A SIMULATION-FREE DIFFUSION-BASED METHOD FOR SOLVING NONLINEAR PDES

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep learning methods based on backward stochastic differential equations (BSDEs) provide an effective approach to solving nonlinear partial differential equations (PDEs), particularly in high dimensions. By exploiting stochastic representations of PDE solutions, these methods encode second-order differential operators through the dynamics of diffusion processes, avoiding the explicit computation of second-order derivatives. Their application, however, typically requires sequential path simulation and, on bounded domains, accounting for random exit times, which introduces additional numerical error.

A key observation is that the stochastic representation extends beyond the Markovian setting, as Itô's formula remains valid for general continuous semimartingales. We exploit this observation using Brownian bridges, whose endpoints can be prescribed and whose samples can be generated in closed form. Choosing the terminal endpoint on the domain boundary eliminates the need to account for random exit times, while closed-form sampling enables fully parallel generation of the stochastic inputs. This yields a diffusion-based deep-learning method for nonlinear PDEs that avoids sequential path simulation and random exit times, while retaining the advantages of stochastic representations in high dimensions.

1 INTRODUCTION

Partial differential equations (PDEs) are ubiquitous in science and engineering (Evans, 2010). However, classical grid-based discretization methods can become prohibitively expensive in high dimensions (Bellman, 1957) and may be difficult to apply on complex domains (Sawhney & Crane, 2020). Recently, the combination of Monte Carlo methods and deep learning has given rise to a range of alternative approaches designed to address these limitations, most prominently driven by methods based on physics-informed neural networks (PINNs) (Raissi et al., 2019) or backward stochastic differential equations (BSDEs) (Han et al., 2018). These approaches, however, introduce challenges of their own, including the cost and the discretization error of simulating trajectories (Park & Tu, 2025; Seo et al., 2026), slow convergence (Rathore et al., 2024), and difficulties in accurately handling complex geometries (Lee et al., 2026). In this paper, we develop methods aimed at addressing these limitations.

We consider nonlinear boundary value problems on a spatial domain $\Omega \subset \mathbb{R}^d$ and time interval $[0, T]$ of the form

$$(\partial_t + L)V(x,t) + h\big(x,t,V(x,t),\sigma^\top \nabla V(x,t)\big) = 0, \qquad (x,t) \in \Omega \times [0,T), \tag{1a}$$

$$V(x,T) = f(x), \qquad x \in \Omega, \tag{1b}$$

$$V(x,t) = g(x,t), \qquad (x,t) \in \partial\Omega \times [0,T]. \tag{1c}$$

Here, $h \in C(\mathbb{R}^d \times [0,T] \times \mathbb{R} \times \mathbb{R}^d, \mathbb{R})$ specifies the nonlinearity, while $f \in C(\mathbb{R}^d, \mathbb{R})$ and $g \in C(\mathbb{R}^d \times [0,T], \mathbb{R})$ define the terminal and boundary conditions, respectively. The operator L is given by

$$L = \frac{1}{2} \sum_{i,j=1}^{d} (\sigma\sigma^\top)_{ij}(x,t)\, \partial_{x_i}\partial_{x_j} + \sum_{i=1}^{d} b_i(x,t)\, \partial_{x_i}, \tag{2}$$

where $b \in C(\mathbb{R}^d \times [0,T], \mathbb{R}^d)$ and $\sigma \in C(\mathbb{R}^d \times [0,T], \mathbb{R}^{d\times d})$.

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.

The elliptic case is recovered by omitting the time derivative and terminal condition, yielding

$$LV + h\big(\cdot, V, \sigma^\top \nabla V\big) = 0, \qquad V|_{\partial\Omega} = g.$$

Throughout we work under standing regularity assumptions on the domain, the coefficients b, σ and the data f, g, h, which are restated as Assumptions B.1 and B.2 in App. B.1, following Nüsken & Richter (2023, Assumption 1 and Proposition 3.4).

For an approximation $\varphi \approx V$, we can define the pointwise PDE residual by

$$\rho_\varphi(x,t) := (\partial_t + L)\varphi(x,t) + h\big(x, t, \varphi(x,t), \sigma^\top \nabla\varphi(x,t)\big). \tag{3}$$

A standard PINN objective combines this interior residual with penalties enforcing the terminal and boundary conditions (1b)-(1c),

$$\mathcal{L}_{\mathrm{PINN}}(\varphi) = \alpha_{\mathrm{int}}\mathcal{L}_{\mathrm{int}}^{\mathrm{PINN}}(\varphi) + \alpha_T\mathcal{L}_T(\varphi) + \alpha_b\mathcal{L}_b(\varphi), \tag{4}$$

with weights $\alpha_{\mathrm{int}}, \alpha_T, \alpha_b > 0$ and

$$\mathcal{L}_{\mathrm{int}}^{\mathrm{PINN}}(\varphi) := \mathbb{E}[\rho_\varphi(X,t)^2],\ \mathcal{L}_T(\varphi) := \mathbb{E}[(\varphi(X_T,T) - f(X_T))^2],\ \mathcal{L}_b(\varphi) := \mathbb{E}[(\varphi(X_b,t_b) - g(X_b,t_b))^2], \tag{5}$$

where $(X,t) \sim \nu_{\Omega\times[0,T]}$, $X_T \sim \nu_\Omega$ and $(X_b,t_b) \sim \nu_{\partial\Omega\times[0,T]}$ are sampled from probability distributions with full support on $\Omega \times [0,T]$, Ω and $\partial\Omega \times [0,T]$, respectively.

An alternative class of objectives is motivated by observing that L defined in (2) is the infinitesimal generator of the stochastic differential equation

$$\mathrm{d}X_s = b(X_s, s)\,\mathrm{d}s + \sigma(X_s, s)\,\mathrm{d}W_s, \tag{6}$$

where W is classical Brownian motion. Applying Itô's formula to $\varphi(X_s, s)$ along (6) transforms the residual into an integral evaluated along the random trajectory of X. More precisely, over a time window $[t_0, \mathcal{T}]$,

$$\begin{aligned}\int_{t_0}^{\mathcal{T}} \rho_\varphi(X_s, s)\,\mathrm{d}s = \varphi(X_{\mathcal{T}}, \mathcal{T}) - \varphi(X_{t_0}, t_0) - \int_{t_0}^{\mathcal{T}} \sigma^\top \nabla\varphi(X_s, s)\cdot \mathrm{d}W_s \\ + \int_{t_0}^{\mathcal{T}} h\big(X_s, s, \varphi(X_s,s), \sigma^\top\nabla\varphi(X_s,s)\big)\,\mathrm{d}s.\end{aligned} \tag{7}$$

This identity underlies both the BSDE-based loss of E et al. (2017) and the diffusion loss of Nüsken & Richter (2023). It replaces the pointwise PDE residual by a stochastic process formulation and, in particular, eliminates the need to evaluate second-order derivatives explicitly.

This reformulation, however, ties the loss evaluation to the simulation of trajectories of (6), which entails two computational difficulties. First, simulating the SDE is inherently iterative: using, for example, the Euler-Maruyama scheme with time step Δt, evaluating the loss over a time window of length ℓ requires $O(\ell/\Delta t)$ simulation steps. Second, on bounded domains, the loss must account for the occupation of the process and, in the BSDE setting, for the exit time from Ω. Both quantities are generally accessible only through an approximate discretization of the simulated trajectory.

The main observation of this paper is that neither an iterative simulation of the SDE nor a discretized approximation of the trajectory is intrinsic to the underlying argument. The stochastic representation above depends on the process through its quadratic variation, whereas its drift can be treated separately. This makes it possible to replace simulated trajectories by sampling schemes with tractable, closed-form finite-dimensional distributions and the appropriate quadratic variation. In particular, we may exploit exact Brownian bridges, which allow boundary interactions to be resolved analytically while avoiding both iterative path simulation and the evaluation of second-order derivatives.

1.1 Related work

Stochastic representations and BSDE losses. Backward stochastic differential equations (Pardoux & Peng, 1990) give a probabilistic representation of semilinear elliptic and parabolic problems (Pardoux, 1998). Their discrete-time approximation (Bouchard & Touzi, 2004), and its extension to bounded domains, where an exit time has to be resolved (Bouchard & Menozzi, 2009), is the setting our discretization analysis inherits. The same representation is also solved without learning, by multilevel Picard iteration (Hutzenthaler et al., 2020). The deep BSDE loss (E et al., 2017; Han

et al., 2018) turned the representation into a learning objective, later extended to fully nonlinear equations (Beck et al., 2019), while Nüsken & Richter (2021) derive further losses from controlled diffusions and measures on path space. Nüsken & Richter (2023) interpolate between BSDEs and PINNs, and our construction generalizes their diffusion loss to non-Markovian sampling processes. More recent work targets the integration scheme (Park & Tu, 2025) and removes discretization bias and second-order derivatives from training (Seo et al., 2026).

Residual-based losses. Physics-informed neural networks (Raissi et al., 2019) and the deep Galerkin method (Sirignano & Spiliopoulos, 2018) minimize the pointwise PDE residual directly; the deep Ritz method (E & Yu, 2018) instead minimizes a variational energy, which presupposes that the problem admits such a formulation. Their training is delicate: gradient-flow pathologies (Wang et al., 2021), ill-conditioning induced by the differential operator (Rathore et al., 2024) and a neural-tangent-kernel analysis of training failure (Wang et al., 2022) all bear on slow convergence, and the weighting of interior against boundary terms is itself a design problem (van der Meer et al., 2022), which we study empirically through α_b.

Grid-free Monte Carlo solvers. A second line avoids a mesh altogether: walk-on-spheres (Muller, 1956), with modern extensions to complex geometries (Sawhney & Crane, 2020) and neural variants that amortize the walk (Nam et al., 2024). The bridge construction contains it as a special case, which Remark 2.9 makes precise.

Simulating conditioned diffusions. Our sampling process is a conditioned diffusion. Such bridges can be simulated by conditioning (Delyon & Hu, 2006) or, when no closed form is available, by guided proposals, see Schauer et al. (2017).

2 Simulation-free diffusion loss

In this section, we introduce a diffusion-based loss that avoids the evaluation of second-order derivatives while retaining the possibility of simulation-free computation. To this end, we first introduce a suitable class of sampling processes.

Definition 2.1 (Admissible sampling process)**.** Let $\mathbb{Q}$ be a probability measure on $C([0,T],\mathbb{R}^d)$ under which the canonical process $Z=(Z_s)_{s\in[0,T]}$ is a continuous semimartingale satisfying

$$\mathrm{d}\langle Z^i, Z^j\rangle_s = (\sigma\sigma^\top)_{ij}(Z_s,s)\,\mathrm{d}s, \qquad i,j=1,\dots,d, \tag{8}$$

and such that $(Z_{t_0},t_0)\sim\nu_{\Omega\times[0,T]}$, where $\nu_{\Omega\times[0,T]}$ has full support. We call $\mathbb{Q}$ *admissible*, and refer to Z as an *admissible sampling process*.

Note that no restriction is imposed on the finite-variation part of Z. In particular, Z need not be Markovian, and its drift may depend on the entire path, including its terminal value.

Proposition 2.2 (Generalized diffusion residual)**.** *Let $\varphi\in\mathcal{F}$ (the admissible class of App. B.1), let $\mathbb{Q}$ be admissible, and let $\mathcal{T}\in[t_0,T]$ be a stopping time such that $Z_s\in\Omega$ for $s\in[t_0,\mathcal{T}]$. Define*

$$R^{\mathcal{T}}_{t_0}(\varphi;Z) := \varphi(Z_{\mathcal{T}},\mathcal{T}) - \varphi(Z_{t_0},t_0) - \int_{t_0}^{\mathcal{T}} \nabla\varphi(Z_s,s)\cdot \mathrm{d}Z_s \\ + \int_{t_0}^{\mathcal{T}} \Big(b\cdot\nabla\varphi + h\big(\cdot,\cdot,\varphi,\sigma^\top\nabla\varphi\big)\Big)(Z_s,s)\,\mathrm{d}s. \tag{9}$$

Then, $\mathbb{Q}$-almost surely, $R^{\mathcal{T}}_{t_0}(\varphi;Z)=\int_{t_0}^{\mathcal{T}}\rho_\varphi(Z_s,s)\,\mathrm{d}s$.

The proof, given in App. B, is Itô's formula for continuous semimartingales together with (8).

For any admissible sampling measure $\mathbb{Q}$, we define the loss

$$\mathcal{L}^{\mathbb{Q}}(\varphi) = \alpha_{\mathrm{int}}\mathcal{L}^{\mathbb{Q}}_{\mathrm{int}}(\varphi) + \alpha_T\mathcal{L}_T(\varphi) + \alpha_b\mathcal{L}_b(\varphi), \tag{10}$$

where

$$\mathcal{L}^{\mathbb{Q}}_{\mathrm{int}}(\varphi) := \mathbb{E}_{\mathbb{Q}}\left[\big(R^{\mathcal{T}}_{t_0}(\varphi;Z)\big)^2\right] \tag{11}$$

is the interior loss, while $\mathcal{L}_T(\varphi)$ and $\mathcal{L}_b(\varphi)$ are the losses associated with the terminal and boundary conditions defined in (5). The terms are weighted by positive constants $\alpha_{\mathrm{int}},\alpha_T,\alpha_b>0$.

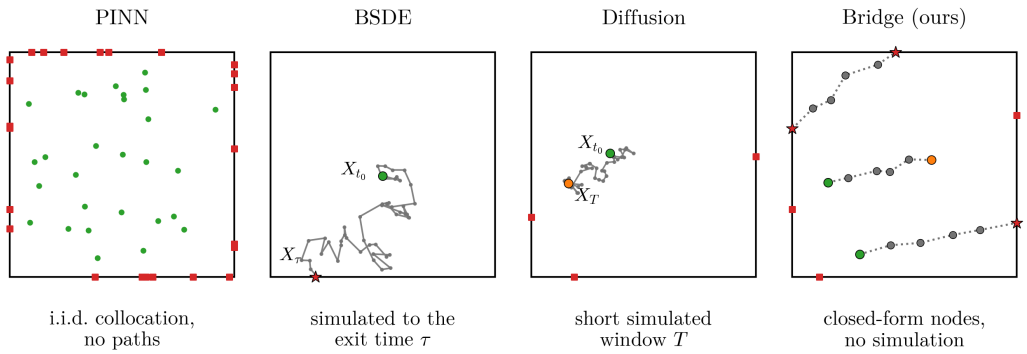

Figure 1: Illustration of the training data used by each loss. Green: interior point; orange: interior endpoint; red: point on $\partial\Omega$ (star: path endpoint, square: independently sampled). Grey: sampled path, simulated (solid) or closed-form bridge (dotted). The last panel shows the three bridge variants from Section 2.2, distinguished by their endpoints.

Remark 2.3 (Recovery of known losses). For $Z = X$ solving the SDE (6), $\mathcal{L}_{\text{int}}^{\mathbb{Q}}$ is exactly the diffusion loss of Nüsken & Richter (2023). When additionally considering the limit $\mathcal{T} \to \tau \wedge T$ with $\tau = \inf\{s : X_s \notin \Omega\}$, one recovers the BSDE loss originally suggested by E et al. (2017), which rests on the BSDE representation, see, e.g., Pardoux (1998).

Proposition 2.4 (Validity). *Let $\varphi \in \mathcal{F}$ and let $\mathbb{Q}$ be admissible. Then φ solves (1a)-(1c) if and only if $\mathcal{L}^{\mathbb{Q}}(\varphi) = 0$. Furthermore, in the short-window limit $\ell \to 0$, writing $\mathcal{T} = (t_0 + \ell) \wedge \tau \wedge T$, it holds*

$$\frac{\mathcal{L}_{\text{int}}^{\mathbb{Q}}(\varphi)}{\ell^2} \longrightarrow \mathcal{L}_{\text{int}}^{\text{PINN}}(\varphi). \tag{12}$$

The proof, given in App. B, is elementary and needs no Markov property. This will matter for the sampling process considered in Section 2.2, which will have a drift that blows up as the terminal time is approached, so the well-posedness argument used for the diffusion loss in Nüsken & Richter (2023, Prop. 3.4) does not apply to them directly.

Remark 2.5 (Vanishing variance at the solution). Let φ solve (1a)-(1c), let $\mathbb{Q}$ be an admissible sampling process. Then $R_{t_0}^{\mathcal{T}}(\varphi; Z) = 0$ $\mathbb{Q}$-almost surely. Hence, for an i.i.d. sample $Z^{(1)}, \dots, Z^{(K)} \sim \mathbb{Q}$, both the Monte Carlo loss $\frac{1}{K}\sum_{k=1}^{K}(R^{(k)})^2$ and its gradient have zero variance, cf. Nüsken & Richter (2021).

In practice, the quantity in (9) must be discretized along the sampling process Z on a time grid $t_0 = s_0 < \cdots < s_n = \mathcal{T}$, and therefore requires the states $Z_{s_0}, \dots, Z_{s_n}$. We call the resulting loss *simulation-free* if the following conditions are satisfied:

(i) $\mathbb{Q}$ satisfies (8);

(ii) the finite-dimensional distributions of $(Z_{s_0}, \dots, Z_{s_n})$ are available in closed form and can be sampled from directly, that is, without iterating a time-stepping scheme;

(iii) $\operatorname{supp}\mathbb{Q}$ contains the region on which $\rho_\varphi \equiv 0$ is to be enforced.

Condition (i) is the restrictive requirement. For state-dependent σ, it essentially requires Z to be obtained from an SDE and therefore precludes the closed-form sampling considered here. We consequently assume throughout the remainder of this section that $\sigma(x,t) = \sigma(t)$, i.e. σ is independent of the state variable. The drift b, on the other hand, may remain arbitrary and nonlinear: it does not enter the quadratic variation condition (8) and appears in (9) only through the pointwise term $b \cdot \nabla\varphi$.

We next introduce two admissible sampling processes satisfying these conditions.

2.1 Brownian motion

The simplest admissible choice is the law $\mathbb{Q}^0$ of the driftless sampling process

$$\mathrm{d}Z_s = \sigma(s)\,\mathrm{d}W_s, \qquad (Z_{t_0}, t_0) \sim \nu_{\Omega\times[0,T]}, \tag{13}$$

whose transition distributions are exactly Gaussian:

$$Z_{s_{j+1}} \mid Z_{s_j} \sim \mathcal{N}\big(Z_{s_j}, \Sigma(s_{j+1}) - \Sigma(s_j)\big), \tag{14}$$

where $\Sigma(t) := \int_0^t \sigma(s)\sigma(s)^\top \,\mathrm{d}s$. Consequently, the states can be sampled exactly, independently of the discretization step size. In particular, no Euler-Maruyama discretization error is incurred. For this choice of sampling process, (9) becomes

$$R_{t_0}^{\mathcal{T}}(\varphi; Z) = \varphi(Z_{\mathcal{T}}, \mathcal{T}) - \varphi(Z_{t_0}, t_0) - \int_{t_0}^{\mathcal{T}} \big(\sigma^\top \nabla\varphi\big)(Z_s, s) \cdot \mathrm{d}W_s + \int_{t_0}^{\mathcal{T}} \Big(b \cdot \nabla\varphi + h\big(\cdot, \cdot, \varphi, \sigma^\top \nabla\varphi\big)\Big)(Z_s, s)\,\mathrm{d}s. \quad (15)$$

Thus, this construction yields the diffusion loss of Nüsken & Richter (2023), augmented by the explicit term $b \cdot \nabla\varphi$. This additional term accounts for the fact that the sampling process itself does not carry the drift b. To our knowledge, existing implementations of diffusion-based and BSDE-based losses simulate a version of the SDE (6), typically by Euler-Maruyama, even when σ is independent of x. In the latter case, however, the driftless process (13) has closed-form finite-dimensional distributions, so its states can be sampled exactly and no iterative SDE simulation is required.

2.2 Brownian bridges

More informative than sampling an unconditioned endpoint is to condition the process on both of its endpoints. For $t_0 < r$, conditioning the driftless process (13) on $Z_{t_0} = z_0$ and $Z_r = z_r$ yields

$$\mathrm{d}Z_s = \sigma(s)\sigma(s)^\top \big(\Sigma(r) - \Sigma(s)\big)^{-1}(z_r - Z_s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad Z_{t_0} = z_0 \quad (16)$$

via Doob's h-transform on the interval $[t_0, r]$; the derivation is given in App. B.3. Only the drift has changed: the diffusion coefficient in (16) is still $\sigma(s)$, so that $\mathrm{d}\langle Z^i, Z^j\rangle_s = (\sigma\sigma^\top)_{ij}(s)\,\mathrm{d}s$, which is exactly (8). The bridge is therefore admissible for any choice of anchors (z_0, z_r) at times (t_0, r), irrespective of the drift b of the underlying dynamics.

Conditioning on a terminal value preserves the Markov property, so the transition distributions are again Gaussian and explicit (App. B.3), now conditioned on the anchor z_r as well as on the current state. The drift of (16) therefore anticipates z_r:

$$Z_{s_{j+1}} \mid (Z_{s_j}, z_r) \sim \mathcal{N}\Big(Z_{s_j} + G_j(z_r - Z_{s_j}),\, (I - G_j)\big(\Sigma(s_{j+1}) - \Sigma(s_j)\big)\Big), \quad (17)$$

where we introduce the shorthand notation

$$G_j := \big(\Sigma(s_{j+1}) - \Sigma(s_j)\big)\big(\Sigma(r) - \Sigma(s_j)\big)^{-1}. \quad (18)$$

The states can therefore be sampled exactly, with no discretization error. Unlike the general stopping time $\mathcal{T}$ in Proposition 2.2, the time r is a fixed, prescribed horizon, specified in advance for (17) to arise from ordinary Gaussian conditioning. We therefore instantiate the residual (9) with $\mathcal{T} = r$. In practice, we consider three variants of Brownian bridges (see Fig. 1 for an illustration):

- **Interior-interior.** Both anchors $z_0, z_r \in \Omega$ are drawn from an interior sampling law. This is the direct analogue of the driftless sampler of Section 2.1, but with a prescribed rather than a free endpoint.
- **Interior-boundary (exit bridge).** Choosing $z_r \in \partial\Omega$ places the terminal point of the bridge exactly on the boundary. In this sense, the terminal boundary evaluation required by the BSDE loss is resolved without approximating a hitting time, and $g(z_r, r)$ can be inserted directly in place of $\varphi(Z_{\mathcal{T}}, \mathcal{T})$ in (9).
- **Boundary-boundary.** Drawing both anchors $z_0, z_r \in \partial\Omega$ additionally replaces $\varphi(Z_{s_0}, s_0)$ by the known boundary value $g(z_0, t_0)$. The residual is therefore supervised at both endpoints. This term is used in addition to, rather than as a replacement for, the interior-boundary loss.

In the latter two variants the endpoint anchor is treated as the exit point, which is exact only if the path does not leave Ω before r and return; App. C.3 discusses this effect and how it is controlled. More generally, any $m \geq 1$ such bridges can be *chained* through deterministic anchors, e.g.

$$y_j = x_0 + \frac{j}{m}(z_r - x_0), \qquad j = 1, \ldots, m-1. \quad (19)$$

Each individual bridge satisfies (8), and concatenation preserves the quadratic variation. Proposition 2.2 therefore applies unchanged to the resulting path. The window length $\ell = r - t_0$ and the placement of the anchors are free design choices. In particular, the short-window limit of the latter two variants isolates the local properties that these losses test.

Remark 2.6 (Endpoint-conditioned bridges on unbounded domains). On an unbounded domain there is no boundary to anchor the terminal point to, so the only available variant is interior-interior, and the endpoint law is then typically taken to be that of the process itself. Let us fix $t_0 < s_1 < \cdots < s_n < r$. By construction, (17) is the $\mathbb{Q}^0$-conditional law of $(Z_{s_1}, \ldots, Z_{s_n})$ given Z_r. Hence drawing z_r from Z_r's own $\mathbb{Q}^0$-marginal (13) and then the interior nodes from (17) conditioned on that z_r reproduces $\mathrm{Law}_{\mathbb{Q}^0}(Z_{s_1}, \ldots, Z_{s_n})$ exactly, by the tower property

$$\mathrm{Law}_{\mathbb{Q}^0}(Z_{s_1}, \ldots, Z_{s_n}) = \int_{\mathbb{R}^d} \mathrm{Law}_{\mathbb{Q}^0}\big((Z_{s_1}, \ldots, Z_{s_n}) \mid Z_r = z\big)\, \mathrm{Law}_{\mathbb{Q}^0}(Z_r \in \mathrm{d}z). \tag{20}$$

Proposition 2.7 (Short-window limits of the exit- and boundary-anchored residuals)**.** *Let $\psi := \varphi - V$, let Z be the bridge of (17) with anchors (z_0, z_r) and $\mathcal{T} = r$, $\ell = r - t_0$, and write R^{exit}, R^{bb} for (9) with $\varphi(Z_{\mathcal{T}}, \mathcal{T})$, resp. additionally $\varphi(Z_{t_0}, t_0)$, replaced by the known boundary value(s). Then, as $\ell \to 0$,*

(i) if $z_r \in \partial\Omega$ and ρ_φ and ψ are bounded on $\overline{\Omega} \times [0, T]$, then for any *anchor law*

$$\mathbb{E}\big[(R^{\mathrm{exit}})^2\big] = \mathbb{E}\big[\psi(z_r, r)^2\big] + O(\ell); \tag{21}$$

(ii) if additionally $z_0 \in \partial\Omega$ with $\|z_r - z_0\| = O(\sqrt{\ell})$, then

$$\mathbb{E}\big[(R^{\mathrm{bb}})^2\big] = \mathbb{E}\Big[\big(\nabla\psi(z_0, t_0) \cdot (z_r - z_0)\big)^2\Big] + O(\ell^{3/2}). \tag{22}$$

In particular, if the anchor pair satisfies $\mathbb{E}\big[(z_r - z_0)(z_r - z_0)^\top \mid z_0\big] = \ell\,(\sigma\sigma^\top)(z_0, t_0) + o(\ell)$, then $\ell^{-1}\mathbb{E}[(R^{\mathrm{bb}})^2] \to \mathbb{E}\big[\|\sigma^\top \nabla\psi(Z_{t_0}, t_0)\|^2\big]$.

Proof. See App. B. □

Remark 2.8 (Limiting interpretation of the anchored residuals). Both limits collapse onto the same object, the boundary mismatch $\varphi - g$ on $\partial\Omega$, in two different norms. In (i) it is measured in L^2. In (ii), the chord $z_r - z_0$ becomes tangential as $z_0 \to z_r$, while $\psi|_{\partial\Omega} = \varphi - g$, so the limit is a tangential H^1-seminorm of the same mismatch: stronger than a boundary MSE, but carrying no normal-derivative, and hence no PDE, information. Together with Proposition 2.4, this orders the terms,

$$\underbrace{\mathbb{E}[(R^{\mathrm{exit}})^2]}_{O(1)} \gg \underbrace{\mathbb{E}[(R^{\mathrm{bb}})^2]}_{O(\ell)} \gg \underbrace{\mathcal{L}^{\mathbb{Q}}_{\mathrm{int}}(\varphi)}_{O(\ell^2),\ \to \ell^2 \mathcal{L}^{\mathrm{PINN}}_{\mathrm{int}}}. \tag{23}$$

Thus, as $\ell \to 0$, the boundary-fitting terms dominate the PDE residual. The value of the exit anchor is therefore not that it fits the boundary better, but that it is the only construction in which the boundary datum enters the interior term at all, keeping the objective identifiable as $\alpha_b \to 0$. Note also that the isotropic form $\mathbb{E}[\|\sigma^\top \nabla\psi\|^2]$ is not reachable by any implementable sampler: (ii) needs both a full-rank anchor covariance and known values at both anchors, while known values are available only on $\partial\Omega$, where the covariance degenerates, or at $s = T$, where $\ell = 0$. This is one reason the boundary-boundary term is used in addition to, rather than in place of, the interior-boundary one.

Remark 2.9 (Relation to walk-on-spheres). Our bridge framework contains the so-called *walk-on-spheres* algorithm as a special case (Muller, 1956; Sawhney & Crane, 2020). Consider the elliptic case with $h = b = 0$ and $\sigma\sigma^\top = \sigma^2 I$ constant. Choose $z_0 \in \Omega$, let $B = B(z_0, r_0) \subset \Omega$ be the largest inscribed ball, with $r_0 = \mathrm{dist}(z_0, \partial\Omega)$, and take $\mathbb{Q} = \mathbb{P}$, the law of the forward process (6), with $\mathcal{T} = \tau_B$. Then (9) reduces to $R_0^{\tau_B}(\varphi; X) = \varphi(X_{\tau_B}) - \varphi(z_0) - \int_0^{\tau_B} \sigma^\top \nabla\varphi(X_s) \cdot \mathrm{d}W_s$, whose last term is a martingale. Thus, with $\mu_{z_0} = \mathrm{Unif}(\partial B)$,

$$\int_{\partial B} \varphi(y)\, \mathrm{d}\mu_{z_0}(y) - \varphi(z_0) = \mathbb{E}\big[R_0^{\tau_B}(\varphi; X)\big] = \frac{\sigma^2}{2}\, \mathbb{E}\int_0^{\tau_B} \Delta\varphi(X_s)\, \mathrm{d}s. \tag{24}$$

Requiring (24) to vanish gives the mean-value property, which walk-on-spheres applies recursively: μ_{z_0} is explicit and φ enters only through the point evaluation $\varphi(X_{\tau_B})$. Iterating $z_0 \mapsto X_{\tau_B}$ reaches an ε-shell of $\partial\Omega$ after $O(\log \varepsilon^{-1})$ hops and, up to the $O(\varepsilon)$ bias incurred by stopping there, estimates $V(z_0)$ from independent samples without a parametric φ or optimization. *Neural walk-on-spheres* (Nam et al., 2024) amortizes the same recursion by regressing a neural network onto it. Both constructions are limited to linear problems, however: once h depends nonlinearly on φ or $\nabla\varphi$, the recursion no longer closes.

For convenience, Algorithm 1 in App. C summarizes one gradient step, covering both the driftless sampler of Section 2.1 and the bridge of Section 2.2. Its memory footprint is $O(K(n+1)d)$ in the batch size K, the number of sub-steps n and the dimension d, with n a free parameter; $n = 1$ is admissible for the bridge. The cost compares favorably with both families the loss interpolates: the $n+1$ nodes are drawn in closed form rather than by sequential time stepping, so that no trajectory has to be integrated or stored, and only first derivatives of φ are required, avoiding the $O(d)$ second-derivative passes that $\Delta\varphi$ costs in the PINN residual.

3 Experiments

In this section, we illustrate the performance of the bridge loss on three numerical examples and compare it with classical BSDE losses, PINNs, and the diffusion loss. We consider elliptic problems with challenging boundary conditions or dynamics, as well as a high-dimensional parabolic problem. While PINN and BSDE losses can perform well on low-dimensional problems with simple geometries, we argue that bridge and diffusion losses may be better suited to certain challenging high-dimensional settings.

3.1 Challenging two-dimensional geometry: Perforated square

On the unit square perforated by N circular holes, $\Omega = (0,1)^2 \setminus \bigcup_{n=1}^{N} \overline{B(c_n, r)}$, we consider the nonlinear equation

$$\Delta V + \lambda\big(\|\nabla V\|^2 - q\big) = 0 \quad \text{in } \Omega, \qquad V = g \quad \text{on } \partial\Omega, \tag{25}$$

where $\lambda > 0$ and $g, q \in C(\Omega)$, together with further details, are specified in App. D.2 such that the problem admits a closed-form solution for every N. Figs. 2 and 3 summarize the results; every method is given the same wall-clock budget. We observe that the exit bridge degrades by less than a factor of 1.5 as N increases from 16 to 144, whereas the PINN error grows by a factor of four. Fig. 2 illustrates the underlying behavior: at $N = 64$, the PINN error is essentially a constant offset of each hole's level, differing from hole to hole, and is largest on $\partial\Omega$, even at its optimal boundary weight α_b. In contrast, the bridge error is smaller and approximately uniform with respect to the distance to the boundary. With $N + 4$ boundary components, each carrying its own level, the boundary datum varies on a scale that the network cannot easily resolve pointwise, and increasing α_b does not remedy this. The bridge, by contrast, does not fit g on $\partial\Omega$ in isolation: by (9), each interior anchor is tied directly to a boundary value $g(z_r)$, so the datum is carried into the interior by the loss rather than propagated there by the PDE residual alone. The loss of accuracy of soft boundary penalties as the number of components increases has also been reported for perforated domains by Lee et al. (2026). More generally, this behavior is consistent with the gradient imbalance between penalty and residual terms (Wang et al., 2021). Rathore et al. (2024) trace the ill-conditioning of the PINN loss landscape to the differential operator in the residual, which the bridge loss partly replaces by a stochastic integral.

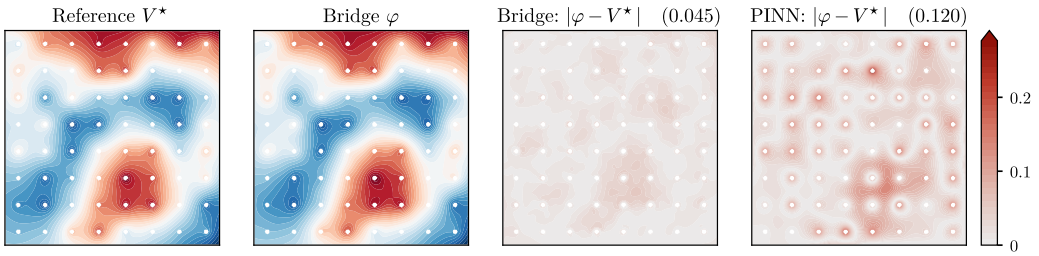

Figure 2: Perforated square with $N = 64$ holes. *Left:* exact solution and bridge approximation. *Right:* pointwise errors of the bridge and PINN methods at their respective optimal boundary weight α_b, again shown on a shared scale; relative L^2 errors are given in parentheses.

3.2 Challenging dynamics: Metastable chain committor

We consider a high-dimensional metastable potential obtained by placing M Gaussian wells along a two-dimensional plane. Writing $x = (y, z) \in \mathbb{R}^2 \times \mathbb{R}^{d-2}$ and letting $a_1, \dots, a_M \in \mathbb{R}^2$, we define

$$U(x) = -\log \sum_{m=1}^{M} w_m \exp\Big(-\frac{\|y - a_m\|^2 + \|z\|^2}{2s^2}\Big), \qquad \mathrm{d}X_t = -\nabla U(X_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{26}$$

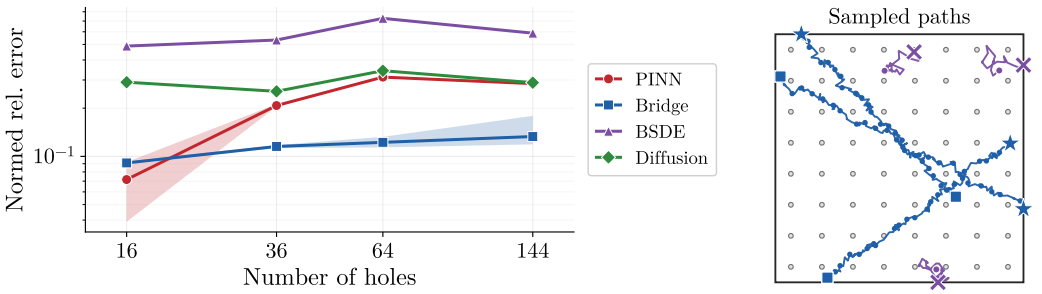

Figure 3: *Left:* converged relative L^2 error divided by that of the nearest-boundary predictor $g(\mathrm{proj}_{\partial\Omega}(x))$, so that 1 corresponds to no improvement over boundary interpolation. Each curve shows a method at its best boundary weight α_b, with the shaded band showing the minimum and maximum over three seeds (most runs agree to within 2%, so the band is thinner than the line). The full dependence on α_b is reported in Fig. 6 in App. D.3. *Right:* exit bridges drawn during training ($N = 64$), each chaining $m = 23$ Brownian bridges from an interior anchor x_0 (square) to an independently sampled $z_b \in \partial\Omega$ (star) through deterministic waypoints (dots), allowing the bridge to navigate around the obstacles. This contrasts with the BSDE method, whose trajectories (purple) hit the boundary immediately and therefore terminate quickly.

see App. D.3 for details. We take the two terminal wells to be absorbing,

$$A = \{x : \|y - a_1\| \le R\}, \qquad B = \{x : \|y - a_M\| \le R\}, \qquad \Omega = \mathbb{R}^d \setminus (A \cup B), \tag{27}$$

and let V denote the committor of B against A, i.e. the probability of reaching B before A,

$$V(x) = \mathbb{P}\big(\tau_B < \tau_A \mid X_0 = x\big), \qquad \tau_S = \inf\{t > 0 : X_t \in S\}. \tag{28}$$

The committor is a central object in transition path theory: it provides an optimal reaction coordinate between two metastable states, from which the transition rate, reaction current, and dominant pathways can be derived (E & Vanden-Eijnden, 2010; Schütte et al., 2023). It is routinely used in applications including protein folding, nucleation, and climate tipping events (Bolhuis et al., 2002; Jacques-Dumas et al., 2023), and in high dimensions it is commonly approximated by neural networks (Khoo et al., 2019; Li et al., 2019). The committor solves the elliptic boundary value problem

$$\tfrac{\sigma^2}{2}\Delta V - \nabla U \cdot \nabla V = 0 \ \text{ in } \Omega, \qquad V|_{\partial A} = 0, \qquad V|_{\partial B} = 1. \tag{29}$$

Fig. 4 shows the results: the bridge loss remains essentially flat in d, with a relative L^2 error of about 0.08 across $d \in \{2, 10, 50, 100\}$, whereas the PINN error grows by an order of magnitude, the diffusion loss is non-monotone in d, and the BSDE loss stays above 0.25. At this fixed wall-clock budget, the two baselines degrade for different reasons. The PINN residual requires $\Delta\varphi$, whose per-sample cost grows with d, and its accuracy depends strongly on α_b (Fig. 7 in App. D.3). The BSDE loss instead needs a realized exit, $Y_\tau = g(X_\tau)$, which on this potential takes $2.5 \cdot 10^4$ sequential steps per parameter update, so at equal wall-clock it makes two orders of magnitude fewer updates than the other losses.

3.3 High-dimensional parabolic problem: Hamilton-Jacobi-Bellman equation

While the preceding examples are elliptic problems on bounded domains, we now turn to a parabolic problem on an unbounded domain, namely the 100-dimensional Hamilton-Jacobi-Bellman equation

$$\partial_t V + \Delta V - \|\nabla V\|^2 = 0 \quad \text{on } \mathbb{R}^d \times [0, T), \qquad V(T, x) = \log\big(\tfrac{1}{2}(1 + \|x\|^2)\big). \tag{30}$$

As established in Remarks 2.3 and 2.6, on an unbounded domain the bridge is used in its interior-interior form with the endpoint drawn from its own marginal, which by (20) makes it coincide with the diffusion loss introduced by Nüsken & Richter (2023); we therefore report the two as a single curve. Crucially, however, we introduce a novel simulation-free formulation of this loss. We use this example to isolate the performance of the windowed residual (9), contrasting it with traditional BSDE-type losses that require evaluating entire trajectories. By operating on smaller subtrajectories instead, our approach allows finer time steps Δt, which substantially reduces discretization bias. To contextualize these improvements, we benchmark against recent techniques that similarly target discretization error, namely improved time-integration schemes (Park & Tu, 2025) and an unbiased, second-order-free estimator (Seo et al., 2026). We also include a PINN baseline that utilizes the exact Laplacian to evaluate the pointwise residual, thereby avoiding time-discretization error entirely. Full numerical details are provided in App. D.4, with the corresponding results illustrated in Fig. 5.

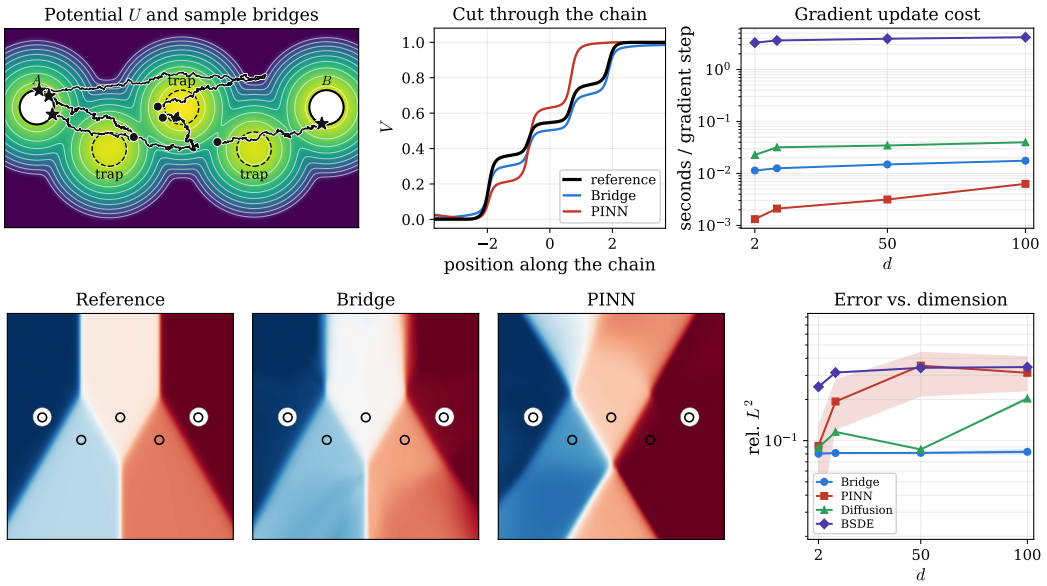

Figure 4: *Top left:* Potential U of (26) on the plane of the centers, with absorbing sets A, B and three traps, and four bridges, chained from an interior anchor (dot) through the intervening traps to the nearer absorbing set (star). *Top middle:* a cut through the chain, comparing against the reference solution. *Top right:* wall-clock cost of one gradient step against d. The PINN residual needs $\Delta\varphi$ and so grows fastest, while the BSDE loss sits two orders of magnitude above the rest because every step simulates each path all the way to $\partial\Omega$. *Bottom:* reference against the bridge and PINN approximations, each at its own best α_b; PINN places the risers but misses the plateau heights. At the right, relative L^2 error of each loss as a function of d at a fixed 3-hour budget, with bands spanning seeds, narrower than the line for bridge and BSDE loss. The bridge remains flat in d and consistent across seeds, whereas the baselines do not. The α_b sweeps are deferred to Fig. 7 in App. D.3.

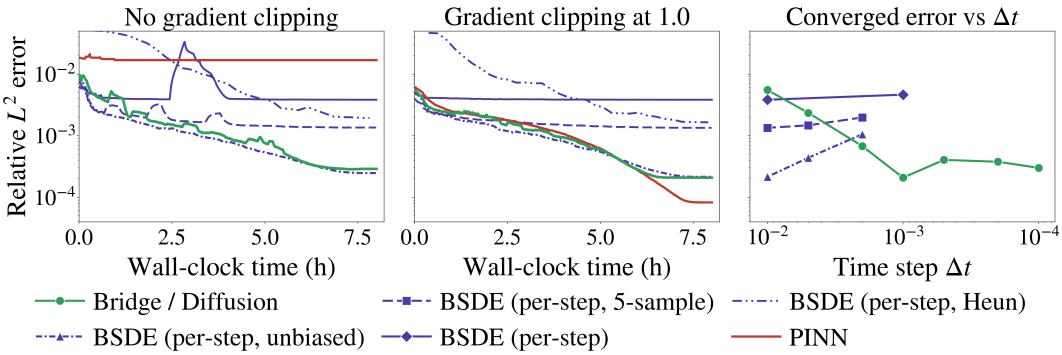

Figure 5: *Left and middle:* Bridge / diffusion works out of the box, whereas PINN requires substantial tuning, including gradient clipping and weight tuning. Bridge / diffusion performs comparably to the unbiased BSDE method of Seo et al. (2026). The Heun method of Park & Tu (2025) requires an eightfold smaller batch size to fit in GPU memory and is significantly slower. *Right:* Converged error as a function of the time step Δt at the fixed window length $\ell = 0.2$. Below $\Delta t = 10^{-3}$ the bias is exhausted: refining Δt shrinks the residual variance and the affordable updates in step, so the error flattens. The trajectory losses worsen instead: an update costs $T/\Delta t$ evaluations against the window's $\ell/\Delta t$, so refining Δt buys proportionally fewer updates.

4 Conclusion

We have generalized the diffusion loss to arbitrary admissible sampling processes, constrained only through their quadratic variation (8). This admits processes that are neither Markovian nor solutions of (6), and in particular Brownian bridges, whose states are available in closed form. Anchoring the terminal point on $\partial\Omega$ then replaces the random exit time by a prescribed horizon and lets the boundary datum enter the interior residual itself, which keeps the objective identifiable as $\alpha_b \to 0$. The resulting loss needs neither sequential path simulation nor second derivatives.

REPRODUCIBILITY STATEMENT

Assumptions and complete proofs are given in Apps. A and B. The loss as it is computed is stated in Algorithm 1 and its discretization analyzed in App. C. App. D reports the experimental protocol and, for each problem, the data, the construction of its reference solution and the settings of every loss compared on it; all three references are deterministic, so reproducing a reported number requires reproducing the solver alone. The implementation will be released with the camera-ready version.

USE OF GENERATIVE AI

In this work, we used generative AI tools for coding assistance in implementing the numerical experiments, and for searching the literature for related work. All generated code was reviewed and tested by the authors against the reference solutions described in App. D, and every reference was checked against its publication record before being cited. The authors take full responsibility for the contents of this paper.

REFERENCES

Christian Beck, Weinan E, and Arnulf Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29:1563–1619, 2019. doi: 10.1007/s00332-018-9525-3.

Richard E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

Peter G. Bolhuis, David Chandler, Christoph Dellago, and Phillip L. Geissler. Transition path sampling: throwing ropes over rough mountain passes, in the dark. *Annual Review of Physical Chemistry*, 53:291–318, 2002.

Bruno Bouchard and Stéphane Menozzi. Strong approximations of BSDEs in a domain. *Bernoulli*, 15(4):1117–1147, 2009. doi: 10.3150/08-BEJ181.

Bruno Bouchard and Nizar Touzi. Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations. *Stochastic Processes and their Applications*, 111(2): 175–206, 2004. doi: 10.1016/j.spa.2004.01.001.

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.

Weinan E and Eric Vanden-Eijnden. Transition-path theory and path-finding algorithms for the study of rare events. *Annual Review of Physical Chemistry*, 61:391–420, 2010.

Weinan E and Bing Yu. The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6:1–12, 2018. doi: 10.1007/s40304-018-0127-z.

Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017. doi: 10.1007/s40304-017-0117-6.

Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 2nd edition, 2010. doi: 10.1090/gsm/019.

Pat Fitzsimmons, Jim Pitman, and Marc Yor. Markovian bridges: construction, Palm interpretation, and splicing. In *Seminar on Stochastic Processes, 1992*, volume 33 of *Progress in Probability*, pp. 101–134. Birkhäuser Boston, Boston, MA, 1993. doi: 10.1007/978-1-4612-0339-1_5.

Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018. doi: 10.1073/pnas.1718942115.

Martin Hutzenthaler, Arnulf Jentzen, Thomas Kruse, Tuan Anh Nguyen, and Philippe von Wurstemberger. Overcoming the curse of dimensionality in the numerical approximation of semilinear parabolic partial differential equations. *Proceedings of the Royal Society A*, 476(2244):20190630, 2020. doi: 10.1098/rspa.2019.0630.

Jean Jacod and Philip Protter. *Discretization of Processes*. Stochastic Modelling and Applied Probability. Springer, Berlin, Heidelberg, 2012. doi: 10.1007/978-3-642-24127-7.

Valérian Jacques-Dumas, René M. van Westen, Freddy Bouchet, and Henk A. Dijkstra. Data-driven methods to estimate the committor function in conceptual ocean models. *Nonlinear Processes in Geophysics*, 30(2):195–216, 2023.

Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. Solving for high-dimensional committor functions using artificial neural networks. *Research in the Mathematical Sciences*, 6(1):1, 2019. doi: 10.1007/s40687-018-0160-2.

Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin, Heidelberg, 1992. doi: 10.1007/978-3-662-12616-5.

Jeeeun Lee, Denis Korolev, Miro Duhovic, and Seong Su Kim. Finite basis physics-informed neural networks with hard constraints for viscous fluid flow in highly perforated domains. *arXiv preprint arXiv:2608.08114*, 2026.

Qianxiao Li, Bo Lin, and Weiqing Ren. Computing committor functions for the study of rare events using deep learning. *The Journal of Chemical Physics*, 151(5):054112, 2019. doi: 10.1063/1.5110439.

Mervin E. Muller. Some continuous Monte Carlo methods for the Dirichlet problem. *Annals of Mathematical Statistics*, 27(3):569–589, 1956.

Hong Chul Nam, Julius Berner, and Anima Anandkumar. Solving Poisson equations using neural walk-on-spheres. In *International Conference on Machine Learning (ICML)*, 2024. arXiv:2406.03494.

Nikolas Nüsken and Lorenz Richter. Solving high-dimensional Hamilton-Jacobi-Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):48, 2021. doi: 10.1007/s42985-021-00102-x.

Nikolas Nüsken and Lorenz Richter. Interpolating between BSDEs and PINNs: deep learning for elliptic and parabolic boundary value problems. *Journal of Machine Learning*, 2(1):31–64, 2023. arXiv:2112.03749.

Étienne Pardoux. Backward stochastic differential equations and viscosity solutions of systems of semilinear parabolic and elliptic PDEs of second order. In *Stochastic Analysis and Related Topics VI*, volume 42 of *Progress in Probability*, pp. 79–127. Birkhäuser Boston, Boston, MA, 1998.

Étienne Pardoux and Shige Peng. Adapted solution of a backward stochastic differential equation. *Systems & Control Letters*, 14(1):55–61, 1990. doi: 10.1016/0167-6911(90)90082-6.

Sungje Park and Stephen Tu. Integration matters for learning PDEs with backward SDEs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, and Madeleine Udell. Challenges in training PINNs: A loss landscape perspective. In *International Conference on Machine Learning (ICML)*, 2024.

L. C. G. Rogers and David Williams. *Diffusions, Markov Processes and Martingales. Volume 2: Itô Calculus*. Cambridge University Press, Cambridge, 2 edition, 2000. doi: 10.1017/CBO9780511805141.

Rohan Sawhney and Keenan Crane. Monte Carlo geometry processing: a grid-free approach to PDE-based methods on volumetric domains. *ACM Transactions on Graphics*, 39(4), 2020.

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.

Christof Schütte, Stefan Klus, and Carsten Hartmann. Overcoming the timescale barrier in molecular dynamics: transfer operators, variational principles and machine learning. *Acta Numerica*, 32: 517–673, 2023.

Jaemin Seo, Surin Lee, and Jae Yong Lee. Unbiased and second-order-free training for high-dimensional PDEs. In *International Conference on Machine Learning (ICML)*, 2026. arXiv:2605.14643.

Justin Sirignano and Konstantinos Spiliopoulos. DGM: a deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018. doi: 10.1016/j.jcp.2018.08.029.

Remco van der Meer, Cornelis W. Oosterlee, and Anastasia Borovykh. Optimally weighted loss functions for solving PDEs with neural networks. *Journal of Computational and Applied Mathematics*, 405:113887, 2022. doi: 10.1016/j.cam.2021.113887.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.

Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why PINNs fail to train: a neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022. doi: 10.1016/j.jcp.2021.110768.

A SETTING AND NOTATION

Let $d \in \mathbb{N}$ and $T \in (0, \infty)$, and let $\Omega \subset \mathbb{R}^d$ be the spatial domain with boundary $\partial\Omega$. We denote by $C(\mathcal{X}, \mathcal{Y})$ the continuous functions from $\mathcal{X}$ to $\mathcal{Y}$. The operators ∇, ∇^2 and Δ act in the spatial variable only, so that Δ denotes the Laplacian throughout and the time step is always written Δt; ∂_t denotes the time derivative, and $a \wedge b := \min(a, b)$. Throughout, $\|\cdot\|$ denotes a norm of a vector or a matrix and $|\cdot|$ the absolute value of a scalar; for a matrix M we write $\|M\|_F := (\sum_{i,j=1}^d M_{ij}^2)^{1/2}$ for the Frobenius norm. We write $\mathcal{N}(\mu, \Sigma)$ for the Gaussian law with mean μ and covariance Σ.

We denote by W a standard d-dimensional Brownian motion; the letter B is reserved for balls and for the metastable set of a committor problem. The process X solves the SDE (6) and has law $\mathbb{P}$, whereas Z denotes the canonical process on $C([0, T], \mathbb{R}^d)$ under a sampling measure $\mathbb{Q}$, with $\mathbb{Q}^0$ the driftless law (13). Here X is the process through which the PDE is represented and Z the process along which the residual is evaluated; the two need not agree. For a continuous semimartingale Z, that is, a process decomposing into a continuous local martingale and a continuous finite-variation part, we write $\langle Z^i, Z^j \rangle$ for the quadratic covariation of its ith and jth components, i.e. the limit in probability of $\sum_{k=0}^{N-1} (Z^i_{u_{k+1}} - Z^i_{u_k})(Z^j_{u_{k+1}} - Z^j_{u_k})$ along refining partitions $0 = u_0 < \cdots < u_N = s$ of $[0, s]$. Condition (8) constrains Z through this object alone and leaves its finite-variation part free, which is what admits the endpoint-conditioned processes of Section 2.2; conditioning on a terminal value is a Doob h-transform, altering only the drift. We write $\tau := \inf\{s : X_s \notin \Omega\}$ for the first exit time and $\mathcal{T}$ for a generic stopping time at which the residual is evaluated, and anchors are drawn from a distribution $\nu_{\Omega \times [0,T]}$ with full support on $\Omega \times [0, T]$.

Concerning the boundary value problem, V denotes the exact solution of (1a)-(1c) and $\varphi \in \mathcal{F}$ a generic approximation of it, with $\psi := \varphi - V$ the approximation error. The operator L is the generator (2) with coefficients b and σ, the function h is the nonlinearity, and f and g are the terminal and boundary data; the diffusion coefficient σ is matrix-valued unless stated otherwise, the isotropic case being written $\sigma\sigma^\top = \sigma^2 I$, where σ also denotes the corresponding scalar; ρ_φ is the pointwise residual (3). The nonlinearity always receives $\sigma^\top \nabla V$ rather than ∇V, since this is the combination appearing in the stochastic representation. The integrated covariance is $\Sigma(t) := \int_0^t \sigma(u)\sigma(u)^\top \mathrm{d}u$, and in the isotropic case $\sigma\sigma^\top(u) = \sigma^2(u) I$ we write $\kappa(t) := \int_0^t \sigma^2(u)\,\mathrm{d}u$, so that $\Sigma = \kappa I$. Finally, $A_\varphi := \sigma^\top \nabla^2 \varphi\, \sigma$ denotes the Hessian of φ conjugated by the diffusion coefficient, the object the discretized residual penalizes implicitly; its trace $\operatorname{tr} A_\varphi$ is the second-order part of $L\varphi$.

A residual is evaluated over a window $[t_0, r]$ of length $\ell := r - t_0$, discretized on a grid $t_0 = s_0 < \cdots < s_n = \mathcal{T}$ with uniform step Δt, so that $\ell = n\Delta t$. The bridge anchors are z_0 and z_r, and m denotes the number of chained bridges, joined at deterministic waypoints y_j. We write $R^{\mathcal{T}}_{t_0}(\varphi; Z)$ for the continuous residual (9), $\widehat{R}$ for its discretization (44) and $\mathcal{E} := \widehat{R} - R$ for the resulting discretization error. The loss (10) decomposes into the interior term $\mathcal{L}^{\mathbb{Q}}_{\mathrm{int}}$ and the terminal and boundary terms $\mathcal{L}_T, \mathcal{L}_b$ of (5), weighted by $\alpha_{\mathrm{int}}, \alpha_T, \alpha_b > 0$; superscripts name the objective and subscripts the term within it, as in $\mathcal{L}^{\mathrm{PINN}}_{\mathrm{int}}$. Batch sizes are k_{int} and k_b.

B PROOFS

B.1 STANDING ASSUMPTIONS

We work throughout under the regularity conditions of Nüsken & Richter (2023), restated here so that the present paper is self-contained. Assumption B.1 is their Assumption 1; the conditions on the coefficients and the data in Assumption B.2 are those of their Proposition 3.4, which is the statement that the diffusion loss characterizes the solution of (1a)-(1c). Throughout, $\mathcal{F} \subset C^{2,1}(\Omega \times [0, T], \mathbb{R}) \cap C(\overline{\Omega} \times [0, T], \mathbb{R})$ denotes the class of admissible approximations, in practice realized by neural networks.

Assumption B.1 (standing assumptions, after Nüsken & Richter, 2023, Assumption 1)**.** The following hold.

1. The domain Ω is either bounded with piecewise smooth boundary, or $\Omega = \mathbb{R}^d$.
2. The boundary value problem (1a)-(1c) admits a unique classical solution $V \in C^{2,1}(\Omega \times [0, T], \mathbb{R}) \cap C(\overline{\Omega} \times [0, T], \mathbb{R})$, and the gradient of V satisfies a polynomial growth condition

in x, that is,

$$\|\nabla V(x,t)\| \le C\big(1+\|x\|^q\big), \qquad (x,t)\in\Omega\times[0,T], \tag{31}$$

for some $C, q > 0$.

3. Complemented with a deterministic initial condition, the SDE (6) admits a unique strong solution, globally in time.

The growth condition in the second item is trivially satisfied whenever Ω is bounded, as for the perforated square of Section 3.1. The other two problems are posed on unbounded domains, where it has to be checked. For the committor of Section 3.2 on $\Omega = \mathbb{R}^d \setminus (A \cup B)$, the reduction (58) gives $V(x) = V_2(y)$ and hence $\nabla V = (\nabla V_2, 0)$, with $V_2 \in [0,1]$ the committor of the two-dimensional problem; its gradient is bounded, by interior estimates away from $\partial\Omega$ and by the confining drift of (26) at infinity. For the parabolic problem of Section 3.3 on $\Omega = \mathbb{R}^d$, differentiating (59) writes $\nabla V(t,x)$ as a weighted average of values of ∇g, and $\|\nabla g(x)\| = 2\|x\|/(1+\|x\|^2) \le 1$, so $C = 1$ there. In both cases $\|\nabla V\|$ is bounded and the condition holds for any $q > 0$.

Assumption B.2 (coefficients and data, after Nüsken & Richter, 2023, Proposition 3.4)**.** The coefficients b and σ are globally Lipschitz continuous in x, uniformly in $t \in [0,T]$, and the data satisfy, for constants $C, p \ge 0$ and all x, y, z and $t \in [0,T]$,

$$|f(x)| \le C(1+\|x\|^p), \qquad |g(x,t)| \le C(1+\|x\|^p), \tag{32}$$

$$|h(t,x,y,z)| \le C(1+\|x\|^p + |y| + \|z\|), \tag{33}$$

$$|h(t,x,y,z) - h(t,x,y',z)| \le C|y-y'|, \tag{34}$$

$$|h(t,x,y,z) - h(t,x,y,z')| \le C\|z-z'\|. \tag{35}$$

B.2 Proofs of the main results

Proof of Proposition 2.2. Itô's formula for continuous semimartingales gives

$$\mathrm{d}\varphi(Z_s,s) = \partial_s\varphi(Z_s,s)\,\mathrm{d}s + \nabla\varphi(Z_s,s)\cdot\mathrm{d}Z_s + \tfrac{1}{2}\sum_{i,j=1}^{d}\partial_{x_i}\partial_{x_j}\varphi(Z_s,s)\,\mathrm{d}\langle Z^i, Z^j\rangle_s.$$

By (8) the last term equals $\big(L\varphi - b\cdot\nabla\varphi\big)(Z_s,s)\,\mathrm{d}s$, so that

$$\mathrm{d}\varphi(Z_s,s) = \big((\partial_s + L)\varphi - b\cdot\nabla\varphi\big)(Z_s,s)\,\mathrm{d}s + \nabla\varphi(Z_s,s)\cdot\mathrm{d}Z_s.$$

Integrating over $[t_0, \mathcal{T}]$ and adding $\int_{t_0}^{\mathcal{T}} h\big(Z_s, s, \varphi(Z_s,s), \sigma^\top\nabla\varphi(Z_s,s)\big)\,\mathrm{d}s$ to both sides gives the claim. □

Proof of Proposition 2.4. Necessity is immediate from Proposition 2.2 and $\rho_\varphi \equiv 0$. Sufficiency: by Proposition 2.2 and $\mathcal{L}^{\mathbb{Q}} = 0$, $\int_{t_0}^{\mathcal{T}} \rho_\varphi(Z_s,s)\,\mathrm{d}s = 0$ almost surely for every anchor. Dividing by $\ell = \mathcal{T} - t_0$ and letting $\ell \to 0$, the Lebesgue differentiation theorem gives $\rho_\varphi \equiv 0$ on $\operatorname{supp}\nu_{\Omega\times[0,T]}$; this argument needs no Markov assumption and covers the bridges of Section 2.2, whose h-transform drift blows up like $1/(r-s)$ as $s \uparrow r$. The short-window limit follows from Proposition 2.2 by dominated convergence, using $\mathcal{T} \to t_0$ a.s. and continuity of ρ_φ. □

Proof of Proposition 2.7. Since $g = V|_{\partial\Omega\times[0,T]}$, (9) gives $R^{\mathrm{exit}} = R^r_{t_0}(\varphi;Z) - \psi(z_r,r)$ and $R^{\mathrm{bb}} = R^r_{t_0}(\varphi;Z) + \psi(z_0,t_0) - \psi(z_r,r)$; note that both anchors are prescribed, so $Z_{t_0} = z_0$ and $Z_r = z_r$ are deterministic.

(i) By Proposition 2.2, $R^r_{t_0}(\varphi;Z) = \int_{t_0}^r \rho_\varphi(Z_s,s)\,\mathrm{d}s$, hence $|R^r_{t_0}(\varphi;Z)| \le \|\rho_\varphi\|_\infty\,\ell$ pathwise. Expanding the square and using boundedness of ψ on $\overline{\Omega}\times[0,T]$ yields (21). No limit $z_0 \to z_r$ and no Markov or path-regularity property of the bridge is used; only that $Z_s \in \Omega$ on $[t_0, r]$, so that ρ_φ is evaluated where it is bounded.

(ii) Applying Proposition 2.2 to φ and to V (for which $\rho_V \equiv 0$) and subtracting gives the exact identity

$$R^{\mathrm{bb}} = -\int_{t_0}^r \nabla\psi(Z_s,s)\cdot\mathrm{d}Z_s + \int_{t_0}^r \big(b\cdot\nabla\psi + h_\varphi - h_V\big)(Z_s,s)\,\mathrm{d}s, \tag{36}$$

where $h_\varphi := h(\cdot,\cdot,\varphi,\sigma^\top\nabla\varphi)$ and likewise for h_V. The $\mathrm{d}s$ integral has a bounded integrand and is $O(\ell)$ pathwise. In the stochastic integral we may *not* discard the finite-variation part of Z: the h-transform drift of (17) blows up like $1/(r-s)$ as $s\uparrow r$ and is exactly what pins the path, so it contributes at the same order as the martingale part. Instead we freeze the integrand and use the pinning directly,

$$\int_{t_0}^{r}\nabla\psi(Z_s,s)\cdot\mathrm{d}Z_s=\nabla\psi(z_0,t_0)\cdot(z_r-z_0)+\int_{t_0}^{r}\big(\nabla\psi(Z_s,s)-\nabla\psi(z_0,t_0)\big)\cdot\mathrm{d}Z_s, \tag{37}$$

the first term being deterministic because $Z_r-Z_{t_0}=z_r-z_0$. For the remainder, a second application of Itô's formula to $\frac12(Z_s-z_0)^\top\nabla^2\psi(z_0,t_0)(Z_s-z_0)$ together with $\mathrm{d}\langle Z\rangle_s=\sigma\sigma^\top\mathrm{d}s$ gives $\frac12(z_r-z_0)^\top\nabla^2\psi(z_0,t_0)(z_r-z_0)-\frac12\int_{t_0}^{r}\operatorname{tr}\big(\nabla^2\psi\,\sigma\sigma^\top\big)\mathrm{d}s$ up to third-order terms, hence a remainder of size $O(\|z_r-z_0\|^2+\ell)=O(\ell)$ in L^2 under $\|z_r-z_0\|=O(\sqrt{\ell})$. The leading term is of size $O(\sqrt{\ell})$, so the cross term in the square is $O(\ell^{3/2})$, which gives (22). The final claim follows by conditioning on z_0 and inserting the assumed anchor covariance. Conversely, if $z_0,z_r\in\partial\Omega$ with $z_0\to z_r$ along a C^1 boundary, then $(z_r-z_0)/\|z_r-z_0\|$ converges into $T_{z_r}\partial\Omega$, so the limiting quadratic form is degenerate in the normal direction and cannot equal $\sigma\sigma^\top$. □

Full statement and proof of Proposition C.1. As in the sketch, with the split $Z=Z^{(0)}+M$ into finite-variation and martingale parts (using (8) for $\mathrm{d}\langle M\rangle_s$), the displayed sum in $\mathcal{E}$ is not itself a martingale once Z carries an h-transform drift, but its finite-variation contribution is $O(\Delta t)$ per interval and hence $o(\Delta t)$ in $\mathbb{E}[\mathcal{E}^2]$ after summing $n=\ell/\Delta t$ intervals and centering, leaving only the martingale contribution computed in the sketch, $\mathbb{E}[\mathcal{E}^2]=\frac12\mathbb{E}\int_{t_0}^{\mathcal{T}}\|A_\varphi(Z_s,s)\|_F^2\,\mathrm{d}s\;\Delta t+o(\Delta t)$. The cross term $2\mathbb{E}[R\mathcal{E}]$ is $O(\ell^2\Delta t)$, not $o(\Delta t)$, so (45) should be read as a joint expansion in ℓ and Δt; it vanishes identically at $\varphi=V$, where $R\equiv0$.

If $\operatorname{tr}A_\varphi$ in the Stratonovich correction (46) is instead estimated by a Hutchinson probe $\varepsilon^\top A_\varphi\varepsilon$, $\varepsilon\sim N(0,I)$, the probe itself contributes a variance $\frac12\sigma^4\|\nabla^2\varphi\|_F^2\,\ell\Delta t$ to $\mathbb{E}[(\widehat{R}^\bullet)^2]$ – a curvature penalty of the same $O(\ell\Delta t)$ order the correction was introduced to remove, so recovering the benefit of (i) requires the exact trace. □

B.3 The Brownian bridge as a Doob h-transform

We derive (16) and (17). The construction is classical: conditioning a Markov process on its terminal value is a Doob h-transform, and bridges of general Markov processes are obtained in this way by Fitzsimmons et al. (1993); see Rogers & Williams (2000) for the h-transform itself. We record the argument for the driftless process (13) because (17) is the form the algorithm samples. Throughout, $\sigma=\sigma(s)$ is state-independent, $\Sigma(t)=\int_0^t\sigma(u)\sigma(u)^\top\mathrm{d}u$, and we abbreviate

$$C(s,u):=\Sigma(u)-\Sigma(s),\qquad s\le u, \tag{38}$$

assuming $C(s,r)$ invertible for $s<r$.

Transition density of the driftless process. The increments of (13) are independent and Gaussian: $Z_u-Z_s=\int_s^u\sigma(v)\,\mathrm{d}W_v$ is centered with covariance $\int_s^u\sigma\sigma^\top=C(s,u)$ and is independent of $\mathcal{F}_s$. Hence

$$p(s,z;u,y)=\big((2\pi)^d\det C(s,u)\big)^{-1/2}\exp\Big(-\tfrac12(y-z)^\top C(s,u)^{-1}(y-z)\Big). \tag{39}$$

Transition law of the bridge. Fix $s<u<r$ and condition on $Z_s=z$ and $Z_r=z_r$. Given $Z_s=z$ we may write $Z_u=z+\xi_1$ and $Z_r=z+\xi_1+\xi_2$ with $\xi_1\sim\mathcal{N}(0,C(s,u))$ and $\xi_2\sim\mathcal{N}(0,C(u,r))$ independent, so that (Z_u,Z_r) is jointly Gaussian with

$$\operatorname{Cov}(Z_u,Z_r)=C(s,u),\qquad \operatorname{Var}(Z_r)=C(s,u)+C(u,r)=C(s,r). \tag{40}$$

The Gaussian conditioning formula therefore gives

$$Z_u\mid(Z_s=z,\,Z_r=z_r)\sim\mathcal{N}\big(z+G(z_r-z),\,(I-G)C(s,u)\big),\qquad G:=C(s,u)\,C(s,r)^{-1}, \tag{41}$$

since the conditional covariance is $C(s,u)-C(s,u)C(s,r)^{-1}C(s,u)=(I-G)C(s,u)$. Taking $s=s_j$ and $u=s_{j+1}$ yields (17). Equivalently, one may start from the Markov property, which gives the conditional density as $p(s,z;u,y)\,p(u,y;r,z_r)/p(s,z;r,z_r)$, and complete the square.

Algorithm 1 Simulation-free diffusion loss (one gradient step, covers both instantiations)

Require: approximation $\varphi \in \mathcal{F}$ with parameters θ, batch size K, sub-steps $n \geq 1$, step size Δt (window $\ell = n\Delta t$), sampler (Section 2.1 or Section 2.2), weights $\alpha_{\mathrm{int}}, \alpha_T, \alpha_b$

1: **for** $k = 1, \ldots, K$ **do**

2: draw anchors $(t_0^{(k)}, z_0^{(k)}) \sim \nu_{\Omega\times[0,T]}$, and $z_r^{(k)}$: free endpoint (Section 2.1) or interior/boundary target (Section 2.2)

3: draw $Z_{s_1}^{(k)}, \ldots, Z_{s_{n-1}}^{(k)}$ in closed form, (13) or (17)

4: evaluate $\varphi, \nabla\varphi$ at the $n+1$ nodes and form

$$\widehat{R}^{(k)} = \varphi(Z_{s_n}^{(k)}, s_n) - \varphi(Z_{s_0}^{(k)}, s_0) - \sum_{j=0}^{n-1} \nabla\varphi(Z_{s_j}^{(k)}, s_j) \cdot \big(Z_{s_{j+1}}^{(k)} - Z_{s_j}^{(k)}\big) + \sum_{j=0}^{n-1} \big(b \cdot \nabla\varphi + h\big)(Z_{s_j}^{(k)}, s_j)\, \Delta t \quad (44)$$

5: **end for**

6: descend on θ along $\alpha_{\mathrm{int}} \frac{1}{K} \sum_{k=1}^{K} (\widehat{R}^{(k)})^2 + \alpha_T \mathcal{L}_T + \alpha_b \mathcal{L}_b$

The bridge SDE. Let $h(s,z) := p(s,z;r,z_r)$. As a transition density in its backward variables, h satisfies the Kolmogorov backward equation $(\partial_s + L_0)h = 0$ with $L_0 = \frac{1}{2}\sum_{i,j=1}^{d} (\sigma\sigma^\top)_{ij}(s)\, \partial_{z_i}\partial_{z_j}$, i.e. h is space-time harmonic for (13). Conditioning on $Z_r = z_r$ is the Doob h-transform associated with this h, whose generator acts as

$$L^h f = \frac{1}{h}\big(L_0(hf) - f L_0 h\big) = L_0 f + \big(\sigma\sigma^\top \nabla \log h\big) \cdot \nabla f, \quad (42)$$

where the middle identity follows by expanding $L_0(hf)$ and using the symmetry of $\sigma\sigma^\top$. From (39),

$$\log h(s,z) = -\tfrac{1}{2}(z_r - z)^\top C(s,r)^{-1}(z_r - z) + c(s), \qquad \nabla_z \log h(s,z) = C(s,r)^{-1}(z_r - z), \quad (43)$$

the normalization $c(s)$ being independent of z. Substituting into (42) gives the drift $\sigma(s)\sigma(s)^\top C(s,r)^{-1}(z_r - Z_s)$, which is (16). Note the order of the factors: $\sigma\sigma^\top$ multiplies $C(s,r)^{-1}$ from the left, and the two commute only in special cases, for instance when σ is constant in time.

Quadratic variation and consistency. By (42) the h-transform adds a first-order term and leaves the second-order part L_0 untouched, so $\mathrm{d}\langle Z^i, Z^j\rangle_s = (\sigma\sigma^\top)_{ij}(s)\,\mathrm{d}s$ as for (8) and (13) holds; this is the admissibility claim of Section 2.2. Finally, the two descriptions agree to first order: as $u \downarrow s$ we have $C(s,u) = \sigma\sigma^\top(s)(u-s) + o(u-s)$, hence $G = \sigma\sigma^\top(s)C(s,r)^{-1}(u-s) + o(u-s)$, so the mean $z + G(z_r - z)$ is the Euler step of (16) and the covariance $(I-G)C(s,u)$ equals $\sigma\sigma^\top(s)(u-s) + o(u-s)$.

C Numerical aspects

This appendix analyzes the discrepancy between the continuous residual (9) and the objective that is minimized in practice. App. C.1 restates the residual as it is computed in practice (Algorithm 1), determines the leading-order effect of evaluating it on a finite grid (Proposition C.1) and draws the consequence for the signal-to-noise ratio (Remark C.2). App. C.2 describes two modifications that remove the resulting bias. App. C.3 turns to the effects that are geometric rather than statistical, and which arise only on a bounded domain.

C.1 The discretized residual

By Proposition 2.2 the continuous residual $R = \int_{t_0}^{\mathcal{T}} \rho_\varphi(Z_s, s)\,\mathrm{d}s$ carries no Itô noise: all stochasticity in $\widehat{R}$ (44) beyond the randomness of the path itself comes from discretizing the stochastic integral on a finite grid. The asymptotics of such discretized stochastic integrals are developed systematically by Jacod & Protter (2012), and Kloeden & Platen (1992) is the standard reference for the numerical discretization of stochastic differential equations; what the following makes explicit is how that error enters the loss (10) rather than a trajectory.

Proposition C.1 (Leading-order discretization effect)**.** *Let $\ell = n\Delta t$ with a uniform step Δt, $A_\varphi := \sigma^\top \nabla^2 \varphi\, \sigma$. Then*

$$\mathbb{E}\big[\widehat{R}^2\big] = \underbrace{\mathbb{E}[R^2]}_{\textit{PDE residual}} + \underbrace{2\,\mathbb{E}[R\,\mathcal{E}]}_{O(\ell^2\Delta t)} + \underbrace{\frac{\Delta t}{2}\,\mathbb{E}\int_{t_0}^{\mathcal{T}} \|A_\varphi(Z_s, s)\|_F^2 \,\mathrm{d}s}_{\textit{implicit Hessian penalty}} + o(\Delta t), \tag{45}$$

where $\mathcal{E} := \widehat{R} - R$.

Proof sketch. Expand $\nabla\varphi(Z_s, s) - \nabla\varphi(Z_{s_j}, s_j) = \nabla^2\varphi(Z_{s_j}, s_j)(Z_s - Z_{s_j}) + O(|s - s_j|)$ and $Z_s - Z_{s_j} = \sigma(W_s - W_{s_j}) + O(|s - s_j|)$; the Itô isometry gives $\mathbb{E}\int_{s_j}^{s_{j+1}} \|A_\varphi(W_s - W_{s_j})\|^2 \mathrm{d}s = \frac{1}{2}\|A_\varphi\|_F^2(\Delta t)^2$ per interval, and summing over $n = \ell/\Delta t$ intervals gives the Hessian term; the Riemann-sum error of the $\mathrm{d}s$-integrals is $O(\Delta t)$ per unit time and enters at order $o(\Delta t)$ after centering. The full statement and the proof can be found in App. B. □

At $\varphi = V$, $R \equiv 0$ and the cross term vanishes identically, so the diffusion loss avoids computing Hessians but does not avoid penalizing them: at finite Δt it minimizes the PDE residual plus $\frac{\Delta t}{2}\|A_\varphi\|_F^2$ integrated along the path, and this Hessian information is traded for a Δt-controlled smoothing bias, not removed.

Remark C.2 (Signal-to-noise). With $\rho_\varphi \approx \rho$ locally constant, $\mathbb{E}[R^2]/\mathbb{E}[\mathcal{E}^2] \approx 2\rho^2 n/\|A_\varphi\|_F^2$. For $n = 1$ the ratio is independent of Δt: shrinking the window does not destroy the signal, it only strengthens the implicit penalty; increasing n at fixed Δt improves the ratio linearly, since the signal accumulates coherently while the discretization error accumulates as a martingale.

C.2 Removing the discretization bias

The $O(\Delta t)$ term in (45) can be removed by either of the two modifications of (44) stated below, where $\widehat{R}_{\Delta t}$ denotes (44) at step Δt and $\widehat{R}^\bullet_{\Delta t}$ its trapezoidal counterpart.

$$\text{(i) trapezoidal:} \qquad \mathbb{E}\big[(\widehat{R}^\bullet_{\Delta t})^2\big] = \mathbb{E}[R^2] + o(\Delta t), \tag{46}$$

$$\text{(ii) independent replicates:} \qquad \mathbb{E}\big[\widehat{R}^{(1)}_{\Delta t}\widehat{R}^{(2)}_{\Delta t}\big] = \mathbb{E}[R^2] - \mathbb{E}\big[\,\mathrm{Var}(R \mid A)\big], \tag{47}$$

(i) replaces the left-point (Itô) rule in (44) by the trapezoidal one, $\frac{1}{2}\big(\nabla\varphi(Z_{s_j}, s_j) + \nabla\varphi(Z_{s_{j+1}}, s_{j+1})\big) \cdot (Z_{s_{j+1}} - Z_{s_j})$, and compensates h by $+\frac{1}{2}\,\mathrm{tr}\,A_\varphi$: it is exact in Δt but needs the Hessian trace at every node (or a Hutchinson probe, which reintroduces a curvature penalty of the same $O(\ell\Delta t)$ order it was meant to remove, see App. B). (ii) draws two independent bridge realizations from the same anchors and replaces the squared loss by their product: it is derivative-free, but $\mathrm{Var}(R \mid A) = O(\ell^3)$ does *not* vanish as $\Delta t \to 0$ and is safe only at $n = 1$.

Both modifications belong to a recent line of work on the discretization bias of BSDE-type losses. Park & Tu (2025) study the role of the integration scheme, which is what (i) changes, while Seo et al. (2026) construct a training objective that is at once unbiased and free of second derivatives, the two properties that (ii) possesses and (i) does not.

modification	cost	derivatives	residual defect
(i) trapezoidal	1 extra $\nabla\varphi$/node + $\mathrm{tr}\,A_\varphi$	Hessian trace	$o(\Delta t)$
(ii) replicates	2 bridge draws	none	$-\mathbb{E}[\mathrm{Var}(R \mid A)] = O(\ell^3)$, Δt-free

Table 1: Two modifications of (44) removing its $O(\Delta t)$ discretization bias.

C.3 Implementation on bounded domains

The effects quantified above originate in the time grid and are present on any domain. On a bounded Ω the exit bridge of Section 2.2 raises three further questions: where the path is allowed to go, how a chain of bridges is routed, and how the window and the number of hops are chosen. These concern the sampling process rather than the residual identity, which by Proposition 2.2 holds for any admissible process; they determine the variance of the estimator and, through the boundary substitution discussed below, its bias. The values used in each experiment are reported in App. D.

Paths that leave and return. By Proposition 2.2 the identity holds for any admissible sampling process, so an excursion of the bridge outside Ω does not by itself invalidate the interior terms; it only evaluates φ where (1a) does not constrain it. What it does affect is the substitution $\varphi(Z_{\mathcal{T}}, \mathcal{T}) \to g(z_r, r)$ that defines the interior-boundary variant of Section 2.2. That substitution attaches the boundary datum to the anchor z_r, and it is exact only if the first exit of the path from Ω occurs at z_r. If the path leaves Ω before r and re-enters, its actual first exit lies elsewhere on $\partial\Omega$ and the datum is attached to the wrong point. The remaining source of bias is therefore not an approximation of the terminal time itself, but the possibility of a leave-and-return excursion.

Three remedies are available, in increasing order of cost. Anchor placement makes the excursions rare: a hop of length $L = \sqrt{d\sigma^2\ell}$ has transverse spread $\sqrt{\sigma^2\ell}/2 = L/(2\sqrt{d})$ about its chord, at its widest in the middle of the hop, so it suffices that the chord keep a clearance of a few multiples of $L/(2\sqrt{d})$ from $\partial\Omega$. Rejection removes the excursions from the sample: each node is redrawn from its own conditional law until it lands in Ω, which is a per-node and not a per-path operation and therefore costs a few extra normal draws rather than a discarded batch. It is not free of consequence, since a node that fails every retry must be frozen at its predecessor, and a frozen node shortens the quadratic variation of that path and so perturbs the sampling law; the frozen fraction should be measured, and a control arm without rejection is the direct way to confirm that it does not carry the result. Truncation at the first exit keeps every sample and attaches the datum to the correct point, at the price of locating the crossing between two consecutive nodes and of the $O(\sqrt{\Delta t})$ discrete-exit bias familiar from Euler-Maruyama schemes.

Routing a chain through a non-convex domain. Chaining m bridges requires intermediate anchors $y_1, \dots, y_{m-1}$, and the uniform choice (19) places them on the segment from x_0 to z_r. That segment lies in Ω when Ω is convex, but not otherwise, and the chain is then routed through the complement of Ω by construction rather than by chance. No amount of node rejection repairs this, since it is the anchors themselves that are outside. The anchors are free, however: Proposition 2.2 constrains them only through the quadratic variation of the bridges that connect them, each hop satisfies (8) whatever the anchor sequence, and routing around an obstacle therefore costs nothing in correctness and only selects a different admissible $\mathbb{Q}$. A simulated path would have to diffuse around the obstacle instead.

The detour should stay in closed form, so that the anchors remain available without a geometric solve. For the perforated square of Section 3.1 it suffices to push a straight-line point that has fallen into a hole radially out to a fixed radius: as long as that radius stays below half the hole pitch, the pushed point lies in Ω for every anchor pair, so that no geometric test is required. On a spherical shell one can interpolate the radius and the direction separately. In both cases what matters is not only that the anchor lies in Ω but by how much: an anchor placed just off an obstacle produces a hop whose transverse spread straddles it. At the settings of App. D the pushed anchors clear the holes by 0.4 times the gap between neighboring holes, about 4.5 times the spread $L/(2\sqrt{d})$, at every N.

Choosing the window and the number of hops. Two numbers are fixed before training: the window ℓ of a single hop and the number m of hops in a chain. When Ω or the coefficients carry a feature scale – the gap between obstacles, the spacing of the wells, the thickness of a shell – a convenient rule is to set the hop length $\sqrt{d\sigma^2\ell}$ to a fixed fraction of that scale and to take m as the median chord $\|z_r - x_0\|$ divided by it, so that a chain reaches $\partial\Omega$ on average. Both numbers are then tied to the geometry and one dimensionless knob is left.

Two diagnostics are worth recording. Since m is determined by a chord length that varies across the batch, a vectorized implementation has to pad every chain to a common number of hops. Padding beyond the required number leaves the sampled path unchanged, whereas an anchor pair requiring more hops than the padding provides yields a chain that terminates in the interior of Ω, in which case the boundary datum is evaluated at a point at which it is not defined. Comparing the realized m with the padding detects this. Second, the bridges are driftless while the problem need not be. The drift is not neglected – it enters (9) through $b \cdot \nabla\varphi$, which is exact at any ℓ – but a sampling law that ignores it is a poor proposal once a single hop is dominated by transport rather than by diffusion, and the ratio of the two over one hop, $\|b\|\,\ell/\sqrt{d\sigma^2\ell} = \|b\|\sqrt{\ell/(d\sigma^2)}$, measures when that happens.

D EXPERIMENTAL DETAILS

This section records what is needed to reproduce the three experiments of Section 3: App. D.1 states the conventions shared by all of them, and the three subsections that follow give, for each problem in turn, the domain and the coefficients, the construction of its reference solution, the anchor laws and the settings of every loss compared on it.

D.1 EXPERIMENTAL DESIGN AND METRICS

The three experiments differ in geometry, in how the reference solution is obtained and in the network used, and those are stated with each of them below. The elements common to all three are collected here.

Controlled factors. A comparison is always between losses on one problem instance at one setting, and within such a cell every arm shares the network architecture, the optimizer and its schedule, the number of collocation points per step and the compute budget. Problem data that are drawn rather than specified – the coefficients of the perforated square, the mixture weights of the chain – are drawn once from a fixed seed and held across arms, weights and training seeds, so that a difference between arms is never a difference between problem instances.

Initialization. Every layer of every network is initialized identically across arms: weights and biases are drawn from $\mathcal{U}(-\kappa,\kappa)$ with $\kappa = (\text{fan-in})^{-1/2}$, and the bias of the output layer is then set to 1 rather than drawn, so that the network starts at a constant rather than near zero. Where the domain is not of order one, the first layer's weights are divided by a constant, which is equivalent to rescaling the input and leaves every loss unchanged; the value is stated with the experiment that uses it.

Matched compute budgets. Throughput differs between the losses by up to an order of magnitude: a PINN residual costs one forward pass and its second derivatives at a single point, whereas a chained bridge evaluates φ and $\nabla\varphi$ at each of its $mn+1$ nodes. Comparing at equal iteration counts would therefore compare unequal budgets. Each cell is instead given the same wall clock on one GPU, and the learning-rate schedule is driven by elapsed time, or by a step horizon projected from the arm's own measured throughput, so that every arm completes its decay within its budget.

Metrics. All reported errors are relative L^2 errors

$$\frac{\|\varphi - V\|_{L^2(\nu)}}{\|V\|_{L^2(\nu)}}, \qquad \|f\|^2_{L^2(\nu)} = \int f(x)^2 \,\nu(\mathrm{d}x), \tag{48}$$

computed as an average over a fixed test set drawn from ν and held across arms and seeds. The evaluation measure ν and the reference V are named with each experiment. The references are deterministic in all three cases – a closed form (App. D.2), a two-dimensional finite-difference solve lifted to $\mathbb{R}^d$ (App. D.3), and a Cole-Hopf reduction to a one-dimensional quadrature (App. D.4) – so none of them contributes Monte Carlo variance to the reported error and the numbers are not bounded below by a sampling floor.

Summary statistic per run. Where a run is reduced to a single number, that number is taken over the last 10% of its evaluations rather than as the minimum over training, so that it reports a converged plateau and not a transient that a stopping rule would have to detect. App. D.2 uses the median of that window and App. D.3 its mean.

Boundary weights and seeds. The weight α_b is a property of a loss rather than of a problem, and the losses compared here do not share an optimal value; fixing one value for all of them would favor whichever loss that value happens to suit. Each loss is therefore swept over its own ladder of α_b at a single seed, and its best cell is re-run at further seeds, three in total unless stated otherwise; selecting on one seed and measuring on the replicates keeps the selection out of the reported number, at the price of leaving the sweeps themselves single-seed. Shaded bands are the minimum and maximum over the seeds at a cell, not a confidence interval; since the range of a sample grows with the number of draws, bands are comparable only at equal seed count. The HJB experiment of App. D.4 carries no boundary term and no sweep.

Remark D.1 ($h \equiv 0$ is the identifiability point of α_b). When $h \equiv 0$, the interior residual alone does not determine V. In that case $\rho_\varphi = L\varphi$, so $\mathcal{L}^{\mathbb{Q}}_{\text{int}}$ is invariant under $\varphi \mapsto \varphi + \eta$ for every L-harmonic η, and V is itself L-harmonic; the interior term cannot separate V from $V + \eta$. At $\alpha_b = 0$ the objective is thus unidentifiable for any loss whose interior term is the only place the data enter – PINN and the losses of Section 2.1 included. The exit bridge of Section 2.2 is the exception: it substitutes $g(z_r, r)$ for $\varphi(Z_{\mathcal{T}}, \mathcal{T})$ inside (9), so the boundary datum enters the interior term itself, and by Proposition 2.7(i) that term tends to $\mathbb{E}[\psi(z_r, r)^2]$, the mean-square boundary mismatch at the exit anchors, as $\ell \to 0$ – rather than to a multiple of the PINN residual. A problem with $h \equiv 0$ therefore isolates a property that the weighted sum $\alpha_{\text{int}}\mathcal{L}_{\text{int}} + \alpha_b\mathcal{L}_b$ otherwise hides, and which Fig. 7 probes as $\alpha_b \to 0$: only a boundary-anchored interior residual remains well posed in that limit.

D.2 Perforated square: geometry and closed-form solution

Geometry. In the domain Ω of Section 3.1 the centers c_n are the nodes of an equidistant $\sqrt{N}\times\sqrt{N}$ grid in $(0,1)^2$ and the radius is $r = J/(2\sqrt{N})$, i.e. the hole diameter is the fraction J of the grid pitch $1/\sqrt{N}$; we use $J = 0.15$ throughout. The perforated *area* fraction is then $\pi J^2/4 \approx 0.018$, independent of N, so only the number of boundary components changes. The holes' share of the total *perimeter* is not constant: it equals $\pi J\sqrt{N}/(4 + \pi J\sqrt{N})$ and grows from 0.32 at $N = 16$ to 0.59 at $N = 144$, so a length-weighted boundary sample places an increasing fraction of its points on holes as N grows.

We construct q and g in (25) so that the solution is explicit. Write $u_k(x) = x - c_k$ and $\rho_k(x) = \|u_k(x)\|$.

Step 1: harmonic building blocks. For $x \neq c_k$,

$$\nabla \log \rho_k = \frac{u_k}{\rho_k^2}, \qquad \Delta \log \rho_k = \nabla \cdot \frac{u_k}{\rho_k^2} = \sum_{i=1}^{d}\Big(\frac{1}{\rho_k^2} - \frac{2(u_k)_i^2}{\rho_k^4}\Big) = \frac{d-2}{\rho_k^2}, \tag{49}$$

which vanishes precisely because $d = 2$: each $\log \rho_k$ is harmonic on $\mathbb{R}^2 \setminus \{c_k\}$. By construction $c_k \in B(c_k, r)$ and $B(c_k, r) \cap \Omega = \emptyset$, so every c_k lies outside $\overline{\Omega}$ and each $\log \rho_k$ is harmonic on an open neighborhood of $\overline{\Omega}$. Placing the singularities inside the holes is the method of fundamental solutions; it is what makes the construction exact for a domain no separable coordinate system fits. Likewise $\operatorname{Re}\zeta^m$ and $\operatorname{Im}\zeta^m$, $\zeta = x_1 + \mathrm{i}x_2$, are harmonic on all of $\mathbb{R}^2$; we take $P = \beta_1(x_1^2 - x_2^2) + \beta_2\, 2x_1x_2 + \beta_3 x_1 + \beta_4 x_2$.

Step 2: the candidate. With coefficients $a_1, \ldots, a_N$ put

$$V^\star(x) = \sum_{k=1}^{N} a_k \log \frac{\rho_k(x)}{r} + P(x), \qquad \nabla V^\star = \sum_{k=1}^{N} a_k \frac{u_k}{\rho_k^2} + \nabla P, \tag{50}$$

so $\Delta V^\star = 0$ in Ω by (49) and linearity. Both $V^\star$ and $\nabla V^\star$ are elementary and evaluable in floating point to machine precision.

Step 3: closing the equation. Choose

$$q := \|\nabla V^\star\|^2, \qquad g := V^\star|_{\partial\Omega}. \tag{51}$$

Then $\Delta V^\star + \lambda(\|\nabla V^\star\|^2 - q) = 0 + 0 = 0$, so $V^\star$ solves (25). Note the normalization $\log(\rho_k/r)$ makes the k-th term vanish on $\partial B(c_k, r)$, so a_k is not read off that hole's own boundary datum; each hole carries a level set by the remaining terms.

Well-posedness. With q, g as in (51) and $\lambda > 0$, the field $V^\star$ of (50) is the *unique* classical solution of (25). Existence is Step 3. For uniqueness, apply the Cole-Hopf transform $\Psi = e^{\lambda V}$, so that $\nabla\Psi = \lambda\Psi\nabla V$ and

$$\Delta\Psi = \lambda\big(\nabla\Psi \cdot \nabla V + \Psi\Delta V\big) = \lambda\Psi\big(\Delta V + \lambda\|\nabla V\|^2\big). \tag{52}$$

Hence V solves (25) if and only if the positive function Ψ solves the *linear* problem

$$\Delta\Psi - \lambda^2 q\Psi = 0 \ \text{ in } \Omega, \qquad \Psi = e^{\lambda g} \ \text{ on } \partial\Omega. \tag{53}$$

Since $q = \|\nabla V^\star\|^2 \geq 0$, the zeroth-order coefficient of (53) is non-positive, so the operator satisfies the maximum principle and the Dirichlet problem (53) has at most one solution on the bounded domain Ω. The map $\Psi \mapsto \lambda^{-1} \log \Psi$ is injective on $\Psi > 0$, and $e^{\lambda V^\star}$ is such a solution, hence $V^\star$ is the only one.

Step 4: fixing the amplitude across N. The coefficients a_k are a random draw, and a fresh one for each N, so the share of $V^\star$ carried by the hole terms rather than by the smooth background P varies with N by luck rather than systematically. Write $\mathrm{sd}(\cdot)$ for the standard deviation under the uniform law on Ω; measured $\mathrm{sd}(V^\star - P)/\mathrm{sd}(P)$ is $0.59,\ 1.24,\ 0.83,\ 1.32$ at $N = 16, 36, 64, 144$. An experiment that puts N on an axis would then confound the number of boundary components with the strength of the structure the solution carries – visibly so: at this raw draw the nearest-boundary predictor of Remark D.2 scores $0.256,\ 0.314,\ 0.328,\ 0.152$, tracking the draw rather than N. We therefore rescale

$$a_k \ \longmapsto\ \varrho\, \frac{\mathrm{sd}(P)}{\mathrm{sd}(V^\star - P)}\, a_k, \qquad \text{so that} \qquad \mathrm{sd}(V^\star - P) = \varrho\, \mathrm{sd}(P), \tag{54}$$

which leaves $V^\star$ harmonic and (25) exact, since (54) only scales coefficients of harmonic functions. Every experiment reported in this paper uses $\varrho = 1$ (measured $1.00,\ 1.00,\ 1.00,\ 1.00$). Under $\varrho = 1$ the nearest-boundary predictor scores $0.364,\ 0.285,\ 0.384,\ 0.133$ at $N = 16, 36, 64, 144$, and these are the values used to normalize Figs. 3 and 6. Numbers obtained at the raw draw are not comparable with them: they belong to a different $V^\star$, hence to different g and q.

Remark D.2 (The fill fraction is not a free parameter). For V harmonic the nearest-boundary predictor $x \mapsto g(\mathrm{proj}_{\partial\Omega}(x))$ has error $O\big(\mathrm{dist}(x, \partial\Omega)\, \|\nabla V\|\big)$, first order in the distance. Perforating a domain *shrinks* that distance everywhere, so raising N at fixed J makes pure boundary interpolation more accurate, not less, and by Proposition 2.7(i) the benchmark degenerates into a regime that measures boundary interpolation rather than PDE solving. Measured relative L^2 error of that predictor at the raw draw: 0.118 ($N = 16$) and 0.142 ($N = 64$) at $J = 0.5$, against 0.254 and 0.331 at $J = 0.15$. Small holes keep the $N + 4$ components while restoring interior far from $\partial\Omega$; the opposing constraint is $\|\nabla V^\star\| \sim |a_k|/r$ as $r \downarrow 0$, by (50).

Coefficients and data. We take $\sigma^2 = 2$, so that $L = \Delta t$, and $b \equiv 0$; the nonlinearity is $h(x, v, z) = \lambda\big(\|z\|^2 - q(x)\big)$ with $\lambda = \frac{1}{2}$ and q as in (51). Note h depends on $\nabla\varphi$ but not on φ: a pointwise dependence would allow q to be inverted for $V^\star$. The coefficients are drawn once, from a fixed seed, and held fixed across methods, weights and training seeds: $a_k \sim \mathcal{U}[-1, 1]/\sqrt{N}$, then re-centered to $\sum_{k=1}^N a_k = 0$ so that $V^\star$ carries no net logarithmic charge, and $\beta \sim \mathcal{U}[-1, 1]^4$ normalized to $\|\beta\| = 1$. The factor $1/\sqrt{N}$ is deliberate: near a hole $\|\nabla V^\star\| \sim |a_k|/r = 2|a_k|\sqrt{N}/J$, so without it the solution's gradient magnitude would grow with N and confound the sweep. Interior points are drawn uniformly on Ω by rejection from the square; boundary points are drawn length-weighted over the four walls and the N circles.

Geometry-dependent quantities. With $J = 0.15$ and $\theta = 0.25$, writing $\mathrm{gap} = 1/\sqrt{N} - 2r$ for the clearance between neighboring holes:

N	16	36	64	144
hole radius r	0.0188	0.0125	0.0094	0.0063
gap	0.2125	0.1417	0.1062	0.0708
hop length $\theta \cdot \mathrm{gap}$	0.0531	0.0354	0.0266	0.0177
window ℓ	$7.1 \cdot 10^{-4}$	$3.1 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$	$7.8 \cdot 10^{-5}$
hops per chain m	12	18	23	32
BSDE step Δt	$9.0 \cdot 10^{-4}$	$4.0 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$	$1.0 \cdot 10^{-4}$

Losses and discretization. All arms use $\alpha_{\mathrm{int}} = 1$, so α_b is the interior-to-boundary ratio, and 200 interior anchors and 1000 boundary points per step. The exit bridge chains m hops, m being the median chord $\|z_b - x_0\|$ divided by the hop length, with 5 sub-steps per hop, so the window of one hop is $\ell = (\theta\, \mathrm{gap})^2/(d\sigma^2)$ and $\sqrt{d\sigma^2\ell}$ equals the hop length. Intermediate anchors are the straight-line points $x_0 \to z_b$, pushed radially out to $r + 0.4\,\mathrm{gap}$ whenever they fall inside a hole; the terminal anchor is z_b itself. Bridge nodes landing outside Ω are resampled from their

own conditional law, up to eight times per node; the worst-placed node, the one just short of z_b, has a first-draw rejection probability of about 0.2, and after the retries a measured $2.2 \cdot 10^{-5}$ of all interior nodes or fewer, across the four hole counts, still fail and are frozen at their predecessor. The simulation-free diffusion loss uses the same window and number of steps. The BSDE arm uses driftless Euler-Maruyama with Δt set so that the per-step displacement is $\mathrm{gap}/5$, at most 1400 steps, and locates the exit by bisecting the crossing segment; since exit is tested at the nodes only, and the step exceeds the hole diameter by the N-independent factor $(1 - J)/(5J) = 1.13$, a path can step across a hole without being stopped; measured at $N = 64$, 32% of paths do so at least once.

Training and evaluation. A fully connected network with three hidden layers of width 64 and $\tanh$ activations, Adam at learning rate 10^{-3} decayed exponentially by a factor 100 over the run, and a budget of 1200 seconds per cell on one GPU. Because throughput differs by an order of magnitude between arms, the budget is wall clock and the decay is driven by elapsed time, so the schedule is comparable across methods. The reported error is the relative L^2 error on $20\,000$ points drawn uniformly on Ω against the closed-form $V^\star$, evaluated in float64, summarized as the median over the last 10% of evaluations so that it is a converged plateau rather than a minimum over training. The boundary weight is swept over $\alpha_b \in \{10^{-4}, 10^{-3}, \dots, 10^{7}\}$ for the PINN, the bridge and the diffusion loss; the BSDE loss has no boundary term. Each method's best α_b is selected on one seed and rerun on two further seeds, with five more at the two cells whose spread exceeds 10% (Fig. 3).

Network capacity versus the objective. The mechanism claimed in Section 3.1 is a property of the objective, so it has to be separated from the alternative explanation that the network is simply too small to represent $N + 4$ distinct hole levels at radius r. We therefore repeated both arms with the width raised from 64 to 256, a factor 15 in parameters, at each method's best α_b:

		width 64	width 256
$N = 16$	PINN, $\alpha_b = 10^{5}$	0.0102	0.0078
	bridge, $\alpha_b = 10^{-1}$	0.0222	0.0224
$N = 64$	PINN, $\alpha_b = 10^{5}$	0.1040	0.0950
	bridge, $\alpha_b = 10^{-2}$	0.0390	0.0383

Quadrupling the width changes the PINN's $N = 64$ error by less than 10%, leaving it a factor 2.5 behind a bridge at width 64, so the dipole-free per-hole offset of Fig. 2 cannot be removed by adding parameters. The bridge is unchanged in the third decimal at both N, which is the converse observation: its floor is set by the estimator rather than by the ansatz, so unlike the PINN it cannot be improved by a larger network – at $N = 16$, where the PINN is already the more accurate of the two, widening increases its margin. Widening also carries a cost under a wall-clock-matched protocol, and not an equal one: throughput falls by a factor 4.7 for the bridge against 1.5 for the PINN, since the bridge evaluates $\sim 10^2$ nodes per anchor where the PINN evaluates a few. These runs predate the amplitude normalization and use the raw coefficient draw rather than $\varrho = 1$, so their absolute values are not comparable with the rest of this appendix; the comparison above is between two widths at one fixed setting, which is what the question requires.

D.3 Metastable chain committor

Centers, weights and σ. The chain of (26) is

$$a_m = r\big(m\varsigma,\ (-1)^m\eta\big) \in \mathbb{R}^2, \qquad \varsigma = \sqrt{1 - 4\eta^2}, \qquad m = 1, \dots, M, \tag{55}$$

so that $\|a_{m+1} - a_m\| = r$ for every m; the centers are then translated to have mean zero, which affects nothing. The offset $\eta > 0$ makes the reaction coordinate turn at every well; with $\eta = 0$ the centers are collinear, V depends on one coordinate alone, and the staircase is not a two-dimensional object. The weights are $w_m \propto e^{\omega_m}$ with $\omega \sim \mathcal{N}(0, \varpi^2 I_M)$ drawn once and mean-centered, so that no well is distinguished except through ω and $\varpi = 0$ recovers the unweighted mixture.

Finally σ is fixed by the barrier. For each pair (i, j) let E_{ij} be the maximum of U along the straight segment $a_i \to a_j$, minus the deeper of the two endpoints; let $E_\star$ be the minimax (bottleneck) value over paths from a_1 to a_M in the complete graph with edge weights E_{ij}, so that a transition may route through any intermediate well rather than only through consecutive ones. Then

$$\sigma^2 := 2E_\star/\beta_\star, \tag{56}$$

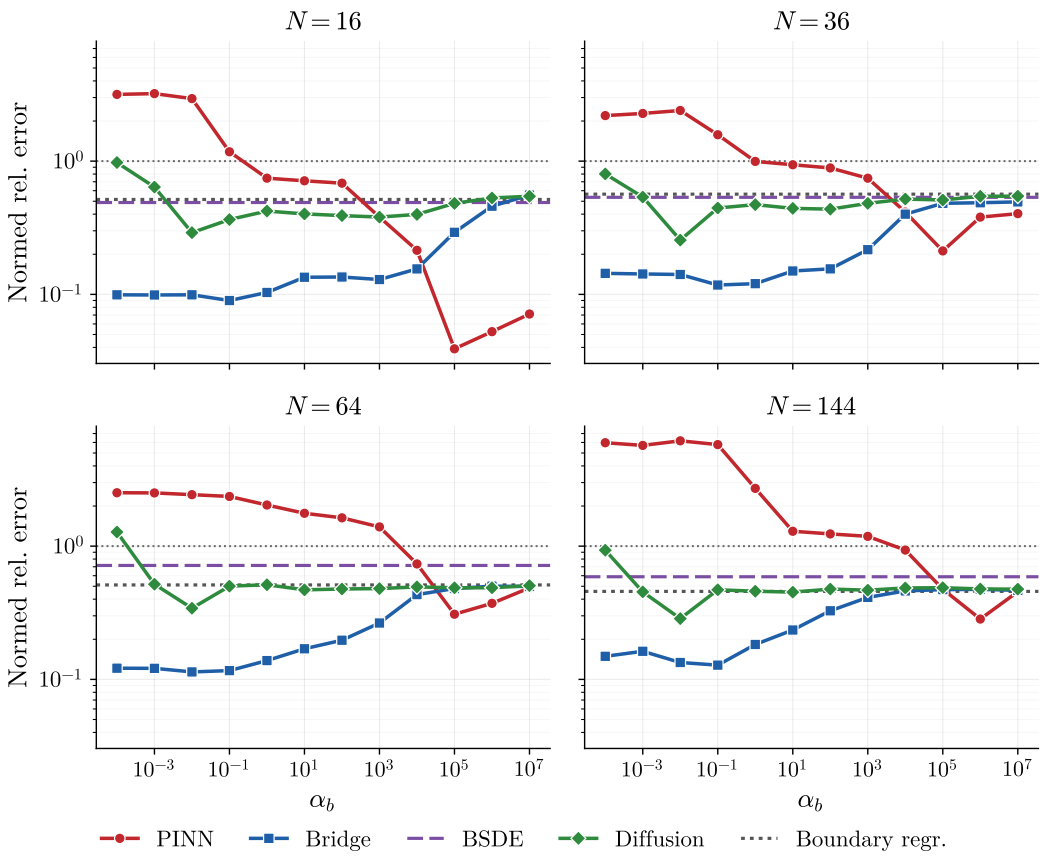

Figure 6: Dependence on the boundary weight α_b, at the settings of Fig. 3 (interior weight fixed to 1, $\varrho = 1$, one seed, equal wall-clock budget per cell). Plotted is the converged relative L^2 error divided by that of the nearest-boundary predictor, so the dotted line at 1 marks no improvement over boundary interpolation; markers are the grid points that were run. All three penalized losses are swept over the same grid, so the curves are directly comparable and every optimum shown is interior; at $N = 144$ the PINN was additionally run at 10^8, which is no better than 10^7. The two losses depend on α_b in opposite directions and cross near $\alpha_b \approx 10^3$ at every N: the PINN needs a large penalty because the boundary datum enters its objective only through that term, whereas the exit bridge already carries the boundary datum $g(z_r, r)$ inside its interior residual (9), which by Proposition 2.7(i) itself tends to a boundary-mismatch term as $\ell \to 0$, so adding the pointwise penalty on top degrades it. BSDE has no boundary term and is drawn as a horizontal level, as is the boundary-regression control: the same network fitted to g on $\partial\Omega$ with no PDE residual at all, i.e. the exact $\alpha_b \to \infty$ limit of every penalized loss. The diffusion loss coincides with that control to within 5% for every $\alpha_b \geq 10^{-1}$, so its apparent insensitivity to the boundary weight is degeneracy rather than robustness: over that range its interior residual contributes nothing. This is a scale effect. At $\alpha_{\text{int}} = \alpha_b = 1$ and $N = 64$ the two terms of the diffusion loss measure $\ell_{\text{int}} = 2.7 \cdot 10^{-5}$ against $\ell_b = 5.6 \cdot 10^{-1}$, because a single increment over the short window ℓ contributes $O(\ell)$ before squaring while the boundary term is $O(\|V^\star\|^2)$; the exit bridge chains m such hops all the way to $\partial\Omega$, which makes its interior residual itself a boundary-datum mismatch and puts it at $\ell_{\text{int}}/\ell_b = 0.99$ with no tuning. The diffusion loss retains one useful setting, $\alpha_b = 10^{-2}$, where it improves on the control by about a third; below it the problem becomes unidentifiable and the error rises monotonically toward the $\alpha_b = 0$ value of 3.7 (not shown).

which holds the barrier at $2E_\star/\sigma^2 = \beta_\star$ independently of every other parameter.

Parameter values. $M = 5$, $\eta = 0.25$, $\varpi = 0.8$, $\beta_\star = 5$, and $(r, s, R) = (1.475,\ 0.25,\ 0.3025)$, giving $E_\star = 4.554$ and $\sigma^2 = 1.822$. The ambient dimension is varied over $d \in \{2, 10, 50, 100\}$ with everything else held fixed.

Exact reduction. Writing $x = (y, z)$ and $\|x - c_m\|^2 = \|y - a_m\|^2 + \|z\|^2$, the factor $e^{-\|z\|^2/2s^2}$ is common to every term of (26), so

$$U(x) = \frac{\|z\|^2}{2s^2} - \log \sum_{m=1}^{M} w_m \exp\Big(- \frac{\|y - a_m\|^2}{2s^2} \Big) \tag{57}$$

splits into an isotropic harmonic well in z and a system in y. The z-marginal is therefore an autonomous Ornstein-Uhlenbeck process independent of Y, and by (27) the sets A, B constrain y alone, so τ_A, τ_B are functions of Y only. Hence

$$V(x) = V_2(y) \qquad \text{exactly, for every } d, \tag{58}$$

with V_2 the committor of the two-dimensional problem. Note the weights do not disturb this: they multiply terms of the sum in y and leave the common z-factor untouched.

Orientation of the plane. The centers are constructed in a plane and then mapped into $\mathbb{R}^d$ by a fixed orthonormal $B \in \mathbb{R}^{d\times 2}$, drawn once from a seed independent of everything else, so the wells span a randomly oriented 2-plane rather than the coordinate plane $\mathrm{span}(e_1, e_2)$. Since the driving noise is isotropic and the potential depends on x only through $B^\top x$ and $\|x\|$, the two are related by a rotation and the notation $x = (y, z) \in \mathbb{R}^2 \times \mathbb{R}^{d-2}$ of (26) is without loss of generality; y is to be read as the component in the plane of the centers. Correspondingly A and B in (27) are cylinders over discs of radius R in that plane, unbounded in the $d - 2$ bath directions. Taking Euclidean balls instead would break (58), since $\|x - c_m\|^2 = \|y - a_m\|^2 + \|z\|^2$ and the bath term then fails to cancel.

Reference and error metric. V_2 is obtained from a five-point finite-difference discretization of (29) in the plane on a 1600×1550 grid over $[-6.06, 6.06] \times [-5.94, 5.80]$, i.e. $h = 7.6 \cdot 10^{-3}$, with a maximum cell Péclet number of 0.42, so no upwinding is applied. It is then lifted to $\mathbb{R}^d$ through (58) and evaluated at 2048 test points drawn from the same interior law the losses sample from. The reported error is $\|\varphi - V\|_{L^2(\nu)} / \|V\|_{L^2(\nu)}$ over those points. There is no Monte Carlo floor, the reduction being exact and the 2-D solve deterministic, but the reference is a discretization of the two-dimensional problem and not a closed form. Each run is summarized by the mean of this error over the last 10% of its evaluations, which is insensitive to the transients that a minimum over training would reward.

Anchors. Interior anchors are drawn on the segment between two centers chosen uniformly at random, jittered by $\mathcal{N}(0, s^2/2)$ per coordinate within the plane and by the bath's own stationary law $\mathcal{N}(0, s^2\sigma^2/2)$ per bath coordinate, which is exact because the bath is the Ornstein-Uhlenbeck process of (57) with rate $1/s^2$. Samples falling in $A \cup B$ are resampled.

Training setup. All methods share a 3-hidden-layer MLP of width 64 with $\tanh$ activations and first-layer weights divided by 0.25, the input rescaling of App. D.1, which the extent of (26)'s domain requires; Adam with $\beta_1 = 0.9$, $\beta_2 = 0.999$, a learning rate of 10^{-3} decayed exponentially to 10^{-5} over the budget, 200 interior and 50 boundary collocation points per step, and a fixed three-hour budget on one GPU. Every cell is run at each method's own ladder of boundary weights α_b (Fig. 7), and the winning weight is then re-run at three seeds for the bridge, the PINN and the BSDE loss and at four for the diffusion loss.

Windows. The bridge and diffusion losses use $n = 250$ steps of $\Delta t = 10^{-4}$, i.e. a window $\ell = n\Delta t = 0.025$, against which $\sqrt{\sigma^2 \ell} = 0.21$ and the well spacing is $r = 1.475$: a single window covers a seventh of the distance between neighboring wells. The bridge loss therefore chains 2 hops per path, routed as described in Fig. 4, and adds an interior-interior term whose endpoint is left unconditioned, so that the residual is enforced in the bulk as well as along the chain; the two interior terms carry equal weight here.

The two arms share this window but not the sampler, and the difference is not only one of implementation. The diffusion loss integrates (26) by Euler-Maruyama over the 250 steps, locating a first exit by bisecting the crossing segment, so its nodes are produced sequentially and its path law is the metastable dynamics itself: paths concentrate inside the wells and cross a barrier only at the rate the

potential allows, which is the event the committor is about. The bridge draws its nodes jointly in closed form from a driftless bridge, so the sampling law is detached from the characteristics of (26) and can place collocation across the barriers directly; the drift is not discarded but enters through the term $b \cdot \nabla\varphi$ of (9). This is unlike the parabolic problem of App. D.4, where $b \equiv 0$ makes the two path laws coincide. The BSDE loss is different in kind, since $Y_\tau = g(X_\tau)$ requires a realized exit: it is run to τ with $\Delta t = 10^{-2}$ and $n = 2.5 \cdot 10^4$, i.e. a horizon $\ell = 250$, at which every trajectory has been absorbed. Shorter horizons do not suffice on this potential – at $\ell = 10$ only 47% of paths have exited, and at $\ell = 100$, 98% – and scoring the unexited remainder against the boundary datum at an interior endpoint is not a valid estimator of (28). Reaching τ costs $2.5 \cdot 10^4$ sequential steps per parameter update, against 250 for the diffusion loss and none for the bridge, whose nodes are drawn jointly; this is the gap visible in the top right panel of Fig. 4.

Insensitivity of the bridge to α_b. The committor has $h \equiv 0$, so Remark D.1 applies here verbatim: the interior residual is invariant under $\varphi \mapsto \varphi + \eta$ for every L-harmonic η, and V is itself L-harmonic, so at $\alpha_b = 0$ any loss whose data enter only through a separate boundary term is unidentifiable. The chain term escapes this because it substitutes the boundary datum for φ at the terminal node of each path, putting g inside the residual itself. That is also why at least two hops are required: with a single hop only 69% of endpoints reach $\partial\Omega$, the rest stopping at a trap center, where g is not defined and the substitution would assert a boundary value in the interior. Intermediate waypoints are the centers of the traps that lie between the anchor's own well and the target, displaced by $\mathcal{N}(0, \varrho^2 I_2)$ within the plane with $\varrho = 2$, i.e. by more than the well spacing $r = 1.475$: routing exactly through the centers would confine every path to a thin tube and the residual is only ever enforced on the support of the sampling law. The terminal waypoint is not jittered, since it must stay on $\partial\Omega$. This mechanism is what Fig. 7 measures: the bridge loss keeps the boundary data as $\alpha_b \to 0$, while the PINN and diffusion losses lose them.

Origin of the parameter values. Dividing a_m, s and R together by a common factor gives $U_\lambda(x) = U_1(\lambda x)$ and hence $V_\lambda(x) = V_1(\lambda x)$, leaving $E_\star$ and σ untouched and changing only the solution's length scale relative to the bridge window $\sqrt{\sigma^2 \ell}$, which is fixed in absolute units. The values above are $(r, s, R) = (5.9, 1, 1.21)$ under that rescaling by 4.

The $M - 2$ interior wells are metastable but absorb nothing, so a path from A to B must cross $M - 1$ barriers in turn and V is a staircase with $M - 1$ risers separated by plateaus, rather than the single riser of a two-state committor. Only y enters (26)-(27) through the centers, so $V(x) = V_2(y)$ exactly for every d; V_2 is computed on a grid and used as a reference with no Monte Carlo floor. The ambient dimension d therefore moves per-step cost and the variance contributed by z alone, leaving the solution fixed.

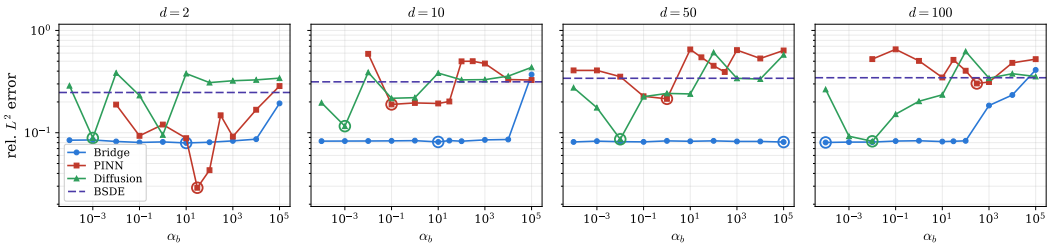

Figure 7: Relative L^2 error against the boundary weight α_b, one panel per ambient dimension, for the chain committor of Section 3.2. Each point is a three-hour cell on one GPU at a single seed; the open marker is that loss's best cell. The BSDE loss carries no boundary term and so is a level rather than a curve. The bridge loss is flat over the six decades $\alpha_b \in [10^{-4}, 10^2]$ at every d, while the baselines vary by a factor of three to eight across the same range and attain their best values at isolated weights; this is the robustness behind the dimension panel of Fig. 4. The α_b grids are not identical across losses, so the curves are comparable only where they overlap.

D.4 High-dimensional HJB

Problem data. We take $d = 100$, $T = 1$, $b \equiv 0$ and $\sigma\sigma^\top = 2\,\mathrm{Id}$, so that $L = \Delta t$ in (2), and $h(x, t, y, z) = -\frac{1}{2}\|z\|^2$; since the third slot of h carries $z = \sigma^\top \nabla V = \sqrt{2}\,\nabla V$, this is $-\|\nabla V\|^2$ and (1a) reduces to (30). The associated process is the driftless $\mathrm{d}X_t = \sqrt{2}\,\mathrm{d}W_t$ started at $X_0 = 0$.

Exact reference by Cole-Hopf. Writing $u = \mathrm{e}^{-V}$ turns the semilinear equation (30) into the heat equation: the quadratic gradient term is exactly the one the exponential change of variables removes, so $\partial_t u + \Delta u = 0$ with $u(T,\cdot) = \mathrm{e}^{-g}$. The Feynman-Kac formula for the driftless process $\mathrm{d}X_t = \sqrt{2}\,\mathrm{d}W_t$ then gives

$$V(t,x) = -\ln \mathbb{E}\big[\mathrm{e}^{-g(x+\sqrt{2(T-t)}\,\xi)}\big], \qquad \xi \sim \mathcal{N}(0,\mathrm{Id}). \tag{59}$$

Reduction to one dimension. For the datum $g(x) = \ln\big(\frac{1}{2}(1+\|x\|^2)\big)$ the integrand is $\mathrm{e}^{-g(y)} = 2/(1+\|y\|^2)$, which depends on y only through $\|y\|^2$. Since $x+\varsigma\xi \sim \mathcal{N}(x,\varsigma^2\mathrm{Id})$ with $\varsigma^2 = 2(T-t)$, the quantity $Q := \|x+\varsigma\xi\|^2/\varsigma^2$ is exactly non-central χ^2 distributed with d degrees of freedom and non-centrality $\lambda = \|x\|^2/\varsigma^2$, so

$$\mathbb{E}\big[\mathrm{e}^{-g(x+\varsigma\xi)}\big] = \mathbb{E}_{Q\sim\chi^2_d(\lambda)}\Big[\tfrac{2}{1+\varsigma^2 Q}\Big], \tag{60}$$

a one-dimensional integral whatever d is. We evaluate it by the substitution $Q = F^{-1}(p)$ with F the non-central χ^2 cdf, i.e. $\mathbb{E}[\phi(Q)] = \int_0^1 \phi(F^{-1}(p))\,\mathrm{d}p$, which keeps the integrand bounded on a fixed interval. Quadrature of the density on $[0,\infty)$ is *not* equivalent in practice: its mass sits far from the origin once $\|x\|$ is not small, and adaptive routines then return a confidently wrong value.

Metric. Following Seo et al. (2026), the error is the mean relative L^2 error over 256 fixed test trajectories sampled at 101 equispaced times, with (59)-(60) supplying the reference at each of the resulting query points. The query points themselves are the driftless paths $X_t = \sqrt{2}\,W_t$ from $X_0 = 0$ on a uniform grid. Note that this reference is deterministic: unlike a per-point Monte Carlo estimate it contributes no variance to the reported error.

Sampling laws. Interior anchors are drawn from the exact forward marginal, $x_0 \sim \mathcal{N}(0,\sigma^2 t_0\,\mathrm{Id})$, with $t_0 \sim \mathcal{U}(0,T-\ell)$ for the windowed residual, so that the window fits before T, and $t_0 \sim \mathcal{U}(0,T)$ for the pointwise PINN residual; terminal points are $x \sim \mathcal{N}(0,\sigma^2 T\,\mathrm{Id})$ at $t = T$. Given (x_0,t_0), the $n = \ell/\Delta t$ increments of a window are drawn jointly and accumulated by a single cumulative sum, with no sequential recursion – this is the simulation-free step. Because $b \equiv 0$ here, the unconditioned window coincides with the path law of the diffusion loss: for a shared random draw the two residuals are pathwise identical (verified to $2\cdot 10^{-7}$), so we pool them and report one curve.

Correspondence with Fig. 5. The curves of Fig. 5 are labeled by loss rather than by reference, so we record the correspondence here. *Bridge / Diffusion* is the windowed residual of Sections 2.1 and 2.2, shown as a single curve because $b \equiv 0$ makes the two path laws coincide. *PINN* is the pointwise residual of Raissi et al. (2019), evaluated with the exact Laplacian. The remaining four are per-step BSDE losses, each integrating to T rather than over a window. *BSDE (per-step)* is the Euler-Maruyama BSDE loss that serves as the common baseline in Park & Tu (2025) and Seo et al. (2026). *BSDE (per-step, Heun)* replaces its integrator by the Heun scheme of Park & Tu (2025). *BSDE (per-step, unbiased)* is the estimator of Seo et al. (2026), which splits each step into $M_1 = M_2 = 5$ independent replicates. *BSDE (per-step, 5-sample)* is the multi-shot variant, which draws $M = 5$ independent one-step replicates and averages their raw errors before squaring. All four are ports of the implementation released with Seo et al. (2026) rather than reimplementations from the papers' prose, so that each baseline is the form its own authors report.

Losses and their settings. The windowed residual uses $\ell = 0.2$ throughout, so refining Δt lengthens the window in steps rather than in time, with interior weight $\alpha_{\mathrm{int}} = 0.1$, terminal weight $\alpha_T = 1$ and batch 512 for both terms; its terminal term matches the value and the gradient. The per-step BSDE baselines run at $\Delta t = 10^{-2}$, hence $T/\Delta t = 100$ steps per trajectory, with $\alpha_{\mathrm{int}} = \alpha_T = 1$, batch 512, 5 replicates for the 5-sample variant and $5{+}5$ for the unbiased one; Heun runs at batch 64. PINN evaluates the exact Laplacian, which costs d second-derivative passes per collocation point, with interior batch 512, terminal batch 50 and a terminal term that matches the value only; its interior weight $\alpha_{\mathrm{int}} = 100$ was selected by a sweep, in contrast to the default weights used everywhere else.

Training protocol. All runs use the same network – a four-layer MLP of width 512 with Mish activations and no skip connections – and Adam with $\mathrm{lr} = 10^{-3}$, $\beta = (0.9, 0.999)$, $\varepsilon = 10^{-8}$,

annealed to zero by a cosine schedule over a step horizon fixed in advance for each run. Comparisons are matched on wall clock rather than on iterations: each configuration is given $8\,\mathrm{h}$ on one H200, and is preceded by a $1200\,\mathrm{s}$ calibration that measures its own throughput, from which the cosine horizon is set to 93% of the projected step count. Where gradient clipping is used it is applied to the global norm, scaling the parameter vector by $\min(1, c/\|g\|_2)$ before the Adam update. Each configuration is run with three seeds, so six for the pooled bridge/diffusion curve; the intermediate Δt values of the right panel are single runs.

D.5 Scope and open directions

We close by recording what the construction assumes and what the experiments do and do not settle. The residual identity holds for any admissible sampling process, but closed-form sampling rests on $\sigma(x,t) = \sigma(t)$: for a state-dependent diffusion coefficient the quadratic variation (8) forces Z back onto an SDE, and whether a simulation-free sampler exists in that case is open. On a bounded domain the exit bridge may leave and re-enter Ω, so the boundary substitution is exact up to the excursion effect of App. C.3; anchor placement makes it rare, and rejection or truncation removes it at a cost in samples or in geometric queries. The analysis of App. C is asymptotic in Δt, and the two modifications of App. C.2 are derived but not measured here, which leaves their practical benefit an open question. The window length ℓ and the anchor law are design choices rather than derived quantities; App. C.3 ties them to the geometry of the problem, which is a rule of thumb and not an optimum.