives their batch sizes, update counts, hardware, precision, and wall-clock measurements. Brownian Smoothing directly generates $N = 64$ equal-weight joint paths in one rollout. Matched FMPE generates the same number after 128 midpoint evaluations of the complete whitened input vector $(X_0, \varepsilon_{0:K-1})$. SDA also returns 64 paths, using 128 predictor evaluations in Table 1; its 256-predictor reference setting is used only in the joint validation below.

LETKS first forms an $E = 256$-member analysis ensemble, using inflation 1.05 and periodic localization radii 5 and 6 for Lorenz–96 and Kolmogorov, respectively. We then draw $N = 64$ analysis members uniformly with replacement and assign them weights $1/N$. The bootstrap particle smoother propagates $P = 4{,}096$ particles for Double Well and $P = 1{,}024$ for the high-dimensional systems, with systematic resampling below an ESS ratio of 0.5. It returns $N = 64$ equal-weight smoothing trajectories: FFBSi draws for Double Well and final-weight genealogical draws for the high-dimensional systems. The normalized filtering and backward weights determine the categorical path-selection probabilities; they are not attached to the selected trajectories after sampling. Thus E and P specify internal inference work, whereas $N = 64$ is the common number of trajectories passed to every quality metric.

Efficiency profiling. Efficiency uses a separate matched profile on one fixed record with evaluation batch size 1. Runtime and VRAM cover the complete sampling call up to the 64 returned paths, but exclude metric computation. They therefore include construction and analysis of all 256 LETKS

members, or propagation, weighting, resampling, and smoothing of all P particle-filter particles, before the final 64-path draw. Quality aggregates instead evaluate only the common $N = 64$ returned paths on the fixed test caches described next.

Evaluation and aggregation. The fixed test caches contain 512 Double Well records and 256 records for each high-dimensional system, generated once with cache seed 100,000 and reused for all methods and seeds. For returned paths $X^{(j)}_{0:K}$, normalized weights w_j, and the simulated path $X^{\mathrm{true}}_{0:K}$, RMSE is computed from the weighted posterior mean over every state and time coordinate. Marginal CRPS is evaluated at the same coordinates and averaged over records. For one record, the weighted joint energy score is

$$\mathrm{ES}_w = \sum_{j=1}^{N} w_j \left\| \mathrm{vec}\Big(X^{(j)} - X^{\mathrm{true}}\Big) \right\|_2 - \frac{1}{2} \sum_{j=1}^{N} \sum_{\ell=1}^{N} w_j w_\ell \left\| \mathrm{vec}\Big(X^{(j)} - X^{(\ell)}\Big) \right\|_2, \qquad \sum_{j=1}^{N} w_j = 1.$$

Double Well uses this score directly. For the high-dimensional systems, the same weighted functional is applied after the fixed projections and RMS normalization defined below. In the reported pipeline every returned set is equally weighted, $w_j = 1/64$; for the particle smoother, likelihood-derived weights enter the preceding FFBSi or genealogical draw rather than the score weights themselves. Both joint scores retain time as part of the scored object.

For each test record i, the Euler innovation residual and its model scale are

$$R_i^2 = \frac{1}{KNd} \sum_{k=0}^{K-1} \sum_{j=1}^{N} \frac{\left\| X_{k+1}^{(i,j)} - X_k^{(i,j)} - b(t_k, X_k^{(i,j)})\Delta_k \right\|_2^2}{\Delta_k}, \qquad R_{0,i}^2 = \frac{\mathrm{tr}(\sigma_i \sigma_i^\top)}{d}. \tag{22}$$

Table 1 reports the record average of $R_i/R_{0,i}$. Under the additive-noise Euler models used here, a value near one matches the stipulated diffusion scale; unlike the other three quality metrics, this is a dynamics diagnostic with target one rather than a lower-is-better posterior score.

A.7 Temporal coupling test

RMSE and marginal CRPS in Table 1 evaluate one-time marginals, whereas the joint score evaluates complete trajectories. We isolate temporal coupling by independently permuting the 64 particle labels at every Euler time for Lorenz–96 at $D = 40$ and the 64^2 Kolmogorov flow. The permutation preserves each empirical time marginal exactly while removing the learned pairing across times. The comparison uses 64 shared records and test paths for each of checkpoints 0–2.

The primary score is a high-dimensional sliced energy score. Let R collect the same $J = 64$ fixed Rademacher projection directions for every time. To remove the mechanical growth of Euclidean norm with trajectory and projection count, we report the RMS-normalized score

$$\mathrm{SES}(Q, z) = \frac{1}{\sqrt{(K+1)J}} \left[\mathbb{E}_{Z \sim Q} \|RZ - Rz\|_2 - \frac{1}{2} \mathbb{E}_{Z, Z' \sim Q} \|RZ - RZ'\|_2 \right].$$

The fixed positive normalization preserves propriety and paired ordering. Each projected object retains every Euler time, so the score evaluates the projected joint trajectory law. We also report lag-one correlation error and the raw Euler innovation residual; for Kolmogorov flow, temporal field autocorrelation provides a system-specific dependence check. The time-shuffled paths serve as a negative control for information carried by the learned cross-time pairing.

Table 5: **Temporal coupling test.** Mean ± sample standard deviation over three frozen checkpoints and 64 shared records per checkpoint. The time-shuffled sample independently permutes particle identities at every Euler time while preserving each empirical time marginal. Lower is better for the proper score and correlation errors; the Euler residual is a transition-breaking diagnostic. "3/3" means the original joint paths improve the paired metric for every checkpoint.

System	Joint-sensitive metric	Original joint paths	Time shuffled	Original better
Lorenz–96, $D = 40$	Normalized sliced energy score	1.5148 ± 0.0393	1.5272 ± 0.0394	3/3
	Lag-one correlation error	0.00298 ± 0.00004	0.45295 ± 0.01777	3/3
	Euler transition residual	0.13162 ± 0.00001	41.05847 ± 1.47441	3/3
Kolmogorov, 64^2	Normalized sliced energy score	0.34508 ± 0.00722	0.34569 ± 0.00725	3/3
	Temporal autocorrelation error	0.000053 ± 0.000007	0.82602 ± 0.00922	3/3
	Euler transition residual	0.004923 ± 0.000001	7.8329 ± 0.1591	3/3

Across both systems, RMSE, spread, coverage, marginal CRPS, and SBC ranks remain unchanged up to a maximum absolute numerical discrepancy of 2.4×10^{-7}. The original coupling lowers the full-trajectory proper score at all three checkpoints. The larger changes in lag-one correlation and field autocorrelation identify the temporal information destroyed by shuffling, while the Euler residual confirms that the shuffled paths break simulator-consistent transitions. Together, these results show that the returned trajectories contain cross-time structure beyond their unchanged time marginals. Exact joint calibration is evaluated separately in the wave and Double Well experiments below.

A.8 Wave posterior and control

This experiment asks whether sequential increment corrections preserve high-dimensional dependence across time and whether the learned control uses observations after the initial state is known. The exact Gaussian posterior answers the first question, and fixed-X_0 interventions answer the second.

Protocol and exact reference. We evaluate the joint posterior on $[0, 2]$ with $K = 256$ Euler steps and observe 64 displacement coordinates at 16 regular times in $(1, 2]$. The Markov state has dimension 512, each transition has 256 standardized Gaussian increments, and the complete trajectory has $512 \times 257 = 131{,}584$ coordinates. The three terminal checkpoints described in Appendix A.2 use 256 training steps and $T \sim \mathrm{Unif}(2, 4)$. Joint validation uses 64 fixed test records and 64 paths; the intervention and stress evaluations use 32 records per checkpoint.

FMPE scope. The matched FMPE construction in Table 1 would integrate a 66,048-dimensional vector $(X_0, \varepsilon_{0:255})$, with each of 128 vector-field evaluations acting across all 256 Euler steps and forcing coordinates. Profiles on one RTX 3090 place the full protocol of 64 paths, 64 records, and three checkpoints beyond the available compute budget. Figure 3 therefore compares Brownian Smoothing with the exact Gaussian posterior; FMPE comparisons remain confined to Table 1.

