FROM SPARSE SENSORS TO FULL-FIELD POSTERIORS: PHYSICS-PRIOR BRIDGES FOR STOCHASTIC PDES

Anonymous authors
Paper under double-blind review

ABSTRACT

Reconstructing physical fields from fixed sparse sensors is challenging when stochastic forcing creates unobserved variations throughout the evolution. We propose Physics Reference Diffusion Bridge (PRBridge), a framework for physical field posterior reconstruction given known stochastic dynamics, and sparse observation trajectories, without full-field training targets. PRBridge models observation-conditioned trajectories in physical time by calibrating a physics-informed Gaussian reference, which enables tractable path densities and unbiased local estimates of path-space KL gradients. Long-window ensemble supervision provide the auxiliary training through evolving cross-covariances between sensor and fields, helping exploit dynamical information about unmeasured structures. A second stage distills the learned terminal distribution into a one-step stochastic sampler. Empirically, we showed the effectiveness of our PRBridge through the experiments on three synthetic stochastic PDE benchmarks and real cylinder-flow measurements. This highlights the insights of PRBridge on providing a route to probabilistic scientific inference from sparse measurements and reduce reliance on fully observed training trajectories.

1 INTRODUCTION

Reconstructing spatially resolved physical fields from sparse sensor measurements support applications across scales, from characterizing intracellular biochemical waves (Martinez-Torres et al., 2026; Mata et al., 2013) to estimating large-scale geophysical flows (Cotter et al., 2020) and weather forecasting (Zhang et al., 2025). With sensors deployed at fixed locations, continuous measurements reveal temporal changes while leaving much of the spatial domain unobserved. Yet stochastic forcing and perturbations leave uncertainty along the evolution even when the governing dynamics like partial differential equations (PDEs) and initial conditions are known. Here, we consider systems governed by known stochastic PDEs and initialized from a known field. For systems governed by stochastic PDEs, the central challenge is to exploit the observation and dynamics to reconstruct reliable spatial structures in unobserved regions while retaining the uncertainty.

Physics-informed learning provides solutions by incorporating governing equations as residual penalties into reconstruction objectives (Raissi et al., 2019; He et al., 2020). However, such deterministic residual-based training objectives alone do not specify the conditional path distribution induced by stochastic forcing. Diffusion-based generative methods offer distributional reconstructions (Ho et al., 2020; Song et al., 2020; Zhou et al., 2024; Bastek et al., 2025), but representative approaches typically rely on full-field training data (Li et al., 2025; Huang et al., 2024). Learning from incomplete data can relax this requirement, but the recovery guarantee depends on training coverage that excludes permanently unobserved locations (Zhou et al., 2026). Recent works reduced setting of dense supervisions (Bu et al., 2026), but the potential challenges arise on the recovery of unmeasured fine-scale structures without exploiting observation histories and long-range spatiotemporal dependencies (Zaki & Wang, 2021; Du et al., 2023) (more related works in Appendix.A). Physical-time diffusion bridges provide a complementary route by conditioning stochastic dynamics, but neural guided bridge training requires differentiation through entire rollout trajectories resulting in training costs grow with the number of integration steps (Yang et al., 2025). These considerations motivate three research questions:

Q1: How can we approximate the physical path posterior from stochastic dynamics and fixed sparse observations, without dense supervision beyond initial conditions? **Q2:** How can we exploit long-range spatiotemporal dependencies to improve the reconstruction of unmeasured fine-scale struc-

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.

tures? **Q3:** How can we amortize bridge inference into an accurate, few even single-step terminal field posterior sampler, while keeping per-instance processing cost low?

To answer above three questions, we propose **Physics Reference Diffusion Bridge (PRBridge)**, a framework for terminal-field posterior reconstruction under fixed sparse sensing, without full-field supervision after initialization. PRBridge defines its bridge over physical time, with intermediate states representing the evolving physical field under stochastic PDE (SPDE), featuring three insights:

Insight #1: Local learning of a globally conditioned path posterior. We introduce a Gaussian-reference parameterization that combines conditioning on the full observation history with local optimization of the physical path-space KL objective. This enables local gradient computation without backpropagation through full physical-time trajectories.

Insight #2: Stochastic dynamics as a source of distributional supervision. We incorporate a long-window ensemble forecast–analysis into auxiliary distributional targets, using known physics to propagate sparse observation information into unmeasured field regions. This provides explicit training signals for dynamics-mediated correlations without requiring full-field supervision.

Insight #3: Terminal inference through compressed bridge distillation. We formulate endpoint reconstruction through a learned conditional transition. This enables single-step terminal-field sampling without the need for full-horizon preparation at inference time.

We evaluate PRBridge on both simulated benchmarks and real-world datasets, assessing reconstruction quality and end-to-end inference efficiency under fixed sparse sensing.

2 PRELIMINARIES

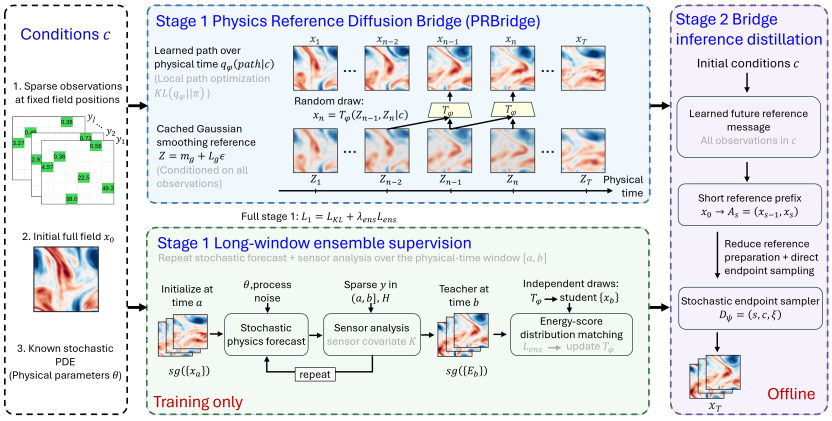


Figure 1: **Overview of PRBridge for sparse SPDE reconstruction. PRBridge localized posterior sampling (top).** PRBridge calibrates joint physics-informed Gaussian smoothing reference samples into calibrated paths through a triangular transformation indexed by physical time. **PRBridge long-windoe ensemble supervision** To better capture the spatial-temporal dependencies of the observations. **Stage 2 (right)** We further proposed a bridge inference distillation stage enabling few even single step sampling from the initial condition to the endpoint at inference time.

Problem setting. As shown in Figure 1, let $x_t \in \mathbb{R}^d$ denote a spatially discretized physical field and we consider the known stochastic dynamics

$$\mathrm{d}x_t = b_\theta(x_t, t)\,\mathrm{d}t + B_\theta(x_t, t)\,\mathrm{d}W_t, \qquad x_{t=0} = x_0, \tag{1}$$

where W_t is a standard Wiener process and B_θ incorporates the prescribed noise amplitude and spatial covariance, allowing spatially correlated forcing. The physical parameters θ and the complete initial field x_0 are known, while the realized process noise is unobserved.

