DIFFUSION BRIDGE VARIATIONAL INFERENCE FOR DEEP GAUSSIAN PROCESSES

Jian Xu School of Mathematics South China University of Technology 2713091379@qq.com Qibin Zhao John Paisley
RIKEN AIP Department of Electrical Engineering
Columbia University

Delu Zeng

South China University of Technology dlzeng@scut.edu.cn

ABSTRACT

Deep Gaussian processes (DGPs) enable expressive hierarchical Bayesian modeling but pose substantial challenges for posterior inference, especially over inducing variables. Denoising diffusion variational inference (DDVI) addresses this by modeling the posterior as a time-reversed diffusion from a simple Gaussian prior. However, DDVI's fixed unconditional starting distribution remains far from the complex true posterior, resulting in inefficient inference trajectories and slow convergence. In this work, we propose Diffusion Bridge Variational Inference (DBVI), a principled extension of DDVI that initiates the reverse diffusion from a learnable, data-dependent initial distribution. This initialization is parameterized via an amortized neural network and progressively adapted using gradients from the ELBO objective, reducing the posterior gap and improving sample efficiency. To enable scalable amortization, we design the network to operate on the inducing inputs $\mathbf{Z}^{(l)}$, which serve as structured, low-dimensional summaries of the dataset and naturally align with the inducing variables' shape. DBVI retains the mathematical elegance of DDVI—including Girsanov-based ELBOs and reverse-time SDEs—while reinterpreting the prior via a Doob-bridged diffusion process. We derive a tractable training objective under this formulation and implement DBVI for scalable inference in large-scale DGPs. Across regression, classification, and image reconstruction tasks, DBVI consistently outperforms DDVI and other variational baselines in predictive accuracy, convergence speed, and posterior quality.

1 Introduction

Deep Gaussian processes (DGPs) Damianou & Lawrence (2013) extend the representational capacity of Gaussian processes (GPs) by composing multiple layers of latent functions, enabling flexible hierarchical Bayesian modeling. However, posterior inference in DGPs is notoriously challenging due to the non-conjugate likelihoods, strong inter-layer dependencies, and the large number of inducing variables required for scalability. Stochastic variational inference (SVI) with inducing points Hensman et al. (2013); Salimbeni & Deisenroth (2017) has become the standard approach, but designing accurate and efficient variational posteriors remains a central bottleneck.

Recently, denoising diffusion variational inference (DDVI) Xu et al. (2024) has been proposed to approximate the inducing-point posterior via the time-reversal of a diffusion stochastic differential equation (SDE) starting from a simple Gaussian prior. By parameterizing the reverse drift with a neural network, DDVI can flexibly capture complex posteriors while retaining scalable SVI training. However, a key limitation of DDVI lies in its reliance on a fixed, unconditional Gaussian distribution as the start of the reverse diffusion. Since the true posterior over inducing variables is typically far from this initial distribution, the reverse-time SDE must traverse a long and complex path to reach the target, resulting in inefficient inference and slow convergence.




Figure 1: **Comparison between DDVI and DBVI.** (Left) DDVI starts from an unconditional Gaussian prior and runs a reverse diffusion SDE towards the posterior. (Right) DBVI starts from an input-conditioned initial distribution and uses an observation-conditioned diffusion bridge SDE, leading to shorter and more efficient inference trajectories.

To address this, in this work, we propose **Diffusion Bridge Variational Inference (DBVI)**, which replaces the unconditional reverse diffusion in DDVI with a *learnable, input-conditioned distribution* that adapts over training. By parameterizing the start point of the diffusion using a neural network, gradients from the ELBO naturally push the initial distribution closer to the posterior. This effectively narrows the inference gap and alleviates the burden on the reverse SDE, yielding more stable and sample-efficient inference. Moreover, conditioning the initial distribution on observations introduces a natural connection to amortized variational inference: the generative structure of DBVI allows each posterior sample to be generated by a single forward pass of a drift network, without requiring per-dataset optimization. Unlike naive amortization that conditions directly on raw inputs—leading to high-dimensional mismatches and overfitting—our design uses the inducing inputs $\mathbf{Z}^{(l)}$ at each layer as input proxies to the amortizer. This enables batch-wise inference while preserving global dataset structure and matching the dimensionality of inducing variables.

Importantly, the theoretical elegance of DDVI—particularly its use of time-reversed SDEs, Girsanov's theorem, and ELBO construction—remains fully preserved in DBVI. We build upon the same diffusion framework, but reinterpret the prior as a *Doob-bridged process* whose start is parameterized by an amortized network. This allows DBVI to inherit the key benefits of DDVI while significantly improving its flexibility and inference efficiency. We derive an evidence lower bound (ELBO) objective for training DBVI within the SVI framework and provide a scalable implementation for large-scale DGPs. Empirical results on regression, classification, and image reconstruction benchmarks demonstrate that DBVI achieves more accurate posterior approximations, faster convergence, and improved predictive performance compared to DDVI and other state-of-the-art DGP inference methods. Our contributions are as follows:

- We propose Diffusion Bridge Variational Inference (DBVI), a novel extension of DDVI that
 replaces the unconditional start of the reverse diffusion with a learnable, input-conditioned
 initial distribution, effectively reducing the inference gap and improving posterior approximation.
- We introduce a bridge-based reinterpretation of the DDVI framework by integrating Doob's h-transform into the variational formulation, while preserving the core machinery of reverse-time SDEs, Girsanov-based ELBO, and SVI scalability.
- We develop a structured amortization strategy that leverages inducing locations $\mathbf{Z}^{(l)}$ as the input to the drift network, enabling batch-wise amortized inference without requiring access to the full dataset or high-dimensional conditioning on raw inputs.
- We provide a scalable implementation of DBVI for deep Gaussian processes and validate its performance on regression, classification, and image reconstruction tasks, showing consistent improvements over DDVI and other state-of-the-art DGP inference methods in terms of accuracy, convergence speed, and sample efficiency.

2 RELATED WORK

Deep Gaussian Process Inference. Early scalable inference for DGPs built on sparse variational Gaussian processes Hensman et al. (2013), extending inducing point methods to multi-layer latent functions Damianou & Lawrence (2013); Salimbeni & Deisenroth (2017). Despite progress with doubly stochastic variational inference (DSVI) Salimbeni & Deisenroth (2017), accurately capturing complex posteriors over inducing variables remains a challenge, especially in high-depth DGPs. Implicit Posterior Variational Inference (IPVI) Yu et al. (2019) instead parameterizes the inducing-point posterior with a neural network and learns it through an adversarial objective, analogous to GAN training Goodfellow et al. (2014). While more expressive than mean-field assumptions, this adversarial formulation is notoriously difficult to optimize and often suffers from instability, which can result in biased posterior estimates.

Diffusion-based Variational Inference. Score-based generative modeling and diffusion probabilistic models Song et al. (2020); Ho et al. (2020) have inspired new VI methods Xu et al. (2024) that represent posteriors as solutions to reverse-time SDEs. DDVI Xu et al. (2024) applies this to DGP inference, parameterizing the reverse drift via a score network. However, its unconditional start distribution can be far from the posterior, requiring long diffusion paths and increasing variance. Our DBVI can be viewed as a strict extension of DDVI: same theoretical framework, but with bridge correction and amortized initialization, yielding both theoretical guarantees and empirical improvements.

Diffusion Bridge Models. Diffusion bridges Zhou et al. (2023); Li et al. (2023) constrain dynamics between fixed endpoints or distributions, enabling more direct and sample-efficient transitions. Observation-conditioned bridges have been explored in Schrödinger bridge formulations Shi et al. (2023) and consistency diffusion models He et al. (2024). Our DBVI adapts this idea to variational inference in DGPs, integrating an *amortized* parameterization Kim et al. (2018); Agrawal & Domke (2021); Margossian & Blei (2023); Ganguly et al. (2023) to map inputs to initial states, reducing the KL gap and improving efficiency.

3 Method

3.1 DEEP GAUSSIAN PROCESSES

Deep Gaussian Processes (DGPs) Damianou & Lawrence (2013) generalize standard Gaussian Processes (GPs) by hierarchically composing multiple GP layers, enabling deep non-linear probabilistic mappings. Let $\mathbf{x} \in \mathbb{R}^d$ be an input and $\mathbf{y} \in \mathbb{R}^p$ the corresponding output. A DGP with L layers defines latent variables $\{\mathbf{f}^{(l)}\}_{l=1}^L$ recursively through:

$$p(\mathbf{f}^{(l)} \mid \mathbf{f}^{(l-1)}) = \mathcal{GP}\left(0, k^{(l)}(\mathbf{f}^{(l-1)}, \mathbf{f}^{(l-1)})\right),\tag{1}$$

where $k^{(l)}$ is the kernel function at layer l with hyperparameters $\gamma^{(l)}$. For scalability, each layer introduces M_l inducing variables $\mathbf{u}^{(l)}$ located at inducing inputs $\mathbf{Z}^{(l)}$, with GP prior:

$$p_{\text{prior}}(\mathbf{u}^{(l)}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{Z}\mathbf{Z}}^{(l)}).$$
 (2)

Assuming conditional independence across layers given inducing variables, the full joint distribution over outputs \mathbf{y} , latent variables $\{\mathbf{f}^{(l)}\}$, and inducing variables $\{\mathbf{u}^{(l)}\}$ factorizes as:

$$p(\mathbf{y}, \mathbf{F}, \mathbf{U}) = \left[\prod_{l=1}^{L} p(\mathbf{f}^{(l)} \mid \mathbf{f}^{(l-1)}, \mathbf{u}^{(l)}) p(\mathbf{u}^{(l)}) \right] p(\mathbf{y} \mid \mathbf{f}^{(L)}),$$
(3)

where the inter-layer conditional distribution $p(\mathbf{f}^{(l)} \mid \mathbf{f}^{(l-1)}, \mathbf{u}^{(l)})$ is given by the sparse GP conditional:

$$p(\mathbf{f}^{(l)} \mid \mathbf{f}^{(l-1)}, \mathbf{u}^{(l)}) = \mathcal{N}\left(\mathbf{K}_{\mathbf{f}^{(l-1)}\mathbf{Z}^{(l)}}\mathbf{K}_{\mathbf{Z}^{(l)}\mathbf{Z}^{(l)}}^{-1}\mathbf{u}^{(l)}, \ \mathbf{K}_{\mathbf{f}^{(l-1)}\mathbf{f}^{(l-1)}} - \mathbf{K}_{\mathbf{f}^{(l-1)}\mathbf{Z}^{(l)}}\mathbf{K}_{\mathbf{Z}^{(l)}\mathbf{Z}^{(l)}}^{-1}\mathbf{K}_{\mathbf{Z}^{(l)}\mathbf{f}^{(l-1)}}\right). \tag{4}$$

Here, we denote $\mathbf{K}_{ab}^{(l)}$ as the kernel matrix at layer l evaluated between sets \mathbf{a} and \mathbf{b} . The input to the first layer is defined as $\mathbf{f}^{(0)} := \mathbf{x}$.

3.2 Denoising Diffusion Variational Inference (DDVI)

Denoising Diffusion Variational Inference (DDVI) Xu et al. (2024) is a recently proposed method for performing posterior inference over inducing variables in Deep Gaussian Processes (DGPs). It draws inspiration from the success of score-based generative modeling and diffusion probabilistic models Ho et al. (2020); Song et al. (2020), and adapts these ideas to variational inference by modeling the variational posterior as the solution to a reverse-time stochastic differential equation (SDE).

Motivation. Traditional variational inference in DGPs typically relies on simple, factorized Gaussian approximations over the inducing variables $\mathbf{U} = \{\mathbf{u}^{(l)}\}_{l=1}^L$, which are often too restrictive to capture the complex, multimodal posterior arising from deep hierarchical GPs. DDVI seeks to construct more expressive variational distributions by modeling them as the terminal distribution of a *reverse-time diffusion process*, effectively defining a flexible transformation from a known fix initial distribution to the posterior.

Variational posterior via reverse diffusion. DDVI defines the variational distribution as the marginal at time t = 1 of a reverse-time SDE $Q_{\phi}(\mathbf{U}_t)$:

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) - g(t)^2 \nabla_{\mathbf{U}_t} \log p_t(\mathbf{U}_t) \right] dt + g(t) d\mathbf{W}_t, \quad t \in [0, 1],$$
(5)

where $\mathbf{U}_0 \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ is the fixed initial distribution, f and g are drift and diffusion coefficients, and $p_t(\mathbf{U}_t)$ is the marginal law at time t under the reverse process. The reverse SDE is not simulated directly; instead, the score function $\nabla_{\mathbf{U}_t} \log p_t(\mathbf{U}_t)$ in the drift is parameterized by a neural network $s_\phi(\mathbf{U}_t,t)$ to approximate the optimal reverse dynamics.

Diffusion-based ELBO. Rather than relying on score estimation (e.g., denoising score matching), DDVI employs a variational diffusion framework to derive a tractable evidence lower bound (ELBO). Specifically, the reverse-time process $\mathbf{U}_{t\in[0,1]}$ is treated as a variational diffusion process, and the ELBO is expressed using the Girsanov formula for likelihood ratios between stochastic processes:

$$\mathcal{L}_{\text{DDVI}} = \mathbb{E}_{\mathbf{U}_{0:1} \sim Q_{\phi}} \left[-\frac{\|\mathbf{U}_{1}\|_{2}^{2}}{2\sigma^{2}} + \log p\left(\mathbf{y} \mid \mathbf{f}^{(L)}\right) \right]$$

$$-\frac{1}{2} \int_{0}^{1} g(t)^{2} \left\| \frac{\mathbf{U}_{t}}{\kappa_{t}} + s_{\phi}(t, \mathbf{U}_{t}) \right\|_{2}^{2} dt$$

$$- \text{KL}\left(\mathcal{N}(0, \sigma^{2}\mathbf{I}) \mid \mathcal{N}(0, \kappa_{1}\mathbf{I})\right) + \log p_{\text{prior}}(\mathbf{U}_{1}),$$
(6)

where Q_{ϕ} denotes the pathwise density of the variational reverse-time SDE, \mathbf{U}_1 is the terminal state of the reverse diffusion SDE using $s_{\phi}(\cdot,t)$, $\mathbf{f}^{(L)}$ denotes the forward inference of the DGP at \mathbf{U}_1 .

Training and implementation. In practice, DDVI jointly trains the variational drift network s_{ϕ} and DGP hyperparameters γ by maximizing $\mathcal{L}_{\text{DDVI}}$ using stochastic gradient descent. Sampling from $Q_{\phi}(\mathbf{U}_t)$ is achieved by solving the reverse SDE from $\mathbf{U}_0 \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, which allows for reparameterized gradients through the sampled trajectories.

Limitations. While DDVI provides a flexible and theoretically grounded approach to inference in deep GPs, it still faces important limitations:

• Unconditional initialization: DDVI relies on a fixed Gaussian start distribution $U_0 \sim \mathcal{N}(0, \sigma^2 I)$ for the reverse diffusion. Since the true posterior over inducing variables is

typically far from this unconditional prior, the reverse-time SDE must traverse a long and complex path to reach the target distribution, leading to inefficient inference, high variance, and slow convergence.

• Lack of amortization: DDVI does not condition its initialization on the observed data, making posterior sampling input-agnostic rather than scalable amortized inference.

These limitations motivate our proposed method, Diffusion Bridge Variational Inference (DBVI), which introduces observation-conditioned diffusion bridges and amortized initializations for and more accurate posterior inference.

DBVI: OBSERVATION-CONDITIONED DIFFUSION BRIDGE

We begin by introducing the data-dependent initialization of DBVI. Instead of starting from a fixed Gaussian prior as in DDVI, we amortize the mean of the initial distribution using a neural network:

$$p_0^{\theta}(\mathbf{U}_0 \mid \mathbf{x}) = \mathcal{N}(\mathbf{U}_0; \, \mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I}), \tag{7}$$

where only the mean $\mu_{\theta}(\mathbf{x})$ depends on the data, while the variance σ^2 is kept fixed. This initialization provides a closer match to the posterior and shortens the diffusion trajectory. To formalize the resulting dynamics, we now turn to a diffusion bridge representation.

Proposition 1 (Forward & reverse SDE under Doob's h-transform). Let the initial constraint be encoded by the Doob h-transform with

$$h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t), \tag{8}$$

Then the forward bridge has drift

$$\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) = f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0), \tag{9}$$

with the same diffusion coefficient q(t). Moreover, the reverse-time bridge SDE is

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) - g(t)^2 s_{cond}(\mathbf{U}_t, t, \mathbf{U}_0) \right] dt + g(t) d\mathbf{W}_t,$$
 (10)

Equivalently, the conditional score equals $s_{cond}(\mathbf{U}_t, t, \mathbf{U}_0) = s(\mathbf{U}_t, t, \mathbf{U}_0) - h(\mathbf{U}_t, t, \mathbf{U}_0)$.

Proposition 1 states that, by introducing Doob's h-transform, we can reinterpret the dynamics as a bridge process, which essentially bends the diffusion toward the posterior endpoint. The forward SDE incorporates an additional drift term that nudges the path toward the target, while the reversetime bridge SDE involves a conditional score function s_{cond} . This result provides the mathematical foundation for DBVI, showing how conditioning on the initialization modifies both forward and reverse dynamics.

Building on this, we next leverage the bridge process trick introduced in DDVI Xu et al. (2024) to characterize the marginal distribution of the bridge process under a linear drift. This formulation yields a tractable Gaussian form for the bridge marginal, which will be crucial for deriving the DBVI training objective.

Proposition 2 (Marginal of Doob-augmented bridge process). Consider the linear forward SDE with Doob bridge correction

$$d\mathbf{U}_t = \left[-\lambda(t) \,\mathbf{U}_t + g(t)^2 \,h(\mathbf{U}_t, t, \mathbf{U}_0) \right] dt + g(t) \,d\mathbf{B}_t, \tag{11}$$

where $h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t)$ is the Doob h-transform. Assume isotropic initialization

$$p_0^{\theta}(\mathbf{U}_0 \mid \mathbf{x}) = \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I}).$$
 (12)

Then for each $t \in [0, 1]$, the marginal law remains Gaussian,

$$p_t(\mathbf{U}^{Bri} \mid \mathbf{x}) = \mathcal{N}(\mathbf{U}^{Bri}; \mathbf{m}_t, \kappa_t \mathbf{I}), \tag{13}$$

where the mean \mathbf{m}_t and variance κ_t satisfy the coupled ODE system

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathbf{m}_{t} = -(\lambda(t) + c(t)) \mathbf{m}_{t} + c(t) a(t) \mu_{\theta}(\mathbf{x}), \quad \mathbf{m}_{0} = \mu_{\theta}(\mathbf{x}),
\frac{\mathrm{d}}{\mathrm{d}t} \kappa_{t} = -2(\lambda(t) + c(t)) \kappa_{t} + g(t)^{2} + 2 c(t) a(t) \sigma^{2}, \quad \kappa_{0} = \sigma^{2},$$
(15)

$$\frac{\mathrm{d}}{\mathrm{d}t} \kappa_t = -2(\lambda(t) + c(t)) \kappa_t + g(t)^2 + 2 c(t) a(t) \sigma^2, \quad \kappa_0 = \sigma^2,$$
(15)

with

$$a(t) = e^{-\Lambda(t)}, \quad \Lambda(t) = \int_0^t \lambda(s) \, \mathrm{d}s,$$

and correction coefficient

$$c(t) = g(t)^2 \, \frac{\sigma^2 \, a(t)^2}{\left(a(t)^2 \sigma^2 + q(t)\right) \, q(t)}, \qquad q(t) = a(t)^2 \int_0^t \frac{g(r)^2}{a(r)^2} \, \mathrm{d}r.$$

In the special case $c(t) \equiv 0$, this reduces to the bridge process trick in DDVI.

Proposition 2 provides a closed-form expression for the bridge marginal, with mean \mathbf{m}_t and variance κ_t determined by the amortized initialization $\mu_{\theta}(\mathbf{x})$. This Gaussian form plays a crucial role in deriving the training objective, as it enables an explicit comparison between the amortized start distribution and the bridge marginal. In particular, it provides the analytical structure needed to express the KL divergence in a tractable score—matching form, which is then incorporated into the DBVI loss. Additional derivations and detailed proofs are provided in the Appendix.

Finally, we combine the above results to obtain the DBVI training loss. In particular, we express the KL divergence in a score–matching form, yielding a tractable ELBO that can be optimized efficiently.

Proposition 3 (DBVI loss with amortized mean). Let $p_0^{\theta}(\mathbf{u} \mid \mathbf{x}) = \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I})$ be the data-dependent start, and let (\mathbf{m}_t, κ_t) be the mean/variance of the reference bridge marginal at time $t \in [0,1]$ (given by Proposition 2 via ODEs in the Doob-augmented case). Then the pathwise KL between the variational reverse bridge Q_{ϕ} and the reference bridge admits a score-matching representation. Consequently, a tractable per-mini-batch ELBO is

$$\ell_{\text{DBVI}}(\theta, \phi, \gamma) = \mathbb{E}_{\mathbf{U}_{0:1} \sim Q_{\phi}} \left[-\log p_{0}^{\theta}(\mathbf{U}_{1}) + \frac{B}{N} \log p(\mathbf{y}_{I} \mid \mathbf{f}^{(L)}) - \frac{1}{2} \int_{0}^{1} g(t)^{2} \left\| \frac{\mathbf{U}_{t} - \mathbf{m}_{t}}{\kappa_{t}} + s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) \right\|_{2}^{2} dt - \text{KL} \left(\mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^{2}\mathbf{I}) \mid | \mathcal{N}(\mathbf{m}_{1}, \kappa_{1}\mathbf{I}) \right) + \log p_{\text{prior}}(\mathbf{U}_{1}) \right],$$
(16)

where B is the batch size, N is the dataset size, $s_{\rm cond} = s_{\phi} - h$ is the conditional score used by the reverse bridge SDE, \mathbf{U}_1 is its terminal state, $\mathbf{f}^{(L)}$ denotes the DGP forward mapping at \mathbf{U}_1 , and γ collects DGP hyperparameters. When $\mu_{\theta}(\mathbf{x}) = \mathbf{0}$ (and thus $\mathbf{m}_t \equiv \mathbf{0}$), the objective recovers the original DDVI loss.

Proposition 3 shows that DBVI departs from DDVI in two essential ways: (i) the initialization is amortized via $\mu_{\theta}(\mathbf{x})$, which induces a time-dependent reference mean \mathbf{m}_t in the bridge marginal, and (ii) the loss involves the conditional score $s_{\text{cond}} = s_{\phi} - h$, explicitly accounting for the bridge correction. When the amortized mean collapses to zero (so that $\mathbf{m}_t \equiv \mathbf{0}$), the objective reduces exactly to the original DDVI loss, recovering DDVI as a special case.

We summarize the full training procedure of DBVI for deep Gaussian processes in Algorithm 1.

3.4 Structure of the Amortized Network μ_{θ}

The amortized network μ_{θ} plays a crucial role in DBVI by producing the parameters of the initial distribution $p_0^{\theta}(\mathbf{U}_0|\mathbf{x})$. In principle, μ_{θ} should take the observed data \mathbf{x} as input and output the corresponding parameters for the inducing variables \mathbf{U} . However, a straightforward amortization design based on the full dataset is computationally infeasible, since storing and processing the entire dataset within a single neural network is prohibitive in both memory and complexity. On the other hand, if one resorts to mini-batch amortization, the network can only observe partial data at each update, which yields biased estimates of the variational parameters.

A second challenge is the mismatch of dimensions. At each layer of a DGP, the input \mathbf{x} typically has dimension $[B,d_{\text{in}}]$ where B is the batch size, while the inducing variables $\mathbf{u}^{(l)}$ live in a space of shape $[M_l,d_{\text{out}}]$ with M_l inducing points. that maps directly from \mathbf{x} to $\mathbf{u}^{(l)}$ is nontrivial: flattening

Algorithm 1 Diffusion Bridge Variational Inference (DBVI) for DGPs

Input: training data X, y; mini-batch size B; forward drift f; diffusion scale g**Params:** DGP hyperparameters γ ; bridge score network s_{ϕ} ; amortizer $\mu_{\theta}(\cdot)$ with fixed variance

Precompute reference bridge marginals (\mathbf{m}_t, κ_t) for $t \in [0, 1]$: numerically integrate the ODEs for (\mathbf{m}_t, κ_t) in Prop. 2 with $a(t) = e^{-\Lambda(t)}$.

Sample mini-batch indices $I \subset \{1, \dots, N\}$ with |I| = B

Amortized start: draw $\mathbf{U}_0 \sim p_{\theta}^0(\cdot \mid \mathbf{X}_I) = \mathcal{N}(\mu_{\theta}(\mathbf{X}_I), \sigma^2 \mathbf{I})$

Initialize the integral accumulator $L_0 \leftarrow 0$

for k=0 to K-1 do

Set $t_s \leftarrow \frac{k}{K}, \quad t_{s+1} \leftarrow \frac{k+1}{K}, \quad \text{draw } \boldsymbol{\epsilon}_{t_s} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ Compute conditional score: $s_{\text{cond}} \big(t_s, \mathbf{U}_{t_s}, \mathbf{U}_0 \big) \leftarrow s_{\phi} \big(t_s, \mathbf{U}_{t_s}, \mathbf{X}_I, \mathbf{y}_I \big) \quad - \quad h \big(t_s, \mathbf{U}_{t_s}, \mathbf{U}_0 \big)$

Reverse bridge SDE step:

 $\mathbf{U}_{t_{s+1}} = \mathbf{U}_{t_s} - f(\mathbf{U}_{t_s}, t_s) \Delta t + g(t_s)^2 s_{\text{cond}}(t_s, \mathbf{U}_{t_s}, \mathbf{U}_0) \Delta t + g(t_s) \sqrt{\Delta t} \epsilon_{t_s}$ Bridge marginal update: obtain $\mathbf{m}_{t_{s+1}}, \kappa_{t_{s+1}}$

Accumulate score-matching term:

$$L_{t_{s+1}} = L_{t_s} + g(t_{s+1})^2 \left\| \frac{\mathbf{U}_{t_{s+1}} - \mathbf{m}_{t_{s+1}}}{\kappa_{t_{s+1}}} + s_{\text{cond}}(t_{s+1}, \mathbf{U}_{t_{s+1}}, \mathbf{U}_0) \right\|_2^2 \Delta t$$

end for

Set $\{\mathbf{u}^{(\ell)}\}_{\ell=1}^L \leftarrow \mathbf{U}_1$

for $\ell = 1$ to L do

Draw $\epsilon^{(\ell)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and compute

$$\mathbf{f}^{(\ell)} \ = \ \mathbf{K}_{\mathbf{F}^{(\ell-1)}\mathbf{Z}}^{(\ell)} \big(\mathbf{K}_{\mathbf{Z}\mathbf{Z}}^{(\ell)}\big)^{-1} \mathbf{u}^{(\ell)} \ + \ \Big(\mathbf{K}_{\mathbf{F}^{(\ell-1)}\mathbf{F}^{(\ell-1)}}^{(\ell)} - \mathbf{K}_{\mathbf{F}^{(\ell-1)}\mathbf{Z}}^{(\ell)} \big(\mathbf{K}_{\mathbf{Z}\mathbf{Z}}^{(\ell)}\big)^{-1} \mathbf{K}_{\mathbf{Z}\mathbf{F}^{(\ell-1)}}^{(\ell)}\Big)^{\frac{1}{2}} \boldsymbol{\epsilon}^{(\ell)}$$

end for

Mini-batch ELBO:

$$\begin{array}{ll} \ell_{\mathrm{DBVI}}(\theta,\phi,\gamma) &= -\log p_{\theta}^{0}\!\!\left(\mathbf{u}_{1} \mid \mathbf{X}_{I}\right) \,+\, \log p_{\mathrm{prior}}\!\!\left(\mathbf{u}_{1}\right) \,+\, \frac{N}{B}\log p\!\!\left(\mathbf{y}_{I} \mid \mathbf{F}^{(L)}\right) \\ &- \frac{1}{2}\,L_{1} \,-\, \mathrm{KL}\!\!\left(\mathcal{N}\!\!\left(\mu_{\theta}\!\left(\mathbf{X}_{I}\right),\sigma^{2}\mathbf{I}\right) \,\middle\|\, \mathcal{N}\!\!\left(\mathbf{m}_{1},\kappa_{1}\mathbf{I}\right)\right) \\ \text{Gradient update of } \ell_{\mathrm{DBVI}}\!\!\left(\theta,\phi,\gamma\right) \\ \mathbf{until} \;\theta,\phi,\gamma \; \text{converge} \end{array}$$




Figure 2: DBVI amortized initialization using inducing inputs $Z^{(l)}$.

the representations would lead to vectors of extremely large size and break the possibility of efficient batch processing.

To address these issues, we propose to employ the inducing points $\mathbf{Z}^{(l)}$ at each layer themselves as the input to the amortizer. This idea is motivated by classical sparse GP theory, where Z is interpreted as a set of representative features of the entire dataset under the joint Gaussian assumption.

Formally, the inducing inputs at layer l, denoted by $\mathbf{Z}^{(l)} \in \mathbb{R}^{M_l \times d_{\mathrm{in}}}$, are naturally aligned with the dimension of the corresponding variational parameters $\mathbf{u}^{(l)}$. To parameterize the initialization, we define a layer-wise neural network $\mu_{\theta}^{(l)}: \mathbb{R}^{d_{\mathrm{in}}} \to \mathbb{R}^{d_{\mathrm{out}}}$ that maps each inducing input $\mathbf{z}^{(l)} \in \mathbb{R}^{d_{\mathrm{in}}}$ to a representation in $\mathbb{R}^{d_{\mathrm{out}}}$. Applying this mapping to all inducing points yields

$$\mu_{\theta}^{(l)}(\mathbf{Z}^{(l)}) \in \mathbb{R}^{M_l \times d_{\text{out}}}.$$

This design has two main advantages: (i) it exploits the knowledge already contained in the inducing locations $\mathbf{Z}^{(l)}$, which evolve adaptively during model training, and (ii) it provides a simple yet effective amortization scheme whose dimension naturally matches the variational variables. In this way, $\mu_{\theta}^{(l)}$ achieves amortization without relying on either the full dataset.




Figure 3: Test RMSE and test mean NLL (with one standard deviation error bars) of deep Gaussian processes with different inference methods (DDVI, IPVI, SGHMC, DSVI, and our proposed DBVI) across 10 benchmark datasets .

4 EXPERIMENTS

We empirically evaluate DBVI against recent state-of-the-art inference methods for Deep Gaussian Processes (DGPs). Our evaluation covers regression on UCI benchmarks, image classification on standard vision datasets, large-scale physics datasets, and an unsupervised reconstruction task. Across these diverse settings, we assess both predictive performance and posterior quality, with particular attention to convergence behavior and scalability. The goal of our experiments is to demonstrate that DBVI consistently improves predictive accuracy and uncertainty estimation while remaining computationally efficient.

4.1 BASELINES AND SETUP

We compare DBVI with the following baselines: **DSVI** (Salimbeni & Deisenroth, 2017), the standard mean-field Gaussian variational approximation; **IPVI** (Yu et al., 2019), which parameterizes the posterior with a neural network trained adversarially; **SGHMC** (Havasi et al., 2018), a sampling-based inference approach; and **DDVI** (Xu et al., 2024), the diffusion-based inference method upon which DBVI builds.

All models use RBF kernels and M=128 inducing points per layer unless otherwise specified. For fairness, we adopt identical initialization and hyperparameter ranges across methods, and optimize using Adam with learning rate 0.01.

4.2 REGRESSION ON UCI BENCHMARKS

We evaluate our method on 10 widely used UCI regression datasets with sample sizes N ranging from a few hundred to over two million, using an 80/20 train/test split. We report root mean squared error (RMSE) and negative log-likelihood (NLL) on the held-out test data, as summarized in Figure 3. Consistent with prior work, we consider deep Gaussian processes with 2-5 layers. The results demonstrate that DBVI consistently outperforms all baseline methods, with particularly pronounced gains on large-scale datasets such as YearMSD and Airline, where unconditional DDVI suffers from slow convergence. By leveraging amortized bridge initialization, DBVI effectively shortens the diffusion path length, resulting in both improved posterior approximation and lower predictive error. Figures 4, 5, and 6 show the test RMSE of a 2-layer DGP trained with DDVI and DBVI during the early stage of optimization.

4.3 IMAGE CLASSIFICATION

We next evaluate DBVI on MNIST, Fashion-MNIST, and CIFAR-10. For CIFAR-10, we adopt ResNet-20 convolutional features as inputs to the DGP classifier. We report both test accuracy and per-iteration runtime in Table 1. Across all three datasets, DBVI consistently surpasses DDVI and other baselines. In particular, on CIFAR-10 with 4-layer DGPs, DBVI achieves an accuracy of 95.68%, slightly higher than DDVI's 95.56%. These findings underscore the advantage of amortized conditioning in handling complex high-dimensional posteriors.

4.4 LARGE-SCALE CLASSIFICATION

We further evaluate DBVI on two large-scale physics datasets, SUSY (5.5M examples) and HIGGS (11M examples). We report AUC scores with 2–5 layer DGPs under random 90/10 train-test splits. As shown in Table 2 in Appendix, DBVI consistently outperforms DDVI across all depths, yielding the best overall AUC values. Compared to SGHMC, DBVI attains comparable or higher performance while being substantially more computationally efficient, highlighting its scalability and effectiveness in modeling complex, high-dimensional posteriors at scale.

4.5 Unsupervised Reconstruction

Finally, we evaluate posterior quality on the Frey Faces dataset using a missing-data reconstruction task. Following prior work, we randomly mask 75% of pixels for a subset of training images and task the models with recovering the originals. Table 3 reports reconstruction RMSE and test log-likelihood. DBVI achieves the lowest RMSE and highest likelihood, surpassing DDVI as well as variational and sampling-based baselines. In qualitative comparisons, reconstructions generated by DBVI are visually sharper and demonstrate better calibrated uncertainty than those of competing methods.

5 CONCLUSION

We introduced Diffusion Bridge Variational Inference (DBVI), a principled extension of DDVI that initiates the reverse diffusion from a learnable, input-conditioned distribution. By incorporating Doob's h-transform into the variational formulation, DBVI preserves the theoretical elegance of diffusion-based inference while substantially reducing the inference gap. Our structured amortization strategy, which conditions on inducing inputs, further enables scalable and data-efficient posterior approximation in deep Gaussian processes. Empirical results across regression, classification, and image reconstruction tasks confirm that DBVI consistently improves over DDVI and other state-of-the-art variational baselines in predictive accuracy, convergence speed, and sample efficiency. These findings highlight the benefits of bridging-based inference for scalable Bayesian learning. Future directions include extending DBVI to broader probabilistic models, integrating with advanced inducing-point strategies for large-scale applications, and exploring theoretical guarantees of diffusion bridges in variational inference.

REFERENCES

- Abhinav Agrawal and Justin Domke. Amortized variational inference for simple hierarchical models. *Advances in Neural Information Processing Systems*, 34:21388–21399, 2021.
- Andreas Damianou and Neil Lawrence. Deep Gaussian processes. In *Conference on Artificial Intelligence and Statistics*, pp. 207–215, 2013.
- Ankush Ganguly, Sanjana Jain, and Ukrit Watchareeruetai. Amortized variational inference: A systematic review. *Journal of Artificial Intelligence Research*, 78:167–215, 2023.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems, 27, 2014.
- Marton Havasi, José Miguel Hernández-Lobato, and Juan José Murillo-Fuentes. Inference in deep Gaussian processes using stochastic gradient Hamiltonian Monte Carlo. In *Conference on Neural Information Processing Systems*, pp. 7517–7527, 2018.
- Guande He, Kaiwen Zheng, Jianfei Chen, Fan Bao, and Jun Zhu. Consistency diffusion bridge models. *Advances in Neural Information Processing Systems*, 37:23516–23548, 2024.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. arXiv preprint arXiv:1309.6835, 2013.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Yoon Kim, Sam Wiseman, Andrew Miller, David Sontag, and Alexander Rush. Semi-amortized variational autoencoders. In *International Conference on Machine Learning*, pp. 2678–2687. PMLR, 2018.
- Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pp. 1952–1961, 2023.
- Charles C Margossian and David M Blei. Amortized variational inference: When and why? *arXiv* preprint arXiv:2307.11018, 2023.
- Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep Gaussian processes. In *Conference on Neural Information Processing Systems*, pp. 4588–4599, 2017.
- Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv* preprint *arXiv*:2011.13456, 2020.
- Jian Xu, Delu Zeng, and John Paisley. Sparse inducing points in deep gaussian processes: Enhancing modeling with denoising diffusion variational inference. In *International Conference on Machine Learning*, pp. 55490–55500. PMLR, 2024.
- Haibin Yu, Yizhou Chen, Bryan Kian Hsiang Low, Patrick Jaillet, and Zhongxiang Dai. Implicit posterior variational inference for deep gaussian processes. *Advances in neural information processing systems*, 32, 2019.
- Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. Denoising diffusion bridge models. *arXiv preprint arXiv:2309.16948*, 2023.

Table 1: Mean test accuracy (%) and training details achieved by DSVI, SGHMC, IPVI, DDVI, and our proposed DBVI on three image classification datasets. Results are shown for 3 and 4 layers as indicated, and runtime is given per iteration.

Data Set	Model	Time3	Iter3	Acc3	Time4	Iter4	Acc4
	DSVI	0.34s	20K	97.17	0.54s	20K	97.41
MNIST	IPVI	0.49s	20K	97.58	0.62s	20K	97.80
	SGHMC	1.14s	20K	97.25	1.22s	20K	97.55
	DDVI	0.38s	20K	98.84	0.50s	20K	99.01
	DBVI (ours)	0.41s	20K	99.02	0.55s	20K	99.10
Fashion	DSVI	0.34s	20K	87.45	0.50s	20K	87.99
	IPVI	0.48s	20K	88.23	0.61s	20K	88.90
	SGHMC	1.21s	20K	86.88	1.25s	20K	87.08
	DDVI	0.40s	20K	90.36	0.55s	20K	90.85
	DBVI (ours)	0.43s	20K	90.53	0.60s	20K	91.07
CIFAR-10	DSVI	0.43s	20K	91.47	0.66s	20K	91.79
	IPVI	0.62s	20K	92.79	0.78s	20K	93.52
	SGHMC	8.04s	20K	92.62	8.61s	20K	92.94
	DDVI	0.45s	20K	95.23	0.69s	20K	95.56
	DBVI (ours)	0.49s	20K	95.42	0.74s	20K	95.68

Table 2: Test AUC values for large-scale classification datasets. Uses random 90% / 10% training and test splits.

		SUSY	HIGGS
	N	5,500,000	11,000,000
	D	18	28
	L=2	0.876	0.830
DSVI	L = 3	0.877	0.837
M = 128	L = 4	0.878	0.841
	L = 5	0.878	0.846
	L=2	0.879	0.843
IPVI	L = 3	0.882	0.847
M = 128	L=4	0.883	0.850
	L = 5	0.883	0.852
	L=2	0.879	0.842
SGHMC	L = 3	0.881	0.846
M = 128	L=4	0.883	0.850
	L = 5	0.884	0.853
	L=2	0.883	0.849
DDVI	L = 3	0.885	0.852
M = 128	L=4	0.887	0.856
	L = 5	0.886	0.857
	L=2	0.885	0.851
DBVI (ours)	L=3	0.887	0.854
M = 128	L=4	0.889	0.858
	L = 5	0.889	0.859

A PROOF OF PROPOSITIONS

Proposition 4 (Forward & reverse SDE under Doob's h-transform). Let the initial constraint be encoded by the Doob h-transform with

$$h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t), \tag{17}$$

Then the forward bridge has drift

$$\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) = f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0), \tag{18}$$

with the same diffusion coefficient g(t). Moreover, the reverse-time bridge SDE is

$$d\mathbf{U}_{t} = \left[f(\mathbf{U}_{t}, t) - g(t)^{2} s_{cond}(\mathbf{U}_{t}, t, \mathbf{U}_{0}) \right] dt + g(t) d\mathbf{W}_{t},$$
(19)

Equivalently, the **conditional score** equals $s_{cond}(\mathbf{U}_t, t, \mathbf{U}_0) = s(\mathbf{U}_t, t, \mathbf{u}_o) - h(\mathbf{U}_t, t, \mathbf{U}_0)$.

Table 3: Mean RMSE and NLL achieved by DSVI, SGHMC, IPVI, DDVI, and our proposed DBVI on the GPLVM data recovery task (Frey Faces). Standard deviation is shown in parentheses. Runtime is given per iteration.

Data Set	Model	Time	Iter	RMSE	NLL
Frey Faces	DSVI IPVI SGHMC DDVI	0.32s 0.42s 1.13s 0.36s	20K 20K 20K 20K	8.32 (0.20) 7.91 (0.40) 7.95 (0.30) 7.64 (0.20)	1.49 (0.02) 1.33 (0.02) 1.36 (0.03) 1.17 (0.01)
	DBVI (ours)	0.40s	20K	7.52 (0.18)	1.12 (0.01)

Proof. **Setup.** Consider the forward diffusion on [0, 1],

$$d\mathbf{U}_t = f(\mathbf{U}_t, t) dt + g(t) d\mathbf{B}_t, \tag{20}$$

where \mathbf{B}_t is a standard Brownian motion, and g(t) > 0 is scalar. Assume standard regularity (existence/uniqueness, non-degeneracy, smooth strictly positive densities). Let p_t denote the time-t marginal density of equation 20. The (backward) Kolmogorov operator is

$$(\mathcal{L}_t \varphi)(u) = f(u, t) \cdot \nabla \varphi(u) + \frac{g(t)^2}{2} \Delta \varphi(u). \tag{21}$$

(i) Doob h-transform with an initial-point constraint. Define the space-time positive function

$$H(t,u) := p(\mathbf{U}_0 \mid \mathbf{U}_t = u). \tag{22}$$

By the Markov property and Chapman–Kolmogorov, H is space–time harmonic for $\partial_t + \mathcal{L}_t$, i.e.

$$\partial_t H(t, u) + (\mathcal{L}_t H)(t, u) = 0. (23)$$

The Doob H-transform of equation 20 is the time-inhomogeneous Markov process with generator

$$(\mathcal{L}_t^H \varphi)(u) = H(t, u)^{-1} \left(\mathcal{L}_t(H\varphi) \right)(u). \tag{24}$$

Expanding $\mathcal{L}_t(H\varphi)$ (see remarks below) and using equation 23 to cancel the $\partial_t H$ term yields, for smooth φ ,

$$\mathcal{L}_{t}^{H}\varphi(u) = f(u,t)\cdot\nabla\varphi(u) + \frac{g(t)^{2}}{2}\Delta\varphi(u) + g(t)^{2}\operatorname{\nabla log}H(t,u)\cdot\nabla\varphi(u). \tag{25}$$

Therefore, the transformed process is again an Itô diffusion with the same diffusion coefficient g(t) and drift

$$\tilde{f}(u, t, \mathbf{U}_0) = f(u, t) + g(t)^2 \nabla_u \log H(t, u) = f(u, t) + g(t)^2 \nabla_u \log p(\mathbf{U}_0 \mid u).$$
 (26)

With the proposition's definition

$$h(\mathbf{U}_t, t, \mathbf{U}_0) := \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t), \tag{27}$$

we obtain the forward bridge SDE

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0) \right] dt + g(t) d\mathbf{B}_t,$$
 (28)

i.e. $\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) = f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0)$ with unchanged diffusion g(t).

(ii) Reverse-time SDE of the initial-point bridge. Let $\tilde{p}_t(u) := p(\mathbf{U}_t = u \mid \mathbf{U}_0)$ denote the time-t conditional density of the bridge. By the Haussmann-Pardoux time-reversal formula for diffusions with isotropic diffusion matrix $a(t) = g(t)^2 I$ (so $\nabla \cdot a \equiv 0$), the time-reversed bridge $\overleftarrow{\mathbf{U}}_t = \mathbf{U}_{1-t}$ solves

$$d\overleftarrow{\mathbf{U}}_{t} = \left[\widetilde{f}(\overleftarrow{\mathbf{U}}_{t}, 1 - t, \mathbf{U}_{0}) - g(1 - t)^{2}\nabla\log\widetilde{p}_{1 - t}(\overleftarrow{\mathbf{U}}_{t})\right]dt + g(1 - t)\,d\overleftarrow{\mathbf{W}}_{t}.$$
 (29)

Re-indexing $t \mapsto 1 - t$ and renaming the Brownian motion gives the reverse–time SDE in forward orientation:

$$d\mathbf{U}_t = \left[\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) - g(t)^2 \nabla \log \tilde{p}_t(\mathbf{U}_t)\right] dt + g(t) d\mathbf{W}_t.$$
(30)

Substituting $\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) = f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0)$ yields

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0) - g(t)^2 \nabla \log \tilde{p}_t(\mathbf{U}_t) \right] dt + g(t) d\mathbf{W}_t.$$
 (31)

Define the conditional score

$$s_{\text{cond}}(\mathbf{U}_t, t, \mathbf{U}_0) := \nabla_{\mathbf{U}_t} \log \tilde{p}_t(\mathbf{U}_t) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_t \mid \mathbf{U}_0), \tag{32}$$

to obtain the claimed reverse-time bridge SDE:

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) - g(t)^2 s_{\text{cond}}(\mathbf{U}_t, t, \mathbf{U}_0) \right] dt + g(t) d\mathbf{W}_t.$$
 (33)

(iii) Conditional score identity (for the initial-point constraint). By Bayes' rule,

$$\log p(\mathbf{U}_t \mid \mathbf{U}_0) = \log p(\mathbf{U}_0 \mid \mathbf{U}_t) + \log p_t(\mathbf{U}_t) - \log p(\mathbf{U}_0). \tag{34}$$

Taking $\nabla_{\mathbf{U}_t}$ gives

$$s_{\text{cond}}(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p_t(\mathbf{U}_t) + h(\mathbf{U}_t, t, \mathbf{U}_0).$$
(35)

If we write the unconditional score as $s(\mathbf{U}_t, t) := \nabla_{\mathbf{U}_t} \log p_t(\mathbf{U}_t)$, this is

$$s_{\text{cond}} = s + h. (36)$$

Remarks: explicit derivation for Doob *H*-transform generator. Fix $0 \le s \le t \le 1$. The transformed semigroup acting on a test function φ is

$$(T_{s,t}^H \varphi)(u) := \frac{1}{H(s,u)} \mathbb{E} \left[H(t, \mathbf{U}_t) \varphi(\mathbf{U}_t) \, \middle| \, \mathbf{U}_s = u \right]. \tag{37}$$

The (backward) generator \mathcal{L}_t^H is defined by

$$\partial_t (T_{s,t}^H \varphi)(u) = (T_{s,t}^H \mathcal{L}_t^H \varphi)(u), \quad \text{with} \quad T_{t,t}^H = \text{Id}.$$
 (38)

Applying Itô to the product $H(t, \mathbf{U}_t) \varphi(\mathbf{U}_t)$ under the forward SDE 20 and taking conditional expectations gives

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathbb{E} \left[H(t, \mathbf{U}_t) \, \varphi(\mathbf{U}_t) \, \middle| \, \mathbf{U}_s = u \right] = \mathbb{E} \left[\left(\partial_t + \mathcal{L}_t \right) \left(H(t, \mathbf{U}_t) \, \varphi(\mathbf{U}_t) \right) \, \middle| \, \mathbf{U}_s = u \right]. \tag{39}$$

By the space–time harmonicity of H,

$$(\partial_t + \mathcal{L}_t)H(t, u) = 0, (40)$$

we can expand $(\partial_t + \mathcal{L}_t)(H\varphi)$ and cancel the $\partial_t H$ terms cleanly. Since the backward generator is

$$(\mathcal{L}_t \psi)(u) = f(u, t) \cdot \nabla \psi(u) + \frac{g(t)^2}{2} \Delta \psi(u), \tag{41}$$

the product rules for gradient and Laplacian give

$$\nabla(H\varphi) = H \,\nabla\varphi + \varphi \,\nabla H, \qquad \Delta(H\varphi) = H \,\Delta\varphi + \varphi \,\Delta H + 2 \,\nabla H \cdot \nabla\varphi. \tag{42}$$

Hence

$$\mathcal{L}_t(H\varphi) = f \cdot \nabla(H\varphi) + \frac{g(t)^2}{2} \Delta(H\varphi) \tag{43}$$

$$= H \underbrace{\left(f \cdot \nabla \varphi + \frac{g(t)^2}{2} \Delta \varphi\right)}_{=C_t \cdot \rho} + \varphi \underbrace{\left(f \cdot \nabla H + \frac{g(t)^2}{2} \Delta H\right)}_{=C_t \cdot H} + g(t)^2 \nabla H \cdot \nabla \varphi \tag{44}$$

$$= H \mathcal{L}_t \varphi + \varphi \mathcal{L}_t H + g(t)^2 \nabla H \cdot \nabla \varphi. \tag{45}$$

Therefore,

$$(\partial_t + \mathcal{L}_t)(H\varphi) = (\partial_t H)\varphi + H\partial_t \varphi + H\mathcal{L}_t \varphi + \varphi \mathcal{L}_t H + g(t)^2 \nabla H \cdot \nabla \varphi$$
 (46)

$$= H(\partial_t + \mathcal{L}_t)\varphi + \underbrace{((\partial_t + \mathcal{L}_t)H)}_{=0 \text{ by } equation } \varphi + g(t)^2 \nabla H \cdot \nabla \varphi$$
 (47)

$$= H(\partial_t + \mathcal{L}_t)\varphi + g(t)^2 \nabla H \cdot \nabla \varphi. \tag{48}$$

Dividing by H(t, u) (i.e., multiplying by the scalar $H(t, u)^{-1}$) gives the pointwise identity

$$\frac{1}{H} (\partial_t + \mathcal{L}_t)(H\varphi) = (\partial_t + \mathcal{L}_t)\varphi + g(t)^2 \operatorname{\nabla log} H \cdot \operatorname{\nabla} \varphi.$$
(49)

Comparing with the definition of the transformed semigroup $T^H_{s,t}$ (which removes the explicit ∂_t on the left via the semigroup relation), we identify the *backward generator* of the H-transform at time t as

$$(\mathcal{L}_t^H \varphi)(u) = (\mathcal{L}_t \varphi)(u) + g(t)^2 \operatorname{\nabla log} H(t, u) \cdot \operatorname{\nabla} \varphi(u).$$
 (50)

Since $H(t,u) = p(\mathbf{U}_0 \mid \mathbf{U}_t = u)$, writing $h(\mathbf{U}_t,t,\mathbf{U}_0) := \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t)$ yields the drift correction

$$\tilde{f}(\mathbf{U}_t, t, \mathbf{U}_0) = f(\mathbf{U}_t, t) + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0), \tag{51}$$

with the *same* diffusion coefficient g(t).

Proposition 5 (Marginal of Doob-augmented bridge process). *Consider the linear forward SDE with Doob bridge correction*

$$d\mathbf{U}_t = \left[-\lambda(t) \mathbf{U}_t + g(t)^2 h(\mathbf{U}_t, t, \mathbf{U}_0) \right] dt + g(t) d\mathbf{B}_t,$$

where $h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t)$ is the Doob h-transform. Assume isotropic initialization

$$p_0^{\theta}(\mathbf{U}_0 \mid \mathbf{x}) = \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I}).$$

Then for each $t \in [0, 1]$, the marginal law remains Gaussian,

$$p_t(\mathbf{U}^{Bri} \mid \mathbf{x}) = \mathcal{N}(\mathbf{U}^{Bri}; \mathbf{m}_t, \kappa_t \mathbf{I}),$$

where the mean \mathbf{m}_t and variance κ_t satisfy the coupled ODE system

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{m}_{t} = -(\lambda(t) + c(t))\mathbf{m}_{t} + c(t)a(t)\mu_{\theta}(\mathbf{x}), \quad \mathbf{m}_{0} = \mu_{\theta}(\mathbf{x}),$$
 (52)

$$\frac{\mathrm{d}}{\mathrm{d}t}\kappa_t = -2(\lambda(t) + c(t))\kappa_t + g(t)^2 + 2c(t)a(t)\sigma^2, \quad \kappa_0 = \sigma^2,$$
(53)

with

$$a(t) = e^{-\Lambda(t)}, \quad \Lambda(t) = \int_0^t \lambda(s) \, \mathrm{d}s,$$

and correction coefficient

$$c(t) = g(t)^2 \frac{\sigma^2 a(t)^2}{(a(t)^2 \sigma^2 + q(t)) q(t)}, \qquad q(t) = a(t)^2 \int_0^t \frac{g(r)^2}{a(r)^2} dr.$$

In the special case $c(t) \equiv 0$, this reduces to the bridge process trick in DDVI..

Proof. We work in the isotropic, linear-Gaussian setting adopted in the main text: the base forward SDE is

$$d\mathbf{U}_t = -\lambda(t) \mathbf{U}_t dt + g(t) d\mathbf{B}_t, \qquad \mathbf{U}_0 \sim \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I}), \tag{54}$$

with $a(t)=e^{-\Lambda(t)}$ and $\Lambda(t)=\int_0^t \lambda(s)\,\mathrm{d}s.$ In this OU setting,

$$\mathbb{E}[\mathbf{U}_t] = a(t)\,\mu_\theta(\mathbf{x}), \qquad \text{Cov}(\mathbf{U}_t) = a(t)^2 \sigma^2 \mathbf{I} + q(t)\,\mathbf{I}, \tag{55}$$

where

$$q(t) = a(t)^2 \int_0^t \frac{g(r)^2}{a(r)^2} dr.$$
 (56)

Moreover, the joint law of $(\mathbf{U}_0, \mathbf{U}_t)$ is Gaussian with

$$Cov(\mathbf{U}_0, \mathbf{U}_t) = \sigma^2 a(t) \mathbf{I}. \tag{57}$$

These identities are standard and match the bridge process used in the DDVI paper.

Step 1: Conditional $p(\mathbf{U}_0 \mid \mathbf{U}_t)$. By joint-Gaussian conditioning,

$$p(\mathbf{U}_0 \mid \mathbf{U}_t) = \mathcal{N}(\mathbf{m}_{0|t}, \mathbf{S}_{0|t}), \tag{58}$$

with isotropic (scalar) parameters

$$\boldsymbol{m}_{0|t} = \mu_{\theta}(\mathbf{x}) + K(t) \left(\mathbf{U}_t - a(t)\mu_{\theta}(\mathbf{x}) \right), \qquad \boldsymbol{S}_{0|t} = S(t) \mathbf{I},$$
 (59)

where

$$K(t) = \frac{\sigma^2 a(t)}{a(t)^2 \sigma^2 + q(t)}, \qquad S(t) = \sigma^2 - \frac{\sigma^4 a(t)^2}{a(t)^2 \sigma^2 + q(t)} = \frac{\sigma^2 q(t)}{a(t)^2 \sigma^2 + q(t)}. \tag{60}$$

Step 2: Doob *h***-transform (bridge) term.** Define the Doob correction

$$h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t). \tag{61}$$

Since $p(\mathbf{U}_0 \mid \mathbf{U}_t) = \mathcal{N}(\mathbf{m}_{0|t}, S(t)\mathbf{I})$ and $\mathbf{m}_{0|t}$ is affine in \mathbf{U}_t with coefficient $K(t)\mathbf{I}$, we have

$$\nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t) = \left(\nabla_{\mathbf{U}_t} \boldsymbol{m}_{0|t}\right)^{\top} S(t)^{-1} \left(\mathbf{U}_0 - \boldsymbol{m}_{0|t}\right) = \frac{K(t)}{S(t)} \left(\mathbf{U}_0 - \boldsymbol{m}_{0|t}\right). \tag{62}$$

Hence the Doob-augmented forward SDE reads

$$d\mathbf{U}_t = \left[-\lambda(t) \mathbf{U}_t + g(t)^2 \frac{K(t)}{S(t)} (\mathbf{U}_0 - \mathbf{m}_{0|t}) \right] dt + g(t) d\mathbf{B}_t. \tag{*}$$

Step 3: Mean dynamics. Let $\mathbf{m}_t := \mathbb{E}[\mathbf{U}_t]$. Taking expectations in (\star) and using $\mathbb{E}[\mathbf{U}_0] = \mu_{\theta}(\mathbf{x})$ and $m_{0|t} = \mu_{\theta}(\mathbf{x}) + K(t) (\mathbf{U}_t - a(t)\mu_{\theta}(\mathbf{x}))$ gives

$$\mathbb{E}\left[\mathbf{U}_{0} - \boldsymbol{m}_{0|t}\right] = \mu_{\theta}(\mathbf{x}) - \left(\mu_{\theta}(\mathbf{x}) + K(t)\left(\mathbf{m}_{t} - a(t)\mu_{\theta}(\mathbf{x})\right)\right) = K(t)\left(a(t)\mu_{\theta}(\mathbf{x}) - \mathbf{m}_{t}\right).$$
(63)

Therefore,

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathbf{m}_t = -\lambda(t) \mathbf{m}_t + g(t)^2 \frac{K(t)}{S(t)} K(t) \left(a(t) \mu_{\theta}(\mathbf{x}) - \mathbf{m}_t \right). \tag{64}$$

Introduce the scalar

$$c(t) = g(t)^{2} \frac{K(t)^{2}}{S(t)} = g(t)^{2} \frac{\sigma^{2} a(t)^{2}}{\left(a(t)^{2} \sigma^{2} + q(t)\right) q(t)},$$
(65)

which yields

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathbf{m}_t = -(\lambda(t) + c(t)) \mathbf{m}_t + c(t) a(t) \mu_{\theta}(\mathbf{x}), \qquad \mathbf{m}_0 = \mu_{\theta}(\mathbf{x}). \tag{66}$$

Step 4: Variance dynamics. Let κ_t denote the (isotropic) variance so that $Cov(\mathbf{U}_t) = \kappa_t \mathbf{I}$. From Itô's lemma for $\mathbf{U}_t \mathbf{U}_t^{\mathsf{T}}$ under (\star) and isotropy,

$$\frac{\mathrm{d}}{\mathrm{d}t} \kappa_t = -2\lambda(t) \kappa_t + g(t)^2 + 2g(t)^2 \frac{K(t)}{S(t)} \operatorname{Cov}(\mathbf{U}_t, \mathbf{U}_0 - \mathbf{m}_{0|t})_{\text{scalar}}.$$
 (67)

Using $Cov(\mathbf{U}_t, \mathbf{U}_0) = a(t)\sigma^2\mathbf{I}$ and $\mathbf{m}_{0|t} = \mu_{\theta}(\mathbf{x}) + K(t)(\mathbf{U}_t - a(t)\mu_{\theta}(\mathbf{x}))$, we get

$$Cov(\mathbf{U}_t, \mathbf{m}_{0|t}) = K(t) Cov(\mathbf{U}_t, \mathbf{U}_t) = K(t) \kappa_t \mathbf{I},$$
(68)

hence

$$Cov(\mathbf{U}_t, \mathbf{U}_0 - \mathbf{m}_{0|t}) = a(t)\sigma^2 \mathbf{I} - K(t)\kappa_t \mathbf{I}.$$
 (69)

Plugging in and using $c(t) = g(t)^2 K(t)^2 / S(t)$,

$$\frac{\mathrm{d}}{\mathrm{d}t} \kappa_t = -2\lambda(t) \kappa_t + g(t)^2 + 2 \frac{c(t)}{K(t)} \Big(a(t)\sigma^2 - K(t)\kappa_t \Big) = -2 \Big(\lambda(t) + c(t) \Big) \kappa_t + g(t)^2 + 2 c(t) a(t) \sigma^2,$$
(70)

with $\kappa_0 = \sigma^2$.

Step 5: Gaussian form. The drift in (\star) is affine in $(\mathbf{U}_t, \mathbf{U}_0)$, and the driving noise is Gaussian. Therefore $(\mathbf{U}_t, \mathbf{U}_0)$ remains jointly Gaussian, and the marginal $p_t(\mathbf{U}^{\mathrm{Br}} \mid \mathbf{x})$ is Gaussian with mean \mathbf{m}_t and variance $\kappa_t \mathbf{I}$. This proves the claimed form and the coupled ODE system. In the special case $c(t) \equiv 0$ (no Doob correction), the system reduces to the closed-form bridge process in DDVI paper, i.e., $\mathbf{m}_t = a(t)\mu_\theta(\mathbf{x})$ and $\kappa_t = a(t)^2\sigma^2 + a(t)^2\int_0^t \frac{g(r)^2}{a(r)^2}\mathrm{d}r$.

Proposition 6 (DBVI loss with amortized mean). Let $p_0^{\theta}(\mathbf{u} \mid \mathbf{x}) = \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I})$ be the data-dependent start, and let (\mathbf{m}_t, κ_t) be the mean/variance of the reference bridge marginal at time $t \in [0, 1]$ (given by Proposition 2 via ODEs in the Doob-augmented case). Then the pathwise KL between the variational reverse bridge Q_{ϕ} and the reference bridge admits a score-matching representation. Consequently, a tractable per-mini-batch ELBO is

$$\ell_{\text{DBVI}}(\theta, \phi, \gamma) = \mathbb{E}_{\mathbf{U}_{0:1} \sim Q_{\phi}} \left[-\log p_{0}^{\theta}(\mathbf{U}_{1}) + \frac{B}{N} \log p(\mathbf{y}_{I} \mid \mathbf{f}^{(L)}) - \frac{1}{2} \int_{0}^{1} g(t)^{2} \left\| \frac{\mathbf{U}_{t} - \mathbf{m}_{t}}{\kappa_{t}} + s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) \right\|_{2}^{2} dt - \text{KL} \left(\mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^{2}\mathbf{I}) \mid | \mathcal{N}(\mathbf{m}_{1}, \kappa_{1}\mathbf{I}) \right) + \log p_{\text{prior}}(\mathbf{U}_{1}) \right],$$
(71)

where B is the batch size, N is the dataset size, $s_{\rm cond} = s_{\phi} - h$ is the conditional score used by the reverse bridge SDE, \mathbf{U}_1 is its terminal state, $\mathbf{f}^{(L)}$ denotes the DGP forward mapping at \mathbf{U}_1 , and γ collects DGP hyperparameters. When $\mu_{\theta}(\mathbf{x}) = \mathbf{0}$ (and thus $\mathbf{m}_t \equiv \mathbf{0}$), the objective recovers the original DDVI loss.

Proof. **Setup.** Let Q_{ϕ} be the path law of the reverse-time bridge SDE

$$d\mathbf{U}_t = \left[f(\mathbf{U}_t, t) - g(t)^2 s_{\text{cond}}(t, \mathbf{U}_t, \mathbf{U}_0) \right] dt + g(t) d\mathbf{W}_t, \qquad s_{\text{cond}} := s_{\phi} - h,$$

with $h(\mathbf{U}_t, t, \mathbf{U}_0) = \nabla_{\mathbf{U}_t} \log p(\mathbf{U}_0 \mid \mathbf{U}_t)$. Let P be the forward reference diffusion and P^{Bri} be the auxiliary forward process that shares the same drift/diffusion as P but starts from $p_0^{\theta}(\mathbf{u} \mid \mathbf{x}) = \mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I})$. Assume Novikov's condition so that Girsanov applies.

(A) Path term \Rightarrow score-matching quadratic. By reverse-time Girsanov (see lemma 1), the process KL splits into a path term plus a boundary term; the path term is

$$\mathbb{E}_{\mathbf{U}_{0:1} \sim Q_{\phi}} \left[\log \frac{\mathrm{d}Q_{\phi}}{\mathrm{d}P^{\mathrm{Bri}}} \right] = \frac{1}{2} \int_{0}^{1} g(t)^{2} \,\mathbb{E}_{Q_{\phi}} \left[\| s_{\mathrm{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) - s_{\mathrm{Bri}}(t, \mathbf{U}_{t}) \|_{2}^{2} \right] \mathrm{d}t, \tag{72}$$

where $s_{\text{Bri}}(t, \mathbf{U}_t) := \nabla_{\mathbf{U}_t} \log p_t(\mathbf{U}_t)$ is the (marginal) score of the reference bridge at time t. By Proposition 2, the marginal is Gaussian $p_t(\mathbf{U}_t) = \mathcal{N}(\mathbf{m}_t, \kappa_t \mathbf{I})$, hence

$$s_{\text{Bri}}(t, \mathbf{U}_t) = -\frac{\mathbf{U}_t - \mathbf{m}_t}{\kappa_t}.$$
 (73)

Substituting this into equation 72 yields the integrand

$$\left\| s_{\text{cond}} - s_{\text{Bri}} \right\|_{2}^{2} = \left\| \frac{\mathbf{U}_{t} - \mathbf{m}_{t}}{\kappa_{t}} + s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) \right\|_{2}^{2},$$
(74)

which is exactly the score-matching form stated in the proposition.

(B) Boundary term. Since P and $P^{\rm Bri}$ share dynamics and only differ at the start, the remaining term in the KL decomposition collapses to a boundary difference (as in DDVI):

$$\mathbb{E}_{Q_{\phi}} \left[\log \frac{\mathrm{d}P^{\mathrm{Bri}}}{\mathrm{d}P} \right] = \mathbb{E}_{Q_{\phi}} \left[\log \frac{p_0^{\theta}(\mathbf{U}_1)}{p_1^{\mathrm{Bri}}(\mathbf{U}_1)} \right] = - \mathrm{KL} \left(\mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^2 \mathbf{I}) \| \mathcal{N}(\mathbf{m}_1, \kappa_1 \mathbf{I}) \right), \tag{75}$$

where $p_1^{\text{Bri}} = \mathcal{N}(\mathbf{m}_1, \kappa_1 \mathbf{I})$ follows from Proposition 2.

(C) Collecting the ELBO terms. Adding the model likelihood and prior contributions (as in DDVI; mini-batch scaled by N/B) yields, up to constants independent of (θ, ϕ, γ) ,

$$\ell_{\text{DBVI}}(\theta, \phi, \gamma) = \mathbb{E}_{\mathbf{U}_{0:1} \sim Q_{\phi}} \left[-\log p_{0}^{\theta}(\mathbf{U}_{1}) + \frac{B}{N} \log p(\mathbf{y}_{I} \mid \mathbf{f}^{(L)}) - \frac{1}{2} \int_{0}^{1} g(t)^{2} \left\| \frac{\mathbf{U}_{t} - \mathbf{m}_{t}}{\kappa_{t}} + s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) \right\|_{2}^{2} dt - \text{KL} \left(\mathcal{N}(\mu_{\theta}(\mathbf{x}), \sigma^{2}\mathbf{I}) \parallel \mathcal{N}(\mathbf{m}_{1}, \kappa_{1}\mathbf{I}) \right) + \log p_{\text{prior}}(\mathbf{U}_{1}) \right],$$
(76)

which matches the statement.

(D) Reduction to DDVI. If $\mu_{\theta}(\mathbf{x}) \equiv \mathbf{0}$, then $\mathbf{m}_t \equiv \mathbf{0}$ and the loss reduces to the original DDVI objective (same path integral with $s_{\text{cond}} = s_{\phi} - h$ and the usual boundary terms).

Lemma 1 (Reverse-time Girsanov: path term as score–matching). Fix T = 1. Consider two reverse-time SDEs on \mathbb{R}^d with the same diffusion scale g(t) > 0:

(Q)
$$d\mathbf{U}_t = \left(f(\mathbf{U}_t, t) - g(t)^2 s_{\text{cond}}(t, \mathbf{U}_t, \mathbf{U}_0) \right) dt + g(t) d\mathbf{W}_t^{(Q)}, \tag{77}$$

$$(P^{\text{Bri}}) \ d\mathbf{U}_t = \left(f(\mathbf{U}_t, t) - g(t)^2 s_{\text{Bri}}(t, \mathbf{U}_t) \right) dt + g(t) d\mathbf{W}_t^{(P)}, \tag{78}$$

where $s_{\rm cond} = s_{\phi} - h$ is the conditional (bridge-corrected) score, and $s_{\rm Bri}(t, {\bf u}) = \nabla_{\bf u} \log p_t({\bf u})$ is the marginal score of the reference bridge (with marginal $p_t({\bf u}) = \mathcal{N}({\bf m}_t, \kappa_t {\bf I})$). Assume standard integrability (e.g. Novikov) so that Girsanov applies and both path laws Q_{ϕ} and $P^{\rm Bri}$ are mutually absolutely continuous on \mathcal{F}_T . Then the pathwise Kullback–Leibler divergence decomposes as

$$KL(Q_{\phi} \parallel P^{Bri}) = \underbrace{\frac{1}{2} \int_{0}^{T} g(t)^{2} \mathbb{E}_{Q_{\phi}} \left[\left\| s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) - s_{\text{Bri}}(t, \mathbf{U}_{t}) \right\|_{2}^{2} \right] dt}_{path \ term} + \underbrace{\mathbb{E}_{Q_{\phi}} \left[\log \frac{q_{T}(\mathbf{U}_{T})}{p_{T}(\mathbf{U}_{T})} \right]}_{boundary \ term}, (79)$$

where q_T and p_T are the terminal densities of the two reverse processes (equivalently: the initial densities of the corresponding forward processes). In particular, the path term equals

$$\frac{1}{2} \int_{0}^{1} g(t)^{2} \mathbb{E}_{Q_{\phi}} \left[\| s_{\text{cond}}(t, \mathbf{U}_{t}, \mathbf{U}_{0}) - s_{\text{Bri}}(t, \mathbf{U}_{t}) \|_{2}^{2} \right] dt,$$
 (80)

which is the score-matching quadratic used in equation 72.

Proof. We write both SDEs on a common probability space up to an equivalent change of measure. Let the *reference* path law be $P^{\rm Bri}$ (drift $b_P:=f-g^2s_{\rm Bri}$). Under $P^{\rm Bri}$, define the progressively measurable process

$$\vartheta_t := \frac{b_Q(\mathbf{U}_t, t) - b_P(\mathbf{U}_t, t)}{g(t)} = -g(t) \left(s_{\text{cond}}(t, \mathbf{U}_t, \mathbf{U}_0) - s_{\text{Bri}}(t, \mathbf{U}_t) \right), \tag{81}$$

where $b_Q := f - g^2 s_{\text{cond}}$. By Novikov's condition, the Dôleans–Dade exponential

$$Z_t = \exp\left(\int_0^t \vartheta_s^\top d\mathbf{W}_s^{(P)} - \frac{1}{2} \int_0^t \|\vartheta_s\|^2 ds\right)$$
 (82)

is a true P^{Bri} -martingale. Define a new measure Q_{ϕ} on \mathcal{F}_T by $\frac{\mathrm{d}Q_{\phi}}{\mathrm{d}P^{\mathrm{Bri}}}\big|_{\mathcal{F}_T}=Z_T\cdot\frac{q_T(\mathbf{U}_T)}{p_T(\mathbf{U}_T)}$, i.e. we also correct the endpoint density (boundary term) so that the terminal marginal under Q_{ϕ} is q_T . Girsanov's theorem yields that under Q_{ϕ} ,

$$\mathbf{W}_t^{(Q)} := \mathbf{W}_t^{(P)} - \int_0^t \vartheta_s \, \mathrm{d}s \tag{83}$$

is a Brownian motion and the drift becomes $b_Q = f - g^2 s_{\text{cond}}$, i.e. the reverse SDE for Q_{ϕ} .




Figure 4: Comparison of DDVI and DBVI on the ENERGY dataset (Test RMSE).

Now compute the log Radon–Nikodym derivative and take Q_ϕ -expectation:

$$KL(Q_{\phi} \parallel P^{Bri}) = \mathbb{E}_{Q_{\phi}} \left[\log \frac{dQ_{\phi}}{dP^{Bri}} \right]$$
(84)

$$= \mathbb{E}_{Q_{\phi}} \left[\int_{0}^{T} \vartheta_{t}^{\top} d\mathbf{W}_{t}^{(P)} - \frac{1}{2} \int_{0}^{T} \|\vartheta_{t}\|^{2} dt \right] + \mathbb{E}_{Q_{\phi}} \left[\log \frac{q_{T}(\mathbf{U}_{T})}{p_{T}(\mathbf{U}_{T})} \right]. \tag{85}$$

Use $d\mathbf{W}_t^{(P)} = d\mathbf{W}_t^{(Q)} + \vartheta_t dt$ to rewrite the stochastic integral:

$$\int_0^T \vartheta_t^\top d\mathbf{W}_t^{(P)} = \int_0^T \vartheta_t^\top d\mathbf{W}_t^{(Q)} + \int_0^T \|\vartheta_t\|^2 dt.$$
 (86)

Taking Q_{ϕ} -expectation annihilates the martingale term, hence

$$KL(Q_{\phi} \parallel P^{Bri}) = \frac{1}{2} \mathbb{E}_{Q_{\phi}} \left[\int_{0}^{T} \|\vartheta_{t}\|^{2} dt \right] + \mathbb{E}_{Q_{\phi}} \left[\log \frac{q_{T}(\mathbf{U}_{T})}{p_{T}(\mathbf{U}_{T})} \right].$$
 (87)

Finally substitute $\vartheta_t = -g(t) \left(s_{\text{cond}}(t, \mathbf{U}_t, \mathbf{U}_0) - s_{\text{Bri}}(t, \mathbf{U}_t) \right)$ to get

$$\frac{1}{2} \int_0^T g(t)^2 \mathbb{E}_{Q_{\phi}} \left[\left\| s_{\text{cond}}(t, \mathbf{U}_t, \mathbf{U}_0) - s_{\text{Bri}}(t, \mathbf{U}_t) \right\|_2^2 \right] dt$$
 (88)

as the path term, plus the boundary correction.

Ablation study. We directly compare DDVI and DBVI under matched compute (same T, optimizer, and hardware) and track Test RMSE during early optimization. As shown in 4, DBVI reduces Test RMSE more rapidly and consistently achieves lower error than DDVI. This confirms that conditioning the reverse diffusion on a data-dependent initialization shortens the inference trajectory and improves convergence speed on the ENERGY, CONCRETE and POWER dataset.

B STATEMENT ON THE USE OF LARGE LANGUAGE MODELS

Large language models (LLMs) were used solely for polishing and editing the text of this manuscript.




Figure 5: Comparison of DDVI and DBVI on the CONCRETE dataset (Test RMSE).




Figure 6: Comparison of DDVI and DBVI on the POWER dataset (Test RMSE).