The Euler model is linear Gaussian. Because the dynamics are periodic and S_4 observes one site in every four, a Fourier transform over the 64 coarse cells decomposes filtering and smoothing into independent complex blocks of dimension eight: four displacement phases and four velocity phases. We perform the Kalman filter, Rauch–Tung–Striebel recursion, and posterior transition calculations in FP64 within these blocks. Exact trajectories are drawn by Matheron's update: an unconditional path from the Euler-discretized model is corrected by its observation residual using the exact smoothing gains. The same construction with a degenerate initial covariance gives exact samples from $p(X_{1:K} \mid X_0, y)$. At $N = 8$, the block posterior mean, every marginal and cross-time covariance, fixed-initial posterior, posterior samples, and conditional-mean control agree with direct conditioning of the dense joint Gaussian in FP64 unit tests.

For a smoothed mean μ_k, let F_k denote the exact posterior transition coefficient, defined by $\mathbb{E}[X_{k+1} \mid X_k = x, y] = \mu_{k+1} + F_k(x - \mu_k)$. Subtracting the uncontrolled Euler mean and mapping its velocity component through the diffusion coefficient gives the target control $u_k(x, y)$ identified by Theorem 1. This control identifies the posterior conditional mean shift. Exact finite-step sampling also requires the posterior conditional covariance, so the sampling floor uses the complete

Gaussian posterior. On the 64 test records, $I(X_0; Y) = 58.45$ nats and $I(W_{0:T}; Y \mid X_0) = 191.88$ nats. The observations therefore contain substantial information about forcing after conditioning on the initial state.

Joint metrics. Let $V_{k,i}$ be the exact marginal variance. At time t_k, posterior-mean error is the root mean squared error normalized coordinate-wise by $V_{k,i}$, marginal-standard-deviation error is relative RMSE, and spatial covariance is the relative Frobenius error of the exact Fourier blocks. Cross-time covariance uses the corresponding error in $\mathrm{Cov}(X_k, X_K)$. The cumulative energy is

$$\mathcal{E}_k(X) = \int_0^{t_k} \frac{1}{2} \left(\|v_t\|^2 + \|q_t\|^2 + \kappa \sum_i (q_t^{(i+1)} - q_t^{(i)})^2 \right) dt,$$

and running maximum displacement is $\max_{j \le k,i} |q_j^{(i)}|$. Their empirical one-dimensional Wasserstein distances are normalized by the exact posterior standard deviation at t_k. Checkpoint summaries pool spatial covariance over all k, average cross-time covariance over $(0, K/2)$, $(K/2, K)$, and $(0, K)$, and evaluate both path functionals at K. We also compare the empirical means and variances of 16 fixed Gaussian projections of the complete 131,584-dimensional trajectory with their exact values. Two independent sets of 64 exact paths supply the sampling floor in Figure 3.

Across seeds, learned and exact-floor errors are 0.1628 ± 0.0051 and 0.1249 ± 0.0004 for posterior mean, 0.0944 ± 0.0016 and 0.0891 ± 0.0001 for marginal scale, 0.2769 ± 0.0042 and 0.2643 ± 0.0004 for spatial covariance, and 0.4873 ± 0.0102 and 0.4687 ± 0.0007 for cross-time covariance. Projection-mean errors are 0.1513 ± 0.0035 and 0.1237 ± 0.0019; projection-variance errors are 0.1752 ± 0.0043 and 0.1742 ± 0.0048.

Control interventions. Within each record, all interventions share the same initial samples and Brownian draws. Shuffled control uses another record's observations while retaining the target record's physical parameters. Rows with a learned initial law are evaluated against $p(X_{0:K} \mid y)$; rows supplied with the simulated initial state use $p(X_{1:K} \mid X_0, y)$. Numerical comparisons are made within each reference group.

Table 6: **Wave control interventions.** Mean $\pm$ sample standard deviation over 3/3 frozen checkpoints and 32 shared records per seed; every method returns 64 paths. Lower is better. The true-X_0 rows use the exact conditional posterior $p(X_{1:K} \mid X_0, y)$, whereas the other rows use $p(X_{0:K} \mid y)$.

Initial state / control	Mean	Marg. s.d.	Spatial cov.	Cross-time cov.	Energy	Max. disp.
Learned / learned	0.1624 ± 0.0043	0.0938 ± 0.0015	0.2753 ± 0.0040	0.4908 ± 0.0096	0.2942 ± 0.0274	0.2445 ± 0.0057
Learned / zero	0.7927 ± 0.0064	0.9310 ± 0.0114	0.4786 ± 0.0063	0.7143 ± 0.0124	0.4429 ± 0.0594	0.3524 ± 0.0181
Learned / shuffled record	1.9515 ± 0.0213	0.0997 ± 0.0020	0.2761 ± 0.0041	0.4915 ± 0.0098	1.1290 ± 0.1128	0.4786 ± 0.0425
Independent exact posterior	0.1247 ± 0.0003	0.0892 ± 0.0004	0.2636 ± 0.0005	0.4717 ± 0.0009	0.2334 ± 0.0084	0.2383 ± 0.0049
True X_0 / learned	0.1646 ± 0.0028	0.0908 ± 0.0007	0.2636 ± 0.0024	0.1754 ± 0.0019	0.3351 ± 0.0447	0.2749 ± 0.0052
True X_0 / zero	1.1375 ± 0.0018	0.7526 ± 0.0014	0.6191 ± 0.0005	0.3687 ± 0.0004	0.6595 ± 0.0139	0.5844 ± 0.0479

With true X_0, learned control reduces all eight reported errors, including the two projection errors omitted from Table 6. Seed-stratified, record-paired intervals from 50,000 bootstrap replicates exclude zero in every case. At 64 fixed time queries along each held-out trajectory, direct comparison with u_k gives normalized MSE 0.0392 ± 0.0038, cosine alignment 0.9804 ± 0.0019, and learned-to-target control-energy ratio 0.9460 ± 0.0354. The control log-gain per driving-noise dimension and unit time is 0.3656 ± 0.0023, compared with -1.6967 ± 0.1258 under shuffled observations. The learned control therefore uses record-specific observation information beyond the initial law.

A.9 Double Well posterior

Independent grid posterior. The scalar state permits a reference that retains temporal event dependence in addition to smoothing marginals. We discretize $[-6, 6]$ with 1,024 uniformly spaced points and assign transition mass from the Gaussian Euler density at each target node times the grid spacing. Each row is normalized after recording its missing boundary mass. For every test record, an FP64 forward recursion incorporates the observation likelihoods, and a scaled backward recursion supplies the future likelihood at every Euler step. Their product gives the smoothing mass on the

grid. This reference is generated only for final evaluation after the model and inference settings are fixed.

Let $S_k = \mathbf{1}\{X_k \geq 0\}$ and define the first crossing index by

$$\tau = \min\{k \in \{1, \ldots, K\} : S_k \neq S_0\}, \qquad \tau = K + 1 \text{ if the set is empty.}$$

The forward recursion is augmented with two masses that retain paths whose initial sign is negative or positive and that have not yet crossed. At step k, probability transported to the opposite sign is multiplied by the backward message and accumulated as $\Pr(\tau = k \mid y)$; the surviving terminal mass gives $\Pr(\tau = K + 1 \mid y)$. This recursion yields the positive well probability $p_k^+ = \Pr(X_k > 0 \mid y)$, the crossing probability $c = \Pr(\tau \leq K \mid y)$, and the complete hitting time mass function $h_k = \Pr(\tau = k \mid y)$.

Metrics. For each method, the same quantities are estimated from its weighted joint paths. We report the absolute error in p_k^+ averaged over time, absolute error in c, and the scalar Wasserstein distance between predicted and exact hitting time laws,

$$E_{\text{mode}} = \frac{1}{K+1} \sum_{k=0}^{K} |\widehat{p}_k^+ - p_k^+|, \qquad E_{\text{cross}} = |\widehat{c} - c|, \qquad W_1 = \frac{1}{K+1} \sum_{k=1}^{K} |\widehat{H}_k - H_k|,$$

where $H_k = \sum_{j=1}^{k} h_j$. Crossing Brier score, hitting time CRPS, and trajectory energy score assess the same samples against the test path. The particle smoother's ESS is reported as a numerical diagnostic.