On a physical-time grid $0 = t_0 < \cdots < t_N = T$, observations are collected at $J \le N$ selected times $s_j = t_{n_j}$, where $1 \le n_1 < \cdots < n_J = N$:

$$y_j = Hx_j + \varepsilon_j, \qquad \varepsilon_j \sim \mathcal{N}(0, \Gamma_j), \tag{2}$$

with $\Gamma_j \succ 0$ and measurement errors independent of one another and of process noise. The sparse sensor operator $H \in \mathbb{R}^{m\times d}$, with $m \ll d$, is fixed across time and training and test instances. Let $c = (x_0, \theta, H, y)$ denote the condition, including the initial condition and sparse observations. Given c, our goal is to reconstruct the full terminal field by learning its observation-conditioned distribution. In sparse-only training, full-field ground-truth labels after initialization are not available.

Observation-conditioned diffusion bridges. The Gaussian measurement model in equation 2 induces the likelihood $\ell_j(x) := p(y_j \mid x_{s_j} = x) = \mathcal{N}(y_j; Hx, \Gamma_j)$. Conditional independence of the measurement errors gives a product of likelihoods along a trajectory. Marginalizing the unobserved future states under the physical dynamics yields the backward information function $h(t, x; c) = \mathbb{E}_{\text{phys}}\left[\prod_{j:s_j>t} \ell_j(x_{s_j}) \,\middle|\, x_t = x\right], h(T, x; c) = 1$, which represents the likelihood of future measurements given the current state. Across observation times from s_j^- to s_j, the backward message satisfies $h(s_j^-, x; c) = \ell_j(x) h(s_j, x; c)$. Let $\mathcal{Q}_{s,t}(x, \mathrm{d}x'; c)$ denote the unnormalized Feynman–Kac transition kernel of the physical process, weighted by the observation likelihoods $\prod_{j:s<s_j\le t} \ell_j(x_{s_j})$. The observation-conditioned transition is its time-dependent Doob transform (Heng et al., 2020; Mider et al., 2021):

$$K^\star_{s,t}(x, \mathrm{d}x'; c) = \frac{h(t, x'; c)}{h(s, x; c)} \mathcal{Q}_{s,t}(x, \mathrm{d}x'; c). \tag{3}$$

The backward identity $h(s, x; c) = \int \mathcal{Q}_{s,t}(x, \mathrm{d}x'; c) h(t, x'; c)$, which follows from the Markov property, ensures that $K^\star_{s,t}$ is normalized. Let $a_\theta(x, t) = B_\theta(x, t) B_\theta(x, t)^\top$. Under standard regularity conditions, h defines a time-dependent Doob transform on each interval between observations (Doob, 1957; Mider et al., 2021), and the state process therefore satisfies

$$\mathrm{d}x_t = \Big[b_\theta(x_t, t) + a_\theta(x_t, t) \nabla_x \log h(t, x_t; c)\Big] \,\mathrm{d}t + B_\theta(x_t, t) \,\mathrm{d}W^\star_t, \tag{4}$$

where $W^\star_t$ is Brownian motion under the conditioned law. The additional drift incorporates future observation information, weighted by the physical process-noise covariance. Past observations are reflected in the conditioned distribution of the current state. Thus, backward propagation of observation information guides state evolution forward in physical time.

3 Physics Reference Diffusion Bridge

In Fig. 1, PRBridge learns a physical-time posterior sampler in two stages. Stage 1 is based on local physics-referenced variational bridge learning and long-window ensemble supervision for capturing the dependencies between sensor observations and stochastic dynamics. These two complementary objectives jointly train the sampler. Stage 2 distills the learned bridge into a stochastic endpoint sampler and compresses future reference information to enable efficient single-step inference.

3.1 Localized Physics-referenced posterior sampling

Posterior sampling formulation. The observation-conditioned bridge in equation 4 characterizes the posterior process, but exact simulation requires the generally intractable backward information function h. We instead construct a variational sampler for the posterior of a fixed numerical discretization. Let $p_n(x_n \mid x_{n-1})$ denote its physical transition density and $g_n(x; c) = \prod_{j:n_j=n} \ell_j(x)$, with $g_n = 1$ at unobserved times. Conditioning on the observations gives

$$\pi(x_{1:N} \mid c) \propto \prod_{n=1}^{N} p_n(x_n \mid x_{n-1}) \, g_n(x_n; c), \tag{5}$$

where x_0 is fixed. Our reconstruction target is its terminal marginal $\pi_N(x \mid c)$. Therefore, we introduce a learnable conditional path $q_\phi(x_{1:N} \mid c)$ with terminal marginal $q_{\phi,N}$. Learning the full path provides a variational route to terminal reconstruction.

Proposition 3.1 (Terminal KL control)**.** *For a fixed condition c, assume that $q_\phi(x \mid c)$ is absolutely continuous with respect to $\pi(x \mid c)$ and has finite KL divergence. Then*

$$\mathrm{KL}(q_{\phi,N}(x \mid c) \,\|\, \pi_N(x \mid c)) \le \mathrm{KL}(q_\phi(x \mid c) \,\|\, \pi(x \mid c)) \tag{6}$$

We therefore minimize the expected path KL, $\mathcal{L}_{\mathrm{KL}}(\phi) := \mathbb{E}_c[\mathrm{KL}(q_\phi(x_{1:N} \mid c) \| \pi(x_{1:N} \mid c))]$. The factorization in equation 5 gives the following objective.

Theorem 3.2 (Path-space variational objective)**.** *Under the support and integrability assumptions,*

$$\mathcal{L}_{\mathrm{KL}}(\phi) = \mathbb{E}_{c,\, x_{1:N} \sim q_\phi(\cdot|c)} \left[\log q_\phi(x_{1:N} \mid c) - \sum_{n=1}^{N} \log(p_n(x_n \mid x_{n-1}) g_n(x_n; c))\right] + C_{\text{phys}}, \tag{7}$$

where C_{phys} is independent of ϕ. The expectation is the negative evidence lower bound (ELBO), and minimizing it is equivalent to minimizing $\mathcal{L}_{\mathrm{KL}}$.

The derivation is provided in Appendix B.1. The objective combines physical transition likelihoods, observation likelihoods, and the learned path log-density, without evaluating the backward information function h. The constant C_{phys} can be omitted during optimization.

Physics-informed Gaussian reference sampling Motivated by the use of tractable forward sampling processes in diffusion models and diffusion bridges (Ho et al., 2020; Song et al., 2020; Mider et al., 2021; Zhou et al., 2024), we introduce a physics-informed Gaussian reference smoothing along the time steps. Given a numerical deterministic update like discrete PDE f_n, and a first-order construction around a nominal trajectory $\bar{x}$, there are $A_n^g = Df_n(\bar{x}_{n-1})$ and $\beta_n^g = f_n(\bar{x}_{n-1}) - A_n^g \bar{x}_{n-1}$. Its mean is obtained by smoothing a linearized physical model using the known initial condition and the sparse observations, and a temporally correlated covariance supplies stochastic fluctuations:

$$\widetilde{p}_n(z' \mid z) = \mathcal{N}(z'; A_n^g z + \beta_n^g, Q_n^g), \qquad z_0 = x_0. \tag{8}$$

where surrogate innovation covariance Q_n^g captures a local approximation to the physical evolution (Mider et al., 2021). For Gaussian reference sampling, we retain this mean and prescribe a structured, temporally correlated covariance:$q_g(z_{1:N} \mid c) = \mathcal{N}(m_g(c), \Sigma_g(c))$. Here Σ_g is the designed sampling covariance, rather than the exact smoothing covariance of Equation equation 8. After reference preprocessing, a temporal window n can be sampled directly as $z_n = m_{g,n} + L_n \xi, L_n L_n^\top = \Sigma_{g,n,n}, \xi \sim \mathcal{N}(0, I)$. Local flow models such as nMA use independent Gaussian base variables and finite receptive fields (Ryder et al., 2021). This physics-informed gaussian smooth surrogate serves two advantages: First, the smoothing mean provides a coarse trajectory estimate on both physics laws and observations. Second, Gaussian marginal enables tractable sampling at selected times after reference preparation, without generating the preceding trajectory for temporal dependencies beyond the local window. We then learn an invertible transport through model T_ϕ to calibrate this reference against the physical posterior through the path-KL objective in equation 7 as in Fig.1.

Local calibration with an explicit path density. The Gaussian reference is a coarse approximation and does not generally reproduce the nonlinear physical posterior. We therefore learn a transport $x_{1:N}$ from$(z_{1:N}; c)$ to calibrate its distribution. For an invertible differentiable transport (Rezende & Mohamed, 2015; Ryder et al., 2021),

$$\log q_\phi(x_{1:N} \mid c) = \log q_g(z_{1:N} \mid c) - \log |\det D_z T_\phi(z_{1:N}; c)| . \tag{9}$$

However, an unrestricted path transformation would make the high-dimensional Jacobian and full-path evaluation costly. We instead take the local affine parameterization trick:

$$x_n = T_{\phi,n}(z_{n-1}, z_n; c) = \exp\big(s_{\phi,n}(z_{n-1}, c)\big) \odot z_n + u_{\phi,n}(z_{n-1}, c), \tag{10}$$

with $x_0 = z_0$. The scale and shift networks may use the full preceding reference field and deterministic observation features, but not the current input z_n. This makes the map triangular and invertible while keeping its Jacobian block lower bidiagonal, with diagonal blocks $\mathrm{diag}(\exp s_{\phi,n})$. Consequently, we can remove the dense path Jacobian or a stochastic log-determinant approximation by:

$$\log |\det D_z T_\phi| = \sum_{n=1}^{N} \sum_{i=1}^{d} s_{\phi,n,i}(z_{n-1}, c). \tag{11}$$

Local optimization of the global path objective. Using equation 9, equation 11, and equation 7:

$$\mathcal{L}_{\mathrm{local}}(\phi) = -\mathbb{E}_{c,\, z \sim q_g(\cdot|c)} \left[\sum_{n=1}^{N} \log p_n(x_n \mid x_{n-1}) + \sum_{n=1}^{N} \log g_n(x_n; c) + \sum_{n=1}^{N} \sum_{i=1}^{d} s_{\phi,n,i} \right]. \tag{12}$$

The reference log-density contributes only a parameter-independent expectation, so $\mathcal{L}_{\mathrm{KL}} = \mathcal{L}_{\mathrm{local}} + C_{\mathrm{ref}}$ for a constant C_{ref} independent of ϕ. The objective combines stochastic physical consistency, sensor agreement, and the transport's contribution to path entropy.

Proposition 3.3 (Unbiased local path-KL gradients)**.** *Assume the reference is independent of ϕ and that differentiation and expectation can be interchanged. For $n \geq 2$, $F_{\phi,n}$ depends only on (z_{n-2}, z_{n-1}, z_n). Sample n uniformly from $\{1, \ldots, N\}$ and draw its joint window from q_g. Then*

$$\widehat{\mathcal{L}} = N F_{\phi,n}, \qquad \mathbb{E}\Big[\nabla_\phi \widehat{\mathcal{L}}\Big] = \nabla_\phi \mathcal{L}_{\mathrm{KL}}. \tag{13}$$

At $n = 1$, only z_1 is random because $z_0 = x_0$. $F_{\phi,n}$ denotes the local contribution of $F_{\phi,n} := -\log p_n(x_n \mid x_{n-1}) - \log g_n(x_n; c) - \sum_{i=1}^{d} s_{\phi,n,i}(z_{n-1}, c)$,

The result follows from the additive objective and the local dependence of the transport (Appendix B.2). After reference preparation, training process only uses short reference windows without backpropagating through complete physical rollouts. Terminal samples are obtained directly from $x_N = T_{\phi,N}(z_{N-1}, z_N; c)$ using the final joint reference pair.

3.2 LONG-WINDOW ENSEMBLE SUPERVISION

Dynamic visibility under fixed sensors. Physical evolution can transmit information about unobserved fine-structure structures to subsequent sensor measurements (Zaki & Wang, 2021). We characterize this connection through the observation and terminal responses to a state perturbation.

Definition 3.4 (Dynamic visibility and target recoverability)**.** Fix an intermediate time t_a, with $0 < a < N$. For a common realization of subsequent process noise, let $x_n(u)$ denote the physical numerical evolution initialized at $x_a = u$. Define

$$G_a(u) := \mathrm{col}_{j:n_j>a}\left(Hx_{n_j}(u)\right), \qquad R_a(u) := P_B x_N(u), \tag{14}$$

where P_B selects unobserved coordinates. Assuming differentiability at x_a, write $J_G := DG_a(x_a)$ and $J_R := DR_a(x_a)$. For an admissible state perturbation v, $J_G v$ and $J_R v$ are its first-order observation and terminal-target responses. A perturbation with $Hv = 0$ is dynamically visible to first order if $J_G v \neq 0$. The terminal target is first-order recoverable from the subsequent observations if

$$J_G v = 0 \quad \Longrightarrow \quad J_R v = 0 \qquad \text{for every admissible } v. \tag{15}$$

The response $J_R v$ may describe changes in fine-scale structures, such as wave-front displacements or vortical filaments. When $J_G v \neq 0$, these changes leave a signature in the observation history. When $J_G v = 0$ but $J_R v \neq 0$, the corresponding terminal variation is invisible to first order.

From dynamic sensitivity to long-window ensemble supervision. A perturbation outside sensor locations can become observable after physical evolution. Exploiting this information requires learning how changes in the unobserved field affect measurements over time. Although the reference and path-KL objective incorporate the full observation history, finite-budget local updates may not fully capture these relationships. We therefore construct auxiliary training targets through long-window ensemble forecasting and analysis to exploit spatial-temporal sensor correlations for reconstruction.

Choose a physical-time window $[t_a, t_b]$, with t_b an observation time. Using a frozen parameter $\bar{\phi}$ during target construction, draw M independent initial members from $q_{\bar{\phi},a}(\cdot \mid c)$. Relabel the observations within the window by $\ell = 1, \ldots, J_{a:b}$, with local index $\ell = 0$ corresponding to t_a. Between consecutive observations, propagate each member m through the stochastic physical solver:

$$x_\ell^{\mathrm{f},(m)} = \Phi_{\Delta t_\ell}\left(x_{\ell-1}^{\mathrm{a},(m)}; \eta_\ell^{(m)}\right), \tag{16}$$

where $\eta_\ell^{(m)}$ includes the physical process-noise increments across all solver substeps. At each observation, apply an ensemble Kalman analysis (Evensen, 2003; 2009):

$$K_\ell = C_\ell^{xy}\left(C_\ell^{yy} + \Gamma_\ell\right)^{-1}, \qquad x_\ell^{\mathrm{a},(m)} = x_\ell^{\mathrm{f},(m)} + K_\ell\left(y_\ell + \varepsilon_\ell^{(m)} - Hx_\ell^{\mathrm{f},(m)}\right), \tag{17}$$

where $\varepsilon_\ell^{(m)} \sim \mathcal{N}(0, \Gamma_\ell)$, and C_ℓ^{xy} and C_ℓ^{yy} are the forecast state–sensor and sensor–sensor sample covariances. Each analyzed ensemble initializes the next forecast, so a sensor's influence changes with evolving physical structures. The final analyzed members define detached targets $\widetilde{x}_b^{(m)}$. These are approximate dynamical pseudo-targets constructed without post-initial full-field ground truth.

Distributional training objective. During the training process, we independently draw $L \geq 2$ student samples from $x_b^{(i)} \sim q_{\phi,b}(\cdot \mid c)$ using valid joint Gaussian reference pairs. With $\mathsf{d}(u, v) = \|u - v\|_2/\sqrt{d}$ on the free field coordinates, define

$$\widehat{\mathcal{L}}_{\mathrm{ens}} = \frac{1}{LM}\sum_{i=1}^{L}\sum_{m=1}^{M}\mathsf{d}\left(x_b^{(i)}, \widetilde{x}_b^{(m)}\right) - \frac{1}{2L(L-1)}\sum_{i\neq i'}\mathsf{d}\left(x_b^{(i)}, x_b^{(i')}\right). \tag{18}$$

This energy-score objective provides distributional supervision without arbitrarily pairing student and target samples (Gneiting & Raftery, 2007). Thus, the complete PRBridge Stage 1 objective is

$$\mathcal{L}_1(\phi; \bar{\phi}) = \mathcal{L}_{\mathrm{KL}}(\phi) + \lambda_{\mathrm{ens}}\,\mathbb{E}\left[\widehat{\mathcal{L}}_{\mathrm{ens}}(\phi; \bar{\phi})\right], \tag{19}$$

Gradients propagate through the student samples and the original KL objective; target initialization, physical forecasts, and ensemble analyses are detached.

Sparse observation reconstruction bound. The following result retains terminal variations that are not constrained by the linearized observations.

Theorem 3.5 (Sensor-dependent posterior reconstruction)**.** *Fix c and a generating noise history $\eta^\star$, with stacked observations $y = G(\eta^\star) + \varepsilon$ and $\|\varepsilon\|_2 \leq \delta$. Assume that G and R are twice*

continuously differentiable on the radius-ρ ball around $\eta^\star$, with second-derivative operator norms bounded by M_G and M_R.

Evaluate J_G and J_R at $\eta^\star$ and define

$$\kappa := \left\| J_R J_G^\dagger \right\|_{\mathrm{op}}, \qquad \beta := \left\| J_R \left(I - J_G^\dagger J_G \right) \right\|_{\mathrm{op}}, \tag{20}$$

where $J_G^\dagger$ is the Moore-Penrose pseudo-inverse, and the derivatives act on the admissible noise coordinates.

Let $\eta(x)$ denote the additive noise history associated with a numerical path, and define

$$\mathcal{A}_{\rho,\tau} := \{x_{1:N} : \|\eta(x) - \eta^\star\|_2 \le \rho,\ \|G(\eta(x)) - y\|_2 \le \tau\}. \tag{21}$$

If $\pi(\mathcal{A}_{\rho,\tau} \mid c) \ge 1 - \gamma$, then any learned path law with finite $\mathrm{KL}(q_\phi(x \mid c) \| \pi(x \mid c))$ satisfies

$$\mathbb{P}_{x_{1:N} \sim q_\phi(x|c)} \left(\|P_B x_N - R(\eta^\star)\|_2 > r_{\rho,\tau} \right) \le \gamma + \sqrt{\frac{1}{2} \mathrm{KL}(q_\phi(x \mid c) \| \pi(x \mid c))}, \tag{22}$$

where $r_{\rho,\tau} = \kappa(\tau + \delta) + \beta\rho + \frac{1}{2}\left(M_R + \kappa M_G\right)\rho^2$.

κ measures amplification of observation discrepancies into terminal-field error. $\beta\rho$ accounts for terminal variations in the null space of the linearized observation map; it vanishes under first-order target recoverability. The remaining terms describe nonlinear effects, posterior mass outside the local region, and learned posterior approximation error. The proof is provided in Appendix C.

Remark (Interpretation and limitations). The reconstruction radius can be large when target-sensitive perturbations produce weak sensor responses (large κ), when terminal variations remain invisible to first order (large $\beta\rho$), or when nonlinear effects are substantial (large $(M_R + \kappa M_G)\rho^2$). The probability guarantee also weakens when the physical posterior is poorly localized (large γ) or the learned path law poorly approximates it.

3.3 One-step posterior bridge distillation

Stage 1 relies on Gaussian reference preparation, while long-window ensemble supervision requires physical forecasts and observation analyses during training. We distill the reconstruction distribution into a direct conditional generator, avoiding deployment-time reference preparation. The finite-grid counterpart of the conditioned kernel in Equation equation 3 characterizes the terminal posterior as

$$\pi_N(\mathrm{d}x_N \mid c) = K^\star_{0,N}(x_0, \mathrm{d}x_N; c) = \frac{\mathcal{Q}_{0,N}(x_0, \mathrm{d}x_N; c)}{h_0(x_0; c)}. \tag{23}$$

This motivates learning a single stochastic map across the full physical-time horizon, conditioned on the known initial field and observation history. For stage 2 training, we freeze a Stage 1 teacher with parameters $\bar{\phi}$ and draw terminal targets $V \sim q_{\bar{\phi},N}(x \mid c)$. The student generates $Y = D_\psi(c, \xi)$ and $Y' = D_\psi(c, \xi')$, where $\xi, \xi' \sim \mathcal{N}(0, I)$ are independent of each other and of V conditional on c. Using a fixed, scaled Euclidean field distance d, we minimize the symmetric energy score

$$\mathcal{L}_{\mathrm{dist}}(\psi) = \mathbb{E}\left[\frac{1}{2}\mathsf{d}(Y, V) + \frac{1}{2}\mathsf{d}(Y', V) - \frac{1}{2}\mathsf{d}(Y, Y') \right]. \tag{24}$$

The teacher targets are detached, and both student draws receive gradients. Under finite first moments, the expected score is minimized over distributions by matching the teacher.

At inference, the student takes $c = (x_0, \theta, H, \mathcal{C})$ and fresh Gaussian noise, producing a complete terminal field in one forward evaluation. Neither Gaussian reference preparation nor ensemble forecasting and analysis is required. Repeated evaluations with independent noise provide an ensemble of terminal reconstructions conditioned on the same initial field and sparse observation history.

4 Experiments Setup

Benchmarks. We evaluate terminal-field reconstruction on synthetic stochastic PDE benchmarks and real-world Cylinder PIV measurements from RealPDEBench (Hu et al., 2026). Under the protocols on stochastic simulation with SPDEBench (Zhu et al., 2025), we expanded the synthetic dataset with one-dimensional Burgers, two-dimensional FitzHugh-Nagumo (FHN), and two-dimensional Navier-Stokes (NS). Each equation is simulated with additive, spatially correlated noise on a periodic domain. We generate 10,000 trajectories per equation and partition them by initial condition into training, validation, and test sets, preventing the leakage of initial conditions. For real Cylinder PIV measurements (Hu et al., 2026), we use recording-level splits for both in-distribution evaluation and Reynolds-number extrapolation. Details are provided in Appendix D.

Observation protocol. For each setting, a single spatial sensor operator H is fixed across time, instances, and data splits. The synthetic benchmarks separately use 8, 16, and 32 point sensors for Burgers and 64 (8×8, 0.39% observation), 256 (16×16, 1.56% observation), and 1024 (32×32, 6.25% observation) uniformly distributed sensor lattices for FHN and NS. Measurements are available at all 100 non-initial saved trajectory frames, with independent Gaussian observation noise of standard deviation 0.05 in physical units. FHN additionally includes nonuniform sensor distribution with a permanently unobserved 64×64 central region. For Cylinder, we use 28, 107, or 432 spatial sensors on the real trajectories, each observing both velocity components. The reconstruction task supplies the complete initial frame and the subsequent sparse observation history.

Baselines. We compare PRBridge with numerical data-assimilation methods and deep learning-based reconstruction models. Under our experimental protocol, the sparse-only group includes the Gaussian smoother, ensemble Kalman filter (EnKF) (Evensen, 2009), nMA (Ryder et al., 2021), SDE Matching(Bartosh et al., 2025), NGDB (Yang et al., 2025), and PICSB (Bu et al., 2026). The dense-supervised group includes DiffusionPDE (Huang et al., 2024) and SDA (Rozet & Louppe, 2023), which additionally use full-field simulated data for training. We also evaluate a dense-trained PRBridge version to examine the effect of additional supervision. All methods receive fixed sparse measurements at inference. The main comparison reports PRBridge's full Stage 1 sampler, trained with the local path-space KL objective and long-window ensemble supervision. Stage 2 is assessed separately through matched teacher–student comparisons.

Evaluations. The main comparison reports the full Stage 1 sampler; Stage 2 distillation is evaluated separately through teacher–student comparisons. Each probabilistic method produces $M = 32$ terminal-field samples per test instance unless otherwise specified. We report full-field RMSE and evaluate relative L_2 error, CRPS (Ferro, 2014), and 90% interval coverage (Cov90) at non-sensor locations. Interval width is reported in the distributional diagnostics (see Appendix E).

5 RESULTS

Main comparisons. Table 1 compares reconstruction across three stochastic PDEs and real Cylinder flow. PRBridge achieves the lowest reported RMSE and relative L_2 error within the sparse-only group across all six main-table settings. The advantage also extends to real Cylinder measurements, where PRBridge obtains RMSE 0.2295 and relative L_2 error 0.3076 with only 28 sensors.

The accuracy gains are accompanied by competitive, but not uniformly superior, probabilistic predictions. PRBridge attains the lowest sparse-only CRPS on FHN with 256 sensors, whereas EnKF obtains a lower CRPS on NS with 64 sensors. For the latter setting, PRBridge's Cov90 is closer to the nominal level. Undercoverage nevertheless remains particularly on FHN. Additional full-field supervision reduces PRBridge's point-estimation errors, but does not uniformly improve CRPS.

Ablation studies. Table 2 examines FHN with 256 sensors and NS with 64 sensors. Local KL training improves reconstruction over the Gaussian reference on FHN, but does not improve the reported errors on NS. Long-window ensemble supervision provides additional gains on both systems. The reported short-window ensemble results lie between local KL-only training and Full Stage 1. These observations support using extended physical evolution and sequential observation analysis to supplement local variational learning. One-step distillation retains similar relative L_2 errors but degrades uncertainty coverage. Thus, the distilled stage 2 still lags behind its teachers' distributional quality for reconstructing the coverage and variance. Finally, dense supervision reduces RMSE, indicating remaining benefits of full-field training information.

Further evaluations Fig 2a compares terminal vorticity reconstructions for stochastic Navier-Stokes. The Gaussian reference captures the broad spatial patterns but yields diffuse vortex boundaries and blurry structures. Local KL training improves some local features for sampling, although distortions in vortex geometry and fine-scale structures remain. Adding long-window ensemble supervision produces reconstructions that more closely reproduce the shapes and spatial arrangement of vortical structures and several thin vorticity filaments. The remaining discrepancies at the finest scales also indicate that reconstruction from fixed sparse observations remains challenging under extreme sparse observation rate like 0.39%. For Navier-Stokes with 64 sensors, Figure 2b compares empirical and nominal interval coverage at non-sensor locations. Both the Gaussian reference and local KL-only Stage 1 exhibit undercoverage.

Figure 3 examines complementary aspects of the learned posterior uncertainty and sampling diversities. We examine the spatial and temporal structure of posterior variability on FHN with 256

Table 1: Reconstruction with fixed sparse sensors. Full RMSE, relative L_2, CRPS, and Cov90 over non-sensor locations. Lower errors are better; Cov90 is in percent with target 90. Bold and underline mark the best and second-best values, respectively, within each supervision group and experimental setting. Sparse only means no dense training fields and includes numerical methods. All methods use sparse observations at inference. Runtime columns refer to the right-hand setting.

(a)		**FHN: 64 sensors 0.39%**				**FHN: 256 sensors 1.56%**					
	Method	RMSE	Rel. L_2	CRPS	Cov90	RMSE	Rel. L_2	CRPS	Cov90	Train (GPU·h)	Infer (s/inst, $M = 32$)
Sparse only	Gaussian smoother	0.1714	0.5399	0.0942	63.8	0.0683	0.2156	0.0591	69.9	-	2.272
	EnKF	0.1639	0.5162	0.0828	75.3	0.0573	0.1769	0.0276	78.4	-	14.384
	nMA	0.3086	0.9702	0.2607	11.4	0.1242	0.3928	0.0937	17.3	1.12	0.03
	SDE Matching	0.2302	0.7212	0.1734	15.7	0.2252	0.7052	0.1680	14.6	2.11	0.075
	NGDB	0.1527	0.4841	**0.0749**	**84.9**	0.0515	0.1721	0.0308	**87.3**	30.30	1.184
	PICSB	0.2020	0.6320	0.1341	21.3	0.0899	0.2826	0.0561	39.9	1.05	0.136
	PRBridge (Sparse)	**0.1513**	**0.4814**	0.0759	75.6	**0.0483**	**0.1703**	**0.0250**	77.8	2.76 / 1.88	2.29 / 0.029
Dense supervision	DiffusionPDE	0.2462	0.7716	0.1313	81.5	0.0680	0.2199	0.0577	82.2	1.89	16.982
	SDA	0.2685	0.7621	0.0984	**92.6**	0.0784	0.2457	0.0549	83.2	2.89	15.176
	PalSB	0.2045	0.5388	0.0942	34.4	0.0626	0.2024	0.0593	48.9	1.94	0.148
	PRBridge (Dense)	**0.1490**	**0.4674**	**0.0751**	85.3	**0.0414**	**0.1417**	**0.0252**	**83.6**	2.59	2.284

(b)		**NS: 64 sensors 0.39%**				**NS: 256 sensors 1.56%**					
	Method	RMSE	Rel. L_2	CRPS	Cov90	RMSE	Rel. L_2	CRPS	Cov90	Train (GPU·h)	Infer (s/inst, $M = 32$)
Sparse only	Gaussian smoother	0.5745	0.4048	0.3204	67.3	0.2975	0.2104	0.2062	69.1	-	11.39
	EnKF	0.4902	0.3654	**0.2560**	83.8	0.2836	0.2129	0.1076	**86.8**	-	40.4
	nMA	1.0460	0.7384	0.8057	11.3	1.3213	0.9404	1.0342	10.7	4.08	0.03
	SDE Matching	0.9018	0.6403	0.6926	14.1	0.5462	0.4124	0.5027	14.2	4.14	0.07
	NGDB	0.5186	0.3972	0.2884	73.4	0.2971	0.2215	0.1268	72.5	122.28	1.58
	PICSB	1.3783	0.9370	1.0210	25.0	0.2825	0.2097	**0.0863**	40.7	4.12	0.12
	PRBridge (Sparse)	**0.4764**	**0.3592**	0.2798	**85.2**	**0.2676**	**0.1760**	0.1031	79.8	6.76 / 1.44	11.62 / 0.02
Dense supervision	DiffusionPDE	1.1057	0.7538	0.6136	72.6	0.3822	0.3605	0.4829	75.0	3.89	20.8
	SDA	1.4425	1.3871	0.9406	71.6	0.2892	0.2012	0.1429	80.4	4.03	53.76
	PalSB	0.5781	0.4142	0.4006	22.4	0.2617	0.1891	0.1573	49.5	4.09	0.12
	PRBridge (Dense)	**0.4523**	**0.3332**	**0.2508**	**87.0**	**0.2229**	**0.1592**	**0.1198**	**82.0**	5.29	11.51

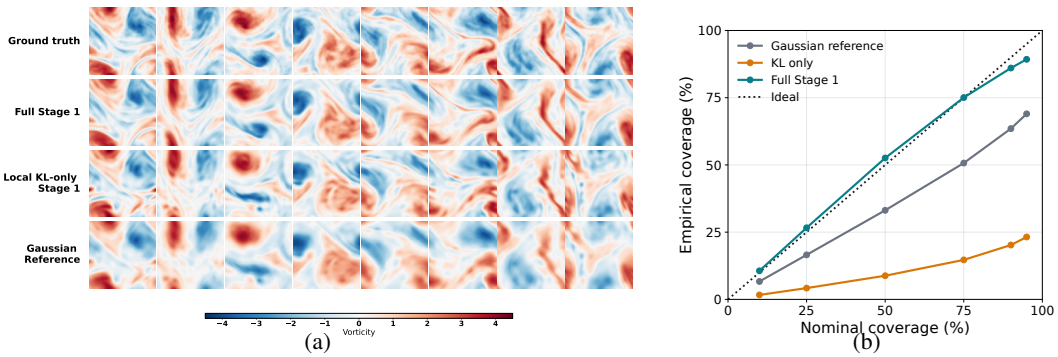
(c)		**Burgers 1D: 16 sensors 3.13%**				**Cylinder: 28 sensors 0.40%**					
	Method	RMSE	Rel. L_2	CRPS	Cov90	RMSE	Rel. L_2	CRPS	Cov90	Train (GPU·h)	Infer (s/inst, $M = 32$)
Sparse only	Gaussian smoother	0.0447	0.2187	0.0512	77.3	0.2524	0.3475	0.1061	69.6	-	4.28
	EnKF	0.0266	0.1316	**0.0151**	**88.1**	0.2369	0.3234	0.1216	72.0	-	17.36
	nMA	0.1542	0.7262	0.1760	13.5	0.2329	0.3201	0.1228	19.3	2.10	0.04
	SDE Matching	0.0669	0.3324	0.0756	31.9	0.2737	0.3768	0.1534	25.8	2.18	0.03
	NGDB	0.0278	0.1386	0.0174	78.3	0.2520	0.3470	**0.1060**	76.8	60.36	0.73
	PICSB	0.0374	0.1657	0.0234	52.7	0.2314	0.3185	0.1163	27.9	2.22	0.31
	PRBridge (Sparse)	**0.0263**	**0.1305**	0.0162	82.7	**0.2295**	**0.3076**	0.1063	**78.1**	5.44 / 1.62	4.33 / 0.02
Dense supervision	DiffusionPDE	0.0351	0.1815	0.0415	87.3	0.2073	0.2703	0.0847	**80.2**	3.28	37.5
	SDA	0.0295	0.1585	0.0863	68.8	0.2136	0.2844	0.0695	79.2	3.28	78.1
	PalSB	0.0510	0.2458	0.0339	36.2	0.2292	0.3162	0.1167	29.7	2.07	0.13
	PRBridge (Dense)	**0.0249**	**0.1302**	**0.0150**	**89.7**	**0.1973**	**0.2496**	**0.0632**	79.6	5.48	4.51

Table 2: Ablation study with fixed uniform sensors.

	FHN: 256 sensors 1.56%						**NS v2: 64 sensors 0.39%**					
Variant	RMSE ↓	Rel. L_2 ↓	CRPS ↓	Cov90	Train (GPU·h)	Infer (s/inst, $M = 32$)	RMSE ↓	Rel. L_2 ↓	CRPS ↓	Cov90	Train (GPU·h)	Infer (s/inst, $M = 32$)
Stage 1 reconstruction												
Gaussian smoother	0.0683	0.2156	0.0591	69.9	-	2.272	0.5745	0.4048	0.3204	67.3	-	11.396
KL only	0.0546	0.1720	0.0308	53.8	2.397	2.291	0.5751	0.4105	0.3868	25.6	4.594	11.451
Stage 1 ensemble only	0.0552	0.1799	0.0426	50.3	2.867	2.291	0.5155	0.3923	0.3086	61.7	6.440	11.714
KL + short ens.[†]	0.0537	0.1717	0.0304	66.2	2.451	2.293	0.5428	0.4004	0.3235	38.5	5.232	11.563
Full Stage 1 (long ens.)	0.0483	0.1703	0.0250	77.8	2.764	2.297	0.4764	0.3592	0.2798	85.2	6.764	11.623
Stage 2 distillation												
KL only → Stage 2	0.0541	0.1726	0.0331	59.5	1.661	0.027	0.5713	0.4104	0.3904	31.9	1.532	0.027
Full Stage 1 → Stage 2	0.0526	0.1711	0.0303	65.6	1.887	0.029	0.4896	0.3628	0.2867	62.8	1.440	0.028
Dense-supervision diagnostic												
Dense training	0.0414	0.1417	0.0252	83.6	2.593	2.284	0.4523	0.3332	0.2508	87.0	5.295	11.511

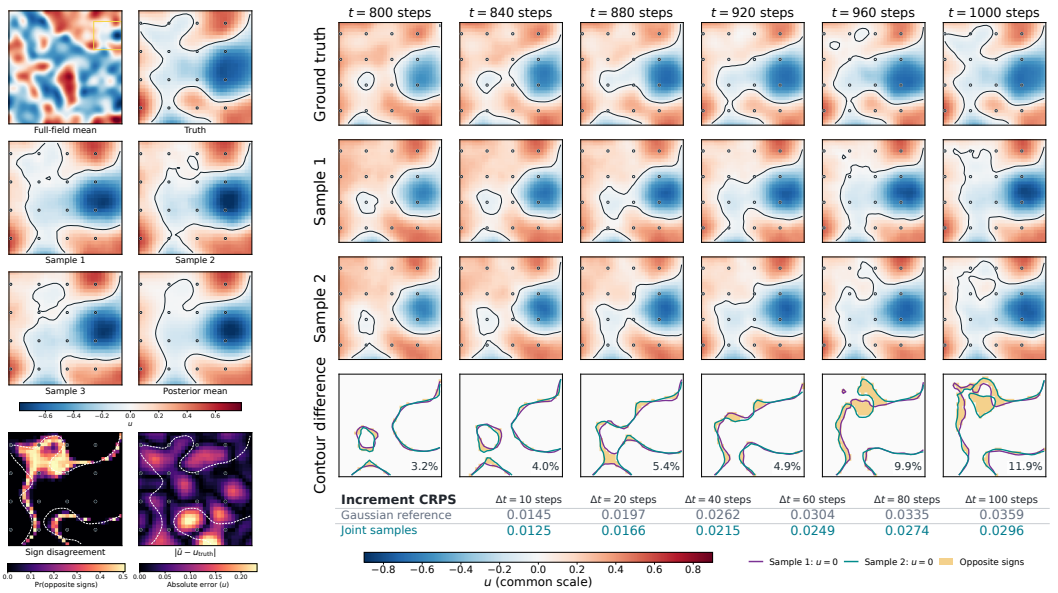
[†]Full PRBridge stage 1 combines the path-space KL objective with long-window ensemble supervision (e.g. 10/20/30); "short ens." uses a short window (e.g. 1/3/5).

sensors. Figure 3a provides a spatial interpretation of the sampled uncertainty. Under the same initial condition and observation history,nposterior samples exhibit localized changes in the positions and shapes of the pattern interfaces. Sign disagreement measures $\Pr[u^{(1)}(\mathbf{r})u^{(2)}(\mathbf{r}) < 0 \mid c]$ for two independent posterior draws at spatial location $\mathbf{r}$. The sign-disagreement probability, computed without using the ground truth, is elevated around portions of the ground-truth zero-level interfaces. This spatial correspondence is consistent with uncertainty in interface placement and illustrates how conditional variability is organized around the reconstructed patterns. Figure 3b shows 2 jointly



(a) (b)

Figure 2: Analysis of stochastic Navier–Stokes reconstruction.**(a) Qualitative effect of long-window ensemble supervision. (b) Coverage 90 of Navier-Stokes, 64 sensors.**



Increment CRPS	Δt = 10 steps	Δt = 20 steps	Δt = 40 steps	Δt = 60 steps	Δt = 80 steps	Δt = 100 steps
Gaussian reference	0.0145	0.0197	0.0262	0.0304	0.0335	0.0359
Joint samples	0.0125	0.0166	0.0215	0.0249	0.0274	0.0296

(a) (b)

Figure 3: **Coverage calibration and spatial structure of uncertainty. (a) Sampling variability from FHN, 256 sensors.** Full Stage 1 samples of the terminal u-field under a shared initial condition and observation history. Open circles indicate sensors and black contours mark the zero level of the fields. **(b) Sampling trajectory from FHN, 256 sensors.** For the same condition and sensor observations, the trained PRBridge is capable of sampling coherent trajectories across timesteps.

sampled trajectories conditioned on the same initial field and observation history. Each row follows one trajectory across physical time, allowing comparison of temporal evolution within a draw and variability across draws. Two different trajectories shows increased difference as steps increase. Compared with Gaussian reference sampling, the CRPS score increase slowly indicating our model could sample more physics meaningful trajectories after learning on the path KL divergence.

6 Conclusion

We presented Physics Reference Diffusion Bridge (PRBridge) for terminal-field reconstruction from fixed sparse sensors under stochastic dynamics. By calibrating an physics-informed Gaussian reference, PRBridge learns a physical-time path distribution through local optimization of a global variational objective. Long-window ensemble supervision exploits evolving state–sensor dependencies to improve the reconstruction of unobserved fine-scale structures. One-step distillation further reduces inference cost, with a trade-off in uncertainty calibration. These findings highlight the potential of solving the problem on lacking dense trajectory labels for scientific field reconstruction.

AI USE STATEMENT

We used AI in the following steps: 1) We used AI to aid or polish writing; 2) We used AI for research ideation or execution; 3) We used AI to draft sections of the paper; 4) We used AI for proving mathematical claims.

ETHICS STATEMENT

This work studies reconstruction of physical fields from sparse observations. Our experiments use simulated SPDE trajectories and experimental fluid-flow measurements and do not involve human participants or personal data. We do not identify specific ethical concerns arising from the present study.

REPRODUCIBILITY STATEMENT

The code of this paper would be updated on this anonymous repo: https://anonymous.4open.science/r/PRBridge-0113.

AUTHOR CONTRIBUTIONS

ACKNOWLEDGMENTS

REFERENCES

Feng Bao, Hristo G Chipilski, Siming Liang, Guannan Zhang, and Jeffrey S Whitaker. Nonlinear ensemble filtering with diffusion models: Application to the surface quasigeostrophic dynamics. *Monthly Weather Review*, 153(7):1155–1169, 2025.

Grigory Bartosh, Dmitry P Vetrov, and Christian Andersson Naesseth. Sde matching: Scalable and simulation-free training of latent stochastic differential equations. *The 43rd International Conference on Machine Learning (ICML)*, 2025.

Jan-Hendrik Bastek, WaiChing Sun, and Dennis Kochmann. Physics-informed diffusion models. In *International Conference on Learning Representations*, volume 2025, pp. 3360–3385, 2025.

Dohyun Bu, Chanho Kim, Seokun Choi, and Jong-Seok Lee. Sparsely-supervised data assimilation via physics-informed schr\" odinger bridge. *arXiv preprint arXiv:2603.22319*, 2026.

Colin Cotter, Dan Crisan, Darryl Holm, Wei Pan, and Igor Shevchenko. Data assimilation for a quasi-geostrophic model with circulation-preserving stochastic transport noise: C. cotter et al. *Journal of Statistical Physics*, 179(5):1186–1221, 2020.

Joseph L Doob. Conditional brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société mathématique de France*, 85:431–458, 1957.

Yifan Du, Mengze Wang, and Tamer A Zaki. State estimation in minimal turbulent channel flow: A comparative study of 4dvar and pinn. *International Journal of Heat and Fluid Flow*, 99:109073, 2023.

Geir Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367, 2003.

Geir Evensen. *Data assimilation: the ensemble Kalman filter*. Springer, 2009.

Christopher AT Ferro. Fair scores for ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society*, 140(683):1917–1923, 2014.

Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

QiZhi He, David Barajas-Solano, Guzel Tartakovsky, and Alexandre M Tartakovsky. Physics-informed neural networks for multiphysics data assimilation with application to subsurface transport. *Advances in Water Resources*, 141:103610, 2020.

Jeremy Heng, Adrian N Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential monte carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Peiyan Hu, Haodong Feng, Hongyuan Liu, Tongtong Yan, Wenhao Deng, Tianrun Gao, Rong Zheng, Haoren Zheng, Chenglei Yu, Chuanrui Wang, et al. Realpdebench: A benchmark for complex physical systems with real-world data. *arXiv preprint arXiv:2601.01829*, 2026.

Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. *Advances in Neural Information Processing Systems*, 37: 130291–130323, 2024.

Zeyu Li, Hongkun Dou, Shen Fang, Wang Han, Yue Deng, and Lijun Yang. Physics-aligned field reconstruction with diffusion bridge. In *International Conference on Learning Representations*, volume 2025, pp. 58864–58900, 2025.

Cristina Martinez-Torres, Wilhelm Stannat, and Jan Szalankiewicz. Nonlinear filtering and spatial asymptotic consistency for spdes observed via spatio-temporal point processes: C. martinez-torres et al. *Stochastics and Partial Differential Equations: Analysis and Computations*, pp. 1–40, 2026.

May Anne Mata, Meghan Dutot, Leah Edelstein-Keshet, and William R Holmes. A model for intracellular actin waves explored by nonlinear local perturbation analysis. *Journal of theoretical biology*, 334:149–161, 2013.

Marcin Mider, Moritz Schauer, and Frank Van der Meulen. Continuous-discrete smoothing of diffusions. *Electronic Journal of Statistics*, 15(2):4295–4342, 2021.

Maziar Raissi, Paris Perdikaris, and George E 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.

Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.

François Rozet and Gilles Louppe. Score-based data assimilation. *Advances in Neural Information Processing Systems*, 36:40521–40541, 2023.

Thomas Ryder, Dennis Prangle, Andrew Golightly, and Isaac Matthews. The neural moving average model for scalable variational inference of state space models. In *Uncertainty in Artificial Intelligence*, pp. 12–22. PMLR, 2021.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

Tim van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014. doi: 10.1109/TIT.2014.2320500.

Pengpeng Xiao, Phillip Si, and Peng Chen. Ld-ensf: synergizing latent dynamics with ensemble score filters for fast data assimilation with sparse observations. In *International Conference on Learning Representations*, volume 2026, pp. 152461–152491, 2026.

Gefan Yang, Frank Van Der Meulen, and Stefan Sommer. Neural guided diffusion bridges. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 71210–71230. PMLR, 13–19 Jul 2025. URL `https://proceedings.mlr.press/v267/yang25af.html`.

Tamer A Zaki and Mengze Wang. From limited observations to the state of turbulence: Fundamental difficulties of flow reconstruction. *Physical Review Fluids*, 6(10):100501, 2021.

Huijun Zhang, Yaxin Liu, Chongyu Zhang, and Ningyun Li. Machine learning methods for weather forecasting: A survey. *Atmosphere*, 16(1):82, 2025.

Kaiwen Zheng, Guande He, Jianfei Chen, Fan Bao, and Jun Zhu. Diffusion bridge implicit models. In *International Conference on Learning Representations*, volume 2025, pp. 81857–81884, 2025.

Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. Denoising diffusion bridge models. In *International Conference on Learning Representations*, volume 2024, pp. 8160–8171, 2024.

Zihan Zhou, Chenguang Wang, Hongyi Ye, Yongtao Guan, and Tianshu Yu. Incomplete data, complete dynamics: A diffusion approach. In *International Conference on Learning Representations*, volume 2026, pp. 10918–10952, 2026.

Yuantu Zhu, Zheyan Li, Dai Shi, Luke Thompson, Oliver Nash, Jose Miguel Lara Rangel, Siran Li, Bingguang Chen, Rongchan Zhu, Qi Meng, et al. Spdebench: An extensive benchmark for learning stochastic pdes. *arXiv preprint arXiv:2505.18511*, 2025.

A Related Works

Physics-informed learning incorporates governing equations into reconstruction objectives (Raissi et al., 2019; He et al., 2020). PalSB and DiffusionPDE combine high-resolution training data with physical constraints for generative reconstruction (Li et al., 2025; Huang et al., 2024). PICSB reduces dense supervision through conditional low-to-high-fidelity transport with physics residuals and hard observation constraints, including globally fixed sensor settings (Bu et al., 2026). PRBridge instead defines its target through known stochastic transition laws and noisy observation likelihoods, learning the physical path posterior without post-initial full-field training targets.

For datat assimilation