Numerical validation. Across all 512 records, the maximum probability normalization error is 1.3×10^{-14}, the maximum prior boundary error is 1.6×10^{-6}, and every observation time aligns exactly with the Euler grid. We recompute 16 evenly spaced record indices on a 2,048-point grid. Relative to that refinement, the 1,024-point reference has mean absolute errors 4.6×10^{-6} in mode probability and 6.3×10^{-5} in crossing probability; its mean hitting time Wasserstein discrepancy is 7.4×10^{-5}. These errors are well below the method discrepancies in Table 7.

Table 7: **Double Well joint posterior evaluation.** Entries are mean ± sample standard deviation over five seeds and 512 fixed records. The first three columns compare directly with the exact grid posterior; the remaining proper scores use the test path. Lower is better. SDA uses the 256-predictor reference setting; Table 1 uses 128 predictor evaluations.

	Error from Exact Posterior			Proper Score		
Method	E_{mode}	E_{cross}	W_1	**Crossing Brier**	**Hitting CRPS**	**Energy**
Brownian Smoothing	0.0061 ± 0.0006	0.0254 ± 0.0080	0.0237 ± 0.0072	0.0506 ± 0.0026	0.0495 ± 0.0020	1.6835 ± 0.0682
FMPE (matched)	0.1344 ± 0.0060	0.1891 ± 0.0149	0.1443 ± 0.0054	0.1128 ± 0.0105	0.0974 ± 0.0057	3.4128 ± 0.2840
SDA	0.2649 ± 0.0091	0.74150 ± 0.00004	0.7505 ± 0.0024	0.7441 ± 0.0001	0.7191 ± 0.0033	7.3436 ± 0.5210
Particle smoother	$\mathbf{0.0034 \pm 0.0000}$	$\mathbf{0.0126 \pm 0.0008}$	$\mathbf{0.0112 \pm 0.0006}$	$\mathbf{0.0471 \pm 0.0012}$	$\mathbf{0.0472 \pm 0.0007}$	$\mathbf{1.6237 \pm 0.0058}$
Exact grid posterior	—	—	—	0.0469	0.0469	—

The particle smoother is closest to the grid reference, consistent with its use of 4,096 particles weighted by the likelihood for each record. Brownian Smoothing has the second smallest discrepancy in every joint quantity using 64 unweighted paths from one forward rollout; its crossing and hitting scores are close to the exact reference, and its trajectory energy score is the second lowest. Matched FMPE remains substantially farther from the exact posterior, while SDA concentrates almost all crossing mass at early times. Seed-stratified paired bootstrap intervals over records exclude zero in favor of Brownian Smoothing for all three exact-posterior discrepancies against matched FMPE and SDA.

The SDA row uses the reference implementation's 256-predictor setting (Rozet & Louppe, 2023) for every record and seed, with shared evaluation batches and sampling seeds. This doubles its inference work relative to Table 1. Posterior-quality comparisons use the reference setting here; matched-budget comparisons use the 128-evaluation results in the main table.

A.10 WAVE ROBUSTNESS

The three wave checkpoints in Table 4 were trained with $K = 256$ and $T \sim \mathrm{Unif}(2, 4)$, then evaluated along three axes. Grid tests set $T = 2$ and $K \in \{128, 256, 512\}$. Horizon tests fix $\Delta = 1/128$, vary $T \in \{2, 3, 4, 5\}$, and place 16 observations in the second half of the interval. Observation tests fix $(T, K) = (2, 256)$ and set $\tau \in \{0.30, 0.15, 0.075\}$, compared with the training range $[0.12, 0.18]$. Each cell uses 32 shared records per checkpoint, its matched Euler posterior, and an independent 64-path sampling floor.

Table 8: **Wave stress test.** Normalized errors, reported as mean $\pm$ sample standard deviation over three checkpoints. Lower is better. "Sampled" denotes $\tau \sim \mathrm{Unif}(0.12, 0.18)$; $T \leq 4$ lies in the training range and $T = 5$ tests extrapolation. Every cell has a matched exact posterior and sampling floor.

Stress axis	T	K	τ	Mean	Marg. s.d.	Cross-time cov.	Energy
Coarser grid	2	128	sampled	0.1635 ± 0.0040	0.0964 ± 0.0018	0.4868 ± 0.0090	0.3359 ± 0.0554
Training grid	2	256	sampled	0.1617 ± 0.0061	0.0947 ± 0.0017	0.4863 ± 0.0101	0.3052 ± 0.0206
Finer grid	2	512	sampled	0.1598 ± 0.0041	0.0929 ± 0.0009	0.4977 ± 0.0109	0.2945 ± 0.0362
In-range horizon	3	384	sampled	0.1605 ± 0.0077	0.0910 ± 0.0005	0.5766 ± 0.0080	0.3103 ± 0.1011
Endpoint horizon	4	512	sampled	0.1740 ± 0.0086	0.0911 ± 0.0004	0.7296 ± 0.0079	0.4653 ± 0.0951
Held-out horizon	5	640	sampled	0.5333 ± 0.0941	0.1067 ± 0.0043	0.9542 ± 0.1033	1.1986 ± 0.6237
Weaker observation	2	256	0.30	0.2427 ± 0.0291	0.1134 ± 0.0124	0.5049 ± 0.0105	0.4900 ± 0.2040
In-range center	2	256	0.15	0.1607 ± 0.0041	0.0939 ± 0.0014	0.4870 ± 0.0100	0.2810 ± 0.0257
Stronger observation	2	256	0.075	0.2561 ± 0.0205	0.1575 ± 0.0093	0.4991 ± 0.0126	0.7091 ± 0.2281

The sampler changes little under the tested twofold grid coarsening and refinement. On observations generated at $K = 512$, the exact $K = 128$ posterior differs from the exact $K = 512$ posterior by 0.0371 in normalized mean and 0.0358 in marginal scale; the corresponding $K = 256$ discrepancies are 0.0122 and 0.0116. The experiment therefore measures stability across these finite grids, while continuous-time convergence remains a separate question.

Posterior-mean and marginal-scale errors remain similar from $T = 2$ through $T = 4$, and cross-time covariance stays close to its exact sampling floor of 0.7219 at $T = 4$. Energy error widens at this endpoint relative to its floor of 0.2159. At the extrapolated horizon $T = 5$, all four errors degrade while the sampling floors remain well behaved. Observation noise outside the training interval also degrades the sampler, especially marginal scale and energy under stronger observations. The model therefore interpolates across its training horizon but does not yet generalize reliably beyond that range.

B RELATED WORK

Smoothing has progressed from closed-form recursions to path samplers and learned conditional models. These methods differ in where they represent posterior corrections and how much work each new record requires during inference and posterior generation for unseen observations.

Gaussian smoothing. The Rauch–Tung–Striebel smoother gave an exact forward–backward solution for linear Gaussian systems (Rauch et al., 1965; Särkkä, 2013). Later Gaussian smoothers extended this idea to nonlinear models by propagating tractable approximations of the filtering and smoothing laws (Archambeau et al., 2007; Särkkä & Sarmavuori, 2013; Ala-Luhtala et al., 2015). These methods are efficient because they summarize uncertainty with means and covariances. The same representation can miss multimodality, skewness, and complex dependence along a trajectory. Gaussian smoothers also run a new forward–backward calculation for every observation record, so they do not amortize inference or efficiently reuse computations across records at test time.

Particle and path-space methods. Particle filters and smoothers replaced a single Gaussian approximation with weighted samples (Carter & Kohn, 1994; Kitagawa, 1996; Doucet et al., 2000; 2001; Godsill et al., 2004; Cappé et al., 2005; Briers et al., 2010; Chopin & Papaspiliopoulos, 2020). This made nonlinear and non-Gaussian posteriors accessible, but long horizons and high dimensions

can require large particle populations. Weight degeneracy and resampling can also reduce ancestral diversity. Path-space MCMC, diffusion bridges, and guided proposals instead update or propose complete conditioned paths (Roberts & Stramer, 2001; Durham & Gallant, 2002; Beskos et al., 2006; Delyon & Hu, 2006; Golightly & Wilkinson, 2008; Papaspiliopoulos et al., 2013; Schauer et al., 2017; van der Meulen & Schauer, 2018; Corstanje et al., 2023). Controlled and twisted particle methods use future observations to improve these proposals (Heng et al., 2020). These methods remain record-specific and do not remove repeated path updates, weighting, resampling, or large sample populations for each new record over long horizons and in high dimensions.

Variational smoothers. Variational methods replace the smoothing posterior with a tractable process and optimize that process for the observed record (Opper & Archambeau, 2009; Maddison et al., 2017; Naesseth et al., 2018; Ryder et al., 2018; Wildner & Koeppl, 2021; Verma et al., 2024). Control-based formulations similarly search for a drift that guides the latent dynamics toward the observations. They connect smoothing to stochastic control and to the conditioned drift induced by Doob's h-transform. Related estimators recover Brownian controls for partially observed elliptic and hypoelliptic diffusions (Clairon & Samson, 2020). These approaches can preserve more dynamical structure than a global Gaussian approximation. They still commonly solve an optimization problem for each record, so they do not provide immediate posterior trajectories or reuse one fitted solution across records without further record-specific optimization at deployment time.

Amortized smoothing. Simulation-based inference moves this repeated work into a training stage. Neural posterior estimators learn conditional distributions from simulated parameter–observation or latent–observation pairs and then reuse the learned model on new records (Papamakarios & Murray, 2016; Cranmer et al., 2020; Radev et al., 2022). Latent neural SDEs amortize a variational initial law and posterior drift while also learning the latent dynamics (Li et al., 2020). Learned particle smoothers amortize proposal construction but retain a population of weighted and resampled particles (Lawson et al., 2022; 2023). Amortization reduces record-specific fitting, yet it does not make trajectory generation cheap when the representation covers the whole path or requires many interacting particles throughout every sampling run for each observation record.

Generative trajectory models. Flows, probability paths, Schrödinger bridges, and score models offer expressive ways to generate complex conditional distributions (De Bortoli et al., 2021; Lipman et al., 2023; Wildberger et al., 2023; Rozet & Louppe, 2023). A global conditional transport can return an unweighted sample in one pass, but its dimension grows with the horizon and state size. Score-based smoothers instead require repeated evaluations and updates of the complete trajectory. Several recent control and bridge methods are close in spirit but solve different problems. APIS fits a control separately for one record and retains importance weights (Ruiz & Kappen, 2017). Neural Guided Diffusion Bridges targets terminally conditioned Euclidean bridges (Yang et al., 2025). Conditional-score and guided-control methods use record-specific objectives or correction mechanisms (Wang & Ganguly, 2026; Eklund et al., 2026). Their conditioning tasks, parameter assumptions, and outputs differ from ours, so we do not compare their published numbers directly. They do not jointly provide offline amortization and equally weighted one-rollout trajectories for general noisy records under the supplied dynamics without additional record-specific correction.

Brownian Smoothing. Brownian Smoothing follows the amortized route while keeping the supplied dynamics inside the sampler. It learns the posterior initial law and sequential mean shifts of the simulator's standardized Gaussian inputs. Model simulations supervise both quantities, so training needs neither posterior samples nor the observation density. At test time, the learned corrections are applied locally as the simulator advances. This avoids a global trajectory transform, repeated full-path guidance, and a weighted particle population. Working in the noise coordinates also supports rectangular or rank-deficient diffusion coefficients without inverting a state covariance. Our matched FMPE baseline applies a global conditional transport to the same input vector, encoder, and prior-predictive training pairs (Wildberger et al., 2023); it isolates the benefit of the sequential factorization. Brownian Smoothing still restricts each conditional increment to an identity-covariance Gaussian mean shift. Theorems 1 and 2 isolate the resulting approximation error and connect the learned correction to Doob's h-transform under refinement of the Euler grid used for sampling.

C PROOFS

C.1 AUXILIARY DISCRETE STATEMENTS

Lemma 1 (Conditional independence of the next increment)**.** *Under the prior-predictive Euler law, for every $k < K$, $\big(X_0, \varepsilon_{0:k-1}, Y_{\le k}\big) \perp\!\!\!\perp \big(\varepsilon_k, Y_{>k}\big) \mid X_k$. Consequently, for every Borel set $A \subseteq \mathbb{R}^r$, versions may be chosen so that*

$$\Pr(\varepsilon_k \in A \mid X_0, \varepsilon_{0:k-1}, Y) = \Pr(\varepsilon_k \in A \mid X_k, Y_{>k}) \quad \textit{a.s.} \tag{23}$$

The same equality holds for conditional expectations whenever ε_k is integrable.

Proposition 4 (Control of path functionals)**.** *For every measurable F on state paths with values in $[0,1]$,*

$$|\mathbb{E}_p F(\Phi(Z)) - \mathbb{E}_q F(\Phi(Z))| \le \sqrt{D_{\mathrm{KL}}(p\|q)/2}. \tag{24}$$

In particular, this applies to crossing and occupation probabilities. If $\bar\tau = \tau/(K+1) \in [0,1]$ is a normalized hitting time, the Wasserstein distance between its two induced laws is also at most $\sqrt{D_{\mathrm{KL}}(p\|q)/2}$. This follows from data processing, Pinsker's inequality, and $W_1 \le \mathrm{TV}$ on $[0,1]$.

C.2 PROOF OF LEMMA 1

Proof. Fix $k < K$, and set $U_k = (X_0, \varepsilon_0, \dots, \varepsilon_{k-1})$. All coordinate and observation spaces are standard Borel, so every conditional distribution used below admits a regular version. The state X_k is a measurable function of U_k. Conditional on $X_k = x$, the future state path is generated from x, the Gaussian increments $\varepsilon_k, \dots, \varepsilon_{K-1}$, and the future observation noises. Those variables are jointly independent of U_k and of the observation noises used no later than t_k. The Markov transition kernels and the conditionally independent observation kernels therefore give

$$\Pr\big(dU_k, dY_{\le k}, d\varepsilon_k, dY_{>k} \mid X_k = x\big) = K_k^-(dU_k, dY_{\le k} \mid x)\, K_k^+(d\varepsilon_k, dY_{>k} \mid x) \tag{25}$$

for suitable probability kernels K_k^- and K_k^+. An observation at t_k belongs to the first kernel: after conditioning on X_k, it depends only on its own observation noise. Thus equation 25 is exactly the asserted conditional independence.

We next derive the conditional-probability identity, including the null-set issue. Let $f : \mathbb{R}^r \to \mathbb{R}$, a, and b be bounded measurable, with a a function of $(U_k, Y_{\le k})$ and b a function of $Y_{>k}$. Equation 25 gives

$$\mathbb{E}[f(\varepsilon_k)a(U_k, Y_{\le k})b(Y_{>k})|X_k] = \mathbb{E}[a(U_k, Y_{\le k}) \mid X_k]\, \mathbb{E}[f(\varepsilon_k)b(Y_{>k})|X_k]\,. \tag{26}$$

Let $s_f(X_k, Y_{>k})$ be a version of $\mathbb{E}[f(\varepsilon_k) \mid X_k, Y_{>k}]$. Applying the defining property of this conditional expectation to the second factor in equation 26, and then conditioning once more on X_k, shows that

$$\mathbb{E}[f(\varepsilon_k) \mid U_k, Y] = s_f(X_k, Y_{>k}) \quad \text{a.s.} \tag{27}$$

Indeed, the equality first holds after multiplication by products $a(U_k, Y_{\le k})b(Y_{>k})$; a monotone-class argument extends it to every bounded $\sigma(U_k, Y)$-measurable test function. Choose a countable convergence-determining class of bounded continuous functions on $\mathbb{R}^r$. Applying equation 27 simultaneously to that class identifies the two regular conditional probability measures outside one common null set, which proves equation 23. Finally, if ε_k is integrable, apply the bounded result to the coordinatewise truncations $f_n(u) = (-n) \vee u_j \wedge n$. Their L^1 convergence and the contraction property of conditional expectation allow $n \to \infty$, proving the corresponding conditional-mean identity. □

C.3 PROOFS OF THE FINITE-GRID RESULTS

Proof. We divide the proof into the exact entropy decomposition, the curvature estimate, and the limiting control identity.

Posterior factorization. First, all conditional second moments of the Gaussian increments exist for almost every record. Indeed, Tonelli's theorem gives

$$\mathbb{E}_Y \mathbb{E}_{p(\,\cdot\,|Y)} \left[\sum_{k=0}^{K-1} \|\varepsilon_k\|^2 \right] = \mathbb{E} \sum_{k=0}^{K-1} \|\varepsilon_k\|^2 = Kr. \tag{28}$$

Fix a record y outside the resulting null set, and use the notation p, q, u_k, v_k from Section 4.1. Let $\pi_k(\cdot \mid x)$ be the conditional distribution of ε_k given $X_k = x$ and $Y_{>k} = y_{>k}$, and let $m_k(x) = \int e\, \pi_k(de \mid x)$. Then $u_k(x) = m_k(x)/\sqrt{\Delta_k}$. The residual in Theorem 1 is

$$R(y) = \sum_{k=0}^{K-1} \mathbb{E}_p D_{\mathrm{KL}}(\pi_k(\cdot \mid X_k) \| \mathcal{N}(m_k(X_k), I_r)) \,. \tag{29}$$

Let $\mathcal{F}_k = \sigma(X_0, \varepsilon_0, \ldots, \varepsilon_{k-1})$, with $\mathcal{F}_0 = \sigma(X_0)$. Lemma 1 states that under p, the conditional distribution of ε_k given $\mathcal{F}_k$ is $\pi_k(\cdot \mid X_k)$. Under q, it is $\gamma_{a_k} = \mathcal{N}(a_k, I_r)$, where $a_k = \sqrt{\Delta_k}\, v_k(X_k)$. This follows directly from equation 5; X_k, and hence a_k, is $\mathcal{F}_k$-measurable. The same observation proves normalization of equation 5: integrate its Gaussian factors successively from ε_{K-1} to ε_0, and then integrate the normalized initial density.

Apply the chain rule for relative entropy first to X_0 and ε_0, then to (X_0, ε_0) and ε_1, and continue through ε_{K-1}. Because all spaces are standard Borel, the preceding regular conditional kernels are valid disintegrations. The chain rule is an identity of nonnegative extended real numbers, so it requires neither absolute continuity nor subtraction of infinite quantities. It yields

$$D_{\mathrm{KL}}(p\|q) = D_{\mathrm{KL}}(p(X_0)\|q(X_0)) + \sum_{k=0}^{K-1} \mathbb{E}_p D_{\mathrm{KL}}(\pi_k(\cdot \mid X_k)\|\gamma_{a_k}) \,. \tag{30}$$

In particular, if the posterior initial marginal is not absolutely continuous with respect to the model initial marginal, the first term and the left-hand side are both $+\infty$. Likewise, a singular conditional increment distribution makes its conditional KL and the corresponding contribution on the left equal to $+\infty$.

Gaussian Pythagorean identity. Let π be any probability measure on $\mathbb{R}^r$ with finite second moment and mean m, and write $\gamma_c = \mathcal{N}(c, I_r)$. We claim that for every $a \in \mathbb{R}^r$,

$$D_{\mathrm{KL}}(\pi\|\gamma_a) = D_{\mathrm{KL}}(\pi\|\gamma_m) + \frac{1}{2}\|a - m\|^2. \tag{31}$$

If $\pi \not\ll \gamma_m$, equivalence of all nondegenerate shifted Gaussian measures implies $\pi \not\ll \gamma_a$; both sides of equation 31 are then $+\infty$. Suppose instead that $\pi \ll \gamma_m$. Since

$$\log \frac{d\gamma_m}{d\gamma_a}(u) = (m - a)^\top u - \frac{1}{2}(\|m\|^2 - \|a\|^2),$$

this logarithm is π-integrable. The Radon–Nikodym chain rule and $\int u\, \pi(du) = m$ give

$$\begin{aligned} D_{\mathrm{KL}}(\pi\|\gamma_a) &= D_{\mathrm{KL}}(\pi\|\gamma_m) + \int \log \frac{d\gamma_m}{d\gamma_a}(u)\, \pi(du) \\ &= D_{\mathrm{KL}}(\pi\|\gamma_m) + (m - a)^\top m - \frac{1}{2}(\|m\|^2 - \|a\|^2) \\ &= D_{\mathrm{KL}}(\pi\|\gamma_m) + \frac{1}{2}\|a - m\|^2, \end{aligned}$$

where the equality remains valid if the first KL is infinite because the added integral is finite. This proves the claim in every case.

Apply equation 31 conditionally with $\pi = \pi_k(\cdot \mid X_k)$, $m = m_k(X_k)$, and $a = a_k$. Since

$$\frac{1}{2}\|a_k - m_k(X_k)\|^2 = \frac{\Delta_k}{2}\|v_k(X_k) - u_k(X_k)\|^2,$$

substitution into equation 30, together with equation 29, proves equation 8, including every $+\infty$ case.

Every term on the right of equation 8 is nonnegative. Gibbs' inequality is an equality precisely when the two initial densities agree almost everywhere. The control-dependent term vanishes precisely when $v_k(X_k) = u_k(X_k)$ almost surely. Thus the initial and quadratic risks have unique minimizers up to the corresponding posterior null sets. When $R(y) < \infty$, these are also the unique minimizers

of the full KL over measurable initial densities and controls; when $R(y) = \infty$, identity equation 8 instead shows that every family member has infinite KL. The quadratic-risk statement remains meaningful in that case and follows directly by conditioning on $S_k = (X_k, Y_{>k})$:

$$\mathbb{E}\left[\frac{\Delta_k}{2}\|v(S_k)\|^2 - \sqrt{\Delta_k}\, v(S_k)^\top \varepsilon_k \middle| S_k\right] = \frac{\Delta_k}{2}\|v(S_k) - u_k(X_k)\|^2 - \frac{1}{2}\|m_k(X_k)\|^2.$$

Conditional Jensen's inequality and equation 28 ensure that the last term is integrable for almost every record.

Conditional maximum likelihood. Assume now the finite entropy and cross-entropy conditions in the theorem. The standard conditional cross-entropy identity gives

$$\mathbb{E}[-\log q(Z \mid Y)] = H + \mathbb{E}_Y D_{\mathrm{KL}}(p(\cdot \mid Y) \| q(\cdot \mid Y)). \tag{32}$$

The assumptions make both expectations finite, so this is an ordinary real-valued decomposition rather than an indeterminate difference. Expanding the Gaussian factors gives pointwise

$$-\log q(Z \mid Y) = \ell(Z, Y) + \frac{1}{2}\sum_{k=0}^{K-1} \|\varepsilon_k\|^2 + \frac{Kr}{2}\log(2\pi).$$

Each increment has its original $\mathcal{N}(0, I_r)$ marginal under the prior-predictive law, even though it is not conditionally standard Gaussian after observing Y. Therefore the expectation of the middle sum is $Kr/2$. Substitution into equation 32 proves equation 33. Specifically,

$$\mathbb{E}[\ell] = H - \frac{Kr}{2}[1 + \log(2\pi)] + \mathbb{E}_Y D_{\mathrm{KL}}(p(\cdot \mid Y) \| q(\cdot \mid Y)), \tag{33}$$

where $H = -\mathbb{E}\log p(Z \mid Y)$.

Future-potential representation. For the remainder of the proof, assume the positive observation densities and regularity stated in the theorem. For every grid point, collect all observations at that point into $G_j(u) = \prod_{m:t_m = t_j} g_m(y_m \mid u)$, with the empty product equal to one. Define the discrete backward information and endpoint potential by

$$H_j(u) = \mathbb{E}\left[\prod_{\ell=j+1}^{K} G_\ell(X_\ell) \middle| X_j = u\right], \qquad \Psi_j = G_j H_j.$$

Here the expectation is under the unconditioned Euler transition law and $H_K \equiv 1$. For fixed x, let $F_k(x, e) = x + b(t_k, x)\Delta_k + \sigma(t_k, x)\sqrt{\Delta_k}\, e$ and $f_{k,x}(e) = \log \Psi_{k+1}(F_k(x, e))$. Before conditioning on the future record, ε_k has standard Gaussian distribution $\gamma_0 = \mathcal{N}(0, I_r)$. Given $X_k = x$ and $\varepsilon_k = e$, the Euler endpoint is $F_k(x, e)$, and integrating all later increments shows that the likelihood of $y_{>k}$ is $\Psi_{k+1}(F_k(x, e))$. Bayes' formula therefore yields

$$\pi_k(de \mid x) = \frac{e^{f_{k,x}(e)}}{\int e^{f_{k,x}(w)}\, \gamma_0(dw)}\, \gamma_0(de). \tag{34}$$

Positivity makes the numerator strictly positive, while existence of the posterior makes the denominator finite and nonzero for the conditioning values under consideration.

We state precisely the analytic meaning of the regularity invoked in the theorem. It suffices that, for prior-predictive almost-everywhere triples (k, y, x), $f = f_{k,x}$ is C^2, the density in equation 34 has a finite second moment, $e^f \in W^{1,1}(\gamma_0)$, and, with $m = \int e\, \pi(de)$, the square root of $d\pi/d\gamma_m$ belongs to $H^1(\gamma_m)$. The first Sobolev condition justifies Gaussian integration by parts by approximation with compactly supported smooth functions; the second is the domain condition for the Gaussian log-Sobolev inequality. These conditions also follow from the usual smooth cutoff argument when e^f, $e^f\|\nabla f\|$, and $e^f\|\nabla f\|^2$ are integrable and the boundary terms of the cutoffs tend to zero.

Since $F_k(x, \cdot)$ is affine and x is held fixed, the ordinary chain rule gives the exact identities

$$\begin{aligned}\nabla_e f(e) &= \sqrt{\Delta_k}\, \sigma(t_k, x)^\top \nabla_x \log \Psi_{k+1}(F_k(x, e)),\\ \nabla_e^2 f(e) &= \Delta_k\, \sigma(t_k, x)^\top \nabla_x^2 \log \Psi_{k+1}(F_k(x, e))\sigma(t_k, x).\end{aligned} \tag{35}$$

The Brownian-direction curvature hypothesis thus implies $\|\nabla_u^2 f(u)\|_{\mathrm{op}} \le L_k$ with $L_k = M\Delta_k$. By the mean-value theorem, ∇f is globally L_k-Lipschitz.

Centered-Gaussian residual. Fix one of the tilted kernels π in equation 34, and let $m = \mathbb{E}_\pi e$. Relative to $\gamma_m = \mathcal{N}(m, I_r)$, its density $r_m = d\pi/d\gamma_m$ is proportional to $\exp(f(e) - m^\top e)$. Gross' Gaussian log-Sobolev inequality states that $\mathrm{Ent}_{\gamma_m}(g^2) \le 2\int \|\nabla g\|^2 d\gamma_m$. Taking $g = \sqrt{r_m}$, using $\int r_m \, d\gamma_m = 1$, and differentiating $\log r_m = f - m^\top e + \mathrm{const}$, we obtain

$$D_{\mathrm{KL}}(\pi\|\gamma_m) \le 2\int \|\nabla\sqrt{r_m}\|^2 d\gamma_m = \frac{1}{2}\mathbb{E}_\pi\|\nabla f(e) - m\|^2. \tag{36}$$

The orientation in equation 36 is essential: the reference measure is the centered-at-the-posterior-mean Gaussian γ_m, exactly the Gaussian appearing in the residual.

It remains to identify the centering in the right-hand side. For the i-th coordinate, Gaussian integration by parts under the standard Gaussian and $e^f \in W^{1,1}(\gamma_0)$ yield

$$\int \partial_i f(e) e^{f(e)} \, \gamma_0(de) = \int \partial_i (e^{f(e)}) \, \gamma_0(de) = \int e_i e^{f(e)} \, \gamma_0(de).$$

After division by the normalizer in equation 34, this becomes

$$\mathbb{E}_\pi \nabla f(e) = \mathbb{E}_\pi e = m. \tag{37}$$

Let e' be an independent copy of e under π. The vector variance identity, the Lipschitz property, and then the same identity for e give

$$\begin{aligned}\mathbb{E}_\pi\|\nabla f(e) - \mathbb{E}_\pi \nabla f(e)\|^2 &= \frac{1}{2}\mathbb{E}\|\nabla f(e) - \nabla f(e')\|^2 \\ &\le \frac{L_k^2}{2}\mathbb{E}\|e - e'\|^2 = L_k^2 \operatorname{tr}\mathrm{Cov}_\pi(e).\end{aligned} \tag{38}$$

Combining equations 36–38, and recalling $L_k = M\Delta_k$, proves the pointwise bound

$$D_{\mathrm{KL}}(\pi_k(\cdot \mid x)\|\mathcal{N}(m_k(x), I_r)) \le \frac{M^2\Delta_k^2}{2} \operatorname{tr}\mathrm{Cov}(\varepsilon_k \mid X_k = x, Y_{>k} = y_{>k}). \tag{39}$$

Average equation 39 first under $p(\cdot \mid y)$ and then over Y. With $S_k = (X_k, Y_{>k})$, the conditional-variance identity gives

$$\begin{aligned}\mathbb{E}\operatorname{tr}\mathrm{Cov}(\varepsilon_k \mid S_k) &= \mathbb{E}\|\varepsilon_k - \mathbb{E}[\varepsilon_k \mid S_k]\|^2 \\ &= \mathbb{E}\|\varepsilon_k\|^2 - \mathbb{E}\|\mathbb{E}[\varepsilon_k \mid S_k]\|^2 \le r.\end{aligned}$$

Summing in k proves the first inequality in equation 9; the second follows from $\sum_k \Delta_k^2 \le |\Delta| \sum_k \Delta_k = T|\Delta|$.

Let $q_{\mathrm{mean}}(\cdot \mid y)$ use the exact posterior initial marginal and the target controls u_k. Equation 8 then gives $D_{\mathrm{KL}}(p(\cdot \mid y)\|q_{\mathrm{mean}}(\cdot \mid y)) = R(y)$. Pinsker's and Jensen's inequalities therefore imply

$$\mathbb{E}_Y\|p(\cdot \mid Y) - q_{\mathrm{mean}}(\cdot \mid Y)\|_{\mathrm{TV}} \le \mathbb{E}_Y\sqrt{R(Y)/2} \le \sqrt{\mathbb{E}_Y R(Y)/2} \le \frac{M}{2}\sqrt{rT|\Delta|}.$$

Limit of the discrete mean control. Equation 37 and the first identity in equation 35 give

$$u_k(x) = \sigma(t_k, x)^\top \mathbb{E}_{\pi_k}\left[\nabla_x \log \Psi_{k+1}(X_{k+1}) | X_k = x\right].$$

Equivalently, applying Gaussian integration by parts before normalization in equation 34,

$$u_k(x) = \sigma(t_k, x)^\top \frac{\int \nabla_x \Psi_{k+1}(F_k(x, e)) \, \gamma_0(de)}{\int \Psi_{k+1}(F_k(x, e)) \, \gamma_0(de)}.$$

This ratio makes the continuous limit precise. Consider grids with $|\Delta| \to 0$ and indices for which $t_k \to t$ and x is fixed. Assume that $\sigma(t_k, x) \to \sigma(t, x)$, that Ψ_{k+1} and its gradient converge locally uniformly near x to $h(t, \cdot)$ and $\nabla h(t, \cdot)$, and that the numerator and denominator integrands above are uniformly integrable under γ_0. Since $F_k(x, e) \to x$ in γ_0-probability, uniform integrability and local uniform convergence give convergence of the two integrals. Strict positivity of $h(t, x)$ then permits division, yielding

$$u_k(x) \longrightarrow \sigma(t, x)^\top \frac{\nabla h(t, x)}{h(t, x)} = u(t, x).$$

At an observation time the same argument applies with the one-sided discrete potentials matching the strictly-future convention. Finally, if an interpolation of the exact Euler posterior converges weakly to the continuous smoothing posterior on a common coupling of records, contraction of total variation under interpolation and the preceding mean-TV estimate show that the exact-mean sampler has the same weak limit in prior-predictive probability. This last implication uses Euler posterior consistency as a separate premise; it does not compare different-grid path laws by KL. □

C.4 Sharpness of the Euler-grid residual rate

Consider $dX_t = dW_t$ with $X_0 \sim \mathcal{N}(0,1)$, and one terminal observation $Y = X_T + \eta$, where $\eta \sim \mathcal{N}(0,\tau^2)$ is independent and $\tau > 0$. On the uniform grid $\Delta = T/K$, conditional on (X_k, Y), write

$$Y - X_k = \sqrt{\Delta}\,\varepsilon_k + \zeta_k, \qquad \zeta_k \sim \mathcal{N}(0, \tau^2 + T - t_{k+1}),$$

with ζ_k independent of ε_k. The scalar Gaussian conditioning formula gives

$$\varepsilon_k \mid X_k, Y \sim \mathcal{N}\left(\frac{\sqrt{\Delta}}{\tau^2 + T - t_k}(Y - X_k), s_k^2\right), \qquad s_k^2 = 1 - \frac{\Delta}{\tau^2 + T - t_k}.$$

The mean-matched Gaussian has the same conditional mean and unit variance. Hence its residual is exactly

$$R_\Delta = \frac{1}{2}\sum_{k=0}^{K-1}\left(s_k^2 - 1 - \log s_k^2\right). \tag{40}$$

Put $a_k = \tau^2 + T - t_k$ and $u_k = \Delta/a_k$. Since $\max_k u_k \le \Delta/\tau^2 \to 0$, the expansion $-u - \log(1-u) = u^2/2 + O(u^3)$ is uniform in k. It follows from equation 40 that

$$R_\Delta = \frac{1}{4}\sum_{k=0}^{K-1}\frac{\Delta^2}{a_k^2} + O\left(\sum_{k=0}^{K-1}\frac{\Delta^3}{a_k^3}\right) = \frac{\Delta}{4}\left\{\Delta\sum_{k=0}^{K-1}\frac{1}{(\tau^2 + T - t_k)^2}\right\} + O(\Delta^2).$$

The expression in braces is a left Riemann sum. Therefore

$$\frac{R_\Delta}{\Delta} \longrightarrow \frac{1}{4}\int_0^T \frac{dt}{(\tau^2 + T - t)^2} = \frac{1}{4}\left(\frac{1}{\tau^2} - \frac{1}{\tau^2 + T}\right) > 0. \tag{41}$$

Thus the $O(|\Delta|)$ residual rate in Proposition 1 is sharp even for a scalar linear–Gaussian smoother.

C.5 Continuous-time setup and assumptions

C.5.1 Posterior construction

Fix a realized record y. Let $\Phi(x_0, w)$ denote the measurable Itô map and let $\mu(dz) = p_0(x_0)\,dx_0\,\mathbb{W}_r(dw)$ be the prior law of $z = (x_0, w)$. For observation densities g_m, define $G(X) = \prod_{m=1}^M g_m(y_m \mid X_{t_m})$ and $C = \mathbb{E}_\mu[G(\Phi(z))]$. The posterior law of the initial state and driving Brownian motion is

$$p(dz \mid y) = \frac{G(\Phi(z))}{C}\,\mu(dz). \tag{42}$$

The backward information function is

$$h(t,x) = \mathbb{E}\left[\prod_{m:\,t_m > t} g_m(y_m \mid X_{t_m})\,\middle|\,X_t = x\right]. \tag{43}$$

Writing $a = \sigma\sigma^\top$ and $\mathcal{A}_t f = b(t,\cdot)^\top\nabla f + \frac{1}{2}\operatorname{tr}(a(t,\cdot)\nabla^2 f)$, the function h satisfies $(\partial_t + \mathcal{A}_t)h = 0$ between observations and $h(t_m^-, x) = g_m(y_m \mid x)h(t_m^+, x)$ at observation times. Its initial density and control are those in equation 10. Under Proposition 2, the posterior state satisfies

$$dX_t = [b(t, X_t) + \sigma(t, X_t)u(t, X_t)]\,dt + \sigma(t, X_t)\,dB_t. \tag{44}$$

If p_X and q_X denote the induced state-path laws, the same proposition and entropy contraction give

$$D_{\mathrm{KL}}(p_X \| q_X) \le D_{\mathrm{KL}}(p(\cdot \mid y) \| q(\cdot \mid y))\,. \tag{45}$$

Assumption 1 (Continuous posterior comparison)**.** *Fix a record y and use C, h, u, v, p, q as in Section 4.2. Assume $0 < C < \infty$, every observation factor is strictly positive, and h is strictly positive and belongs to $C^{1,2}$ on each open interval between consecutive observation times. The prior, target-controlled, and learned-controlled equations are well posed in law; either one-sided value of u may be used at observation times. The initial marginal $p(X_0)$ is absolutely continuous with respect to $q(X_0)$, with finite KL. For $p(X_0)$-almost every x, let Q_x denote the law of equation 11 conditional on $X_0 = x$, and put $d_t = u(t, X_t) - v(t, X_t)$. Assume that $\mathcal{E}(\int_0^\cdot d_t^\top dB_t)_T$ is a uniformly integrable martingale under Q_x, and that $\mathbb{E}_p \int_0^T \|d_t\|^2 dt < \infty$.*

C.6 A POSTERIOR DENSITY-PROCESS LEMMA

Lemma 2 (Posterior density process)**.** *Assume Assumption 1, fix x such that $0 < h(0, x) < \infty$, and let P_x denote the prior law of the driving Brownian motion conditional on $X_0 = x$. With $(\mathcal{F}_t)_{0 \le t \le T}$ denoting the completed canonical filtration, define*

$$H_t = h(t, X_t) \prod_{m:\, t_m \le t} g_m(y_m \mid X_{t_m}), \qquad 0 \le t \le T, \tag{46}$$

where the right-hand value of h is used at every observation time. Then

$$H_t = \mathbb{E}_{P_x}[G(X) \mid \mathcal{F}_t]. \tag{47}$$

In particular, H is a strictly positive uniformly integrable martingale, $H_0 = h(0, x)$, $H_T = G(X)$, and

$$\frac{H_t}{h(0, x)} = \mathcal{E}\left(\int_0^t u(s, X_s)^\top \, dW_s \right). \tag{48}$$

Proof. Fix $t \in [0, T]$. The factors with $t_m \le t$ are $\mathcal{F}_t$-measurable. Conditional independence of the observations given the state path and the Markov property of X imply

$$\begin{aligned} \mathbb{E}_{P_x}[G(X) \mid \mathcal{F}_t] &= \left\{ \prod_{m:\, t_m \le t} g_m(y_m \mid X_{t_m}) \right\} \mathbb{E}_{P_x}\left[\prod_{m:\, t_m > t} g_m(y_m \mid X_{t_m}) \,\middle|\, \mathcal{F}_t \right] \\ &= \left\{ \prod_{m:\, t_m \le t} g_m(y_m \mid X_{t_m}) \right\} h(t, X_t) = H_t. \end{aligned}$$

This proves equation 47. Because $G(X)$ is positive and integrable, its conditional-expectation process is a positive uniformly integrable martingale. The endpoint identities follow from $0 < t_1$ and from the fact that the product defining $h(T, \cdot)$ is empty.

It remains to identify the stochastic logarithm of H. On any open interval containing no observation time, the product in equation 46 is constant. Applying Itô's formula, initially up to a compact-exit localization time, gives

$$\begin{aligned} dh(t, X_t) &= \big(\partial_t h + \mathcal{A}_t h\big)(t, X_t)\, dt + \nabla_x h(t, X_t)^\top \sigma(t, X_t)\, dW_t \\ &= h(t, X_t) u(t, X_t)^\top \, dW_t. \end{aligned}$$

The drift vanishes by the backward Kolmogorov equation, and the last equality follows from strict positivity of h and the definition of u. At an observation time t_m, continuity of X and the jump relation for h give

$$\begin{aligned} H_{t_m -} &= h(t_m^-, X_{t_m}) \prod_{j<m} g_j(y_j \mid X_{t_j}) \\ &= h(t_m^+, X_{t_m}) \prod_{j \le m} g_j(y_j \mid X_{t_j}) = H_{t_m}. \end{aligned}$$

Thus the apparent jumps of h cancel exactly against the newly included likelihood factor. Combining the localized Itô identities over the finitely many observation intervals and then letting the localization levels tend to infinity yields

$$dH_t = H_t \, u(t, X_t)^\top \, dW_t, \qquad H_0 = h(0, x).$$

The unique solution of this scalar stochastic equation is the Doléans–Dade exponential in equation 48. Since its left-hand side is already a uniformly integrable martingale, no separate Novikov condition is needed. □

C.7 Proof of Proposition 2

Proof. We first identify the posterior for a fixed initial state. For every x such that $0 < h(0,x) < \infty$, define Π_x by

$$\frac{d\Pi_x}{dP_x} = \frac{G(X)}{h(0,x)}. \tag{49}$$

The denominator equals $\mathbb{E}_{P_x}[G(X)]$, so equation 49 has unit mass. Lemma 2 identifies its density process as

$$\left.\frac{d\Pi_x}{dP_x}\right|_{\mathcal{F}_t} = \mathcal{E}\left(\int_0^t u(s,X_s)^\top \, dW_s\right).$$

Girsanov's theorem therefore implies that

$$B_t = W_t - \int_0^t u(s,X_s)\, ds$$

is Brownian under Π_x. Substitution into equation 1 yields

$$\begin{aligned} dX_t &= [b(t,X_t) + \sigma(t,X_t)u(t,X_t)]\, dt + \sigma(t,X_t)\, dB_t \\ &= [b(t,X_t) + a(t,X_t)\nabla_x \log h(t,X_t)]\, dt + \sigma(t,X_t)\, dB_t. \end{aligned} \tag{50}$$

We next restore the random initial condition. Let F be any bounded measurable functional of the initial state and driving Brownian path. Disintegrating μ with respect to X_0, using equation 42, and then applying equation 49, we find

$$\begin{aligned} \int F(z)\, p(dz \mid y) &= \frac{1}{C}\int_{\mathbb{R}^d} p_0(x)\, \mathbb{E}_{P_x}[F\, G(X)]\, dx \\ &= \int_{\mathbb{R}^d} \frac{p_0(x)h(0,x)}{C}\mathbb{E}_{\Pi_x}[F]\, dx \\ &= \int_{\mathbb{R}^d} p(x \mid y)\mathbb{E}_{\Pi_x}[F]\, dx. \end{aligned} \tag{51}$$

Thus X_0 has density $p(x \mid y)$, and the conditional distribution of its driving Brownian path given $X_0 = x$ is Π_x. Together with equation 50, this establishes equation 44.

We now prove the relative-entropy identity. Fix an initial state x for which the hypotheses hold. Under Q_x, B is Brownian and X solves equation 11. Recall $d_t = u(t,X_t) - v(t,X_t)$, and define

$$R_x = \mathcal{E}\left(\int_0^\cdot d_t^\top \, dB_t\right)_T. \tag{52}$$

By assumption, R_x has mean one under Q_x. It therefore defines a probability measure $\widetilde{P}_x = R_x Q_x$. Girsanov's theorem shows that

$$\widetilde{W}_t = B_t - \int_0^t d_s\, ds \tag{53}$$

is Brownian under $\widetilde{P}_x$. Since $dB_t = d\widetilde{W}_t + d_t\, dt$, the controlled equation becomes

$$\begin{aligned} dX_t &= [b(t,X_t) + \sigma(t,X_t)\{v(t,X_t) + d_t\}]\, dt + \sigma(t,X_t)\, d\widetilde{W}_t \\ &= [b(t,X_t) + \sigma(t,X_t)u(t,X_t)]\, dt + \sigma(t,X_t)\, d\widetilde{W}_t. \end{aligned} \tag{54}$$

Moreover, the canonical driving path satisfies $W = B + \int_0^\cdot v\, dt = \widetilde{W} + \int_0^\cdot u\, dt$. Thus equation 54 agrees with the fixed-initial posterior equation for both the state and the driving-noise distribution. Well-posedness in law and equation 50 give

$$\widetilde{P}_x = \Pi_x, \qquad \frac{d\Pi_x}{dQ_x} = R_x. \tag{55}$$

Under Π_x, equation 53 reads $dB_t = d\widetilde{W}_t + d_t\,dt$. Taking the logarithm of equation 52 consequently gives

$$\begin{aligned}\log R_x &= \int_0^T d_t^\top\, dB_t - \frac{1}{2}\int_0^T \|d_t\|^2\,dt \\ &= \int_0^T d_t^\top\, d\widetilde{W}_t + \frac{1}{2}\int_0^T \|d_t\|^2\,dt. \end{aligned} \tag{56}$$

The assumed energy bound and equation 51 imply $\mathbb{E}_{\Pi_x}\int_0^T \|d_t\|^2 dt < \infty$ for $p(X_0)$-almost every x. Itô's isometry then makes the stochastic integral in equation 56 a true martingale of mean zero. Hence

$$D_{\mathrm{KL}}(\Pi_x\|Q_x) = \frac{1}{2}\mathbb{E}_{\Pi_x}\int_0^T \|d_t\|^2\,dt. \tag{57}$$

The model and posterior distributions of the initial state and driving noise disintegrate as

$$\begin{aligned} q(dx, dw \mid y) &= q_\phi(x \mid y)\,dx\,Q_x(dw), \\ p(dx, dw \mid y) &= p(x \mid y)\,dx\,\Pi_x(dw). \end{aligned} \tag{58}$$

Define, for $q(\,\cdot \mid y)$-almost every (x, w),

$$R(x, w) = \frac{p(x \mid y)}{q_\phi(x \mid y)} R_x(w). \tag{59}$$

If F is bounded and measurable, then equation 55 and equation 58 give

$$\begin{aligned}\int F R\,dq &= \int_{\mathbb{R}^d} p(x \mid y)\int F(x,w) R_x(w)\,Q_x(dw)\,dx \\ &= \int_{\mathbb{R}^d} p(x \mid y)\int F(x,w)\,\Pi_x(dw)\,dx = \int F\,dp.\end{aligned}$$

Hence $R = dp/dq$. The finite initial KL and energy assumptions justify integrating $\log R$. From equation 59, equation 57, and the posterior disintegration,

$$\begin{aligned} D_{\mathrm{KL}}(p(\,\cdot \mid y)\|q(\,\cdot \mid y)) &= \int_{\mathbb{R}^d} p(x \mid y)\log\frac{p(x \mid y)}{q_\phi(x \mid y)}\,dx + \int_{\mathbb{R}^d} p(x \mid y) D_{\mathrm{KL}}(\Pi_x\|Q_x)\,dx \\ &= D_{\mathrm{KL}}(p(X_0)\|q(X_0)) + \frac{1}{2}\mathbb{E}_p\int_0^T \|u(t, X_t) - v(t, X_t)\|^2\,dt, \end{aligned}$$

which proves equation 12.

It remains to prove entropy contraction. Let $\mathcal{G} = \sigma(\Phi)$, retain $R = dp/dq$, and set $\bar{R} = \mathbb{E}_q[R \mid \mathcal{G}]$. Since the state-path space is standard Borel, the Doob–Dynkin lemma gives a measurable $\bar{r}$ such that $\bar{R} = \bar{r} \circ \Phi$ almost surely. For every bounded measurable state-path functional f,

$$\begin{aligned}\int f\,dp_X &= \mathbb{E}_q[R f(\Phi)] = \mathbb{E}_q[\bar{R} f(\Phi)] \\ &= \int f(x)\bar{r}(x)\,q_X(dx).\end{aligned}$$

Therefore $\bar{r} = dp_X/dq_X$. The function $u \mapsto u\log u$, extended continuously by $0\log 0 = 0$, is convex on $[0, \infty)$. Conditional Jensen's inequality now yields

$$\begin{aligned} D_{\mathrm{KL}}(p_X\|q_X) &= \mathbb{E}_q[\bar{R}\log\bar{R}] \\ &\le \mathbb{E}_q[\mathbb{E}_q[R\log R \mid \mathcal{G}]] \\ &= \mathbb{E}_q[R\log R] = D_{\mathrm{KL}}(p(\,\cdot \mid y)\|q(\,\cdot \mid y))\,. \end{aligned}$$

This is equation 45 and completes the proof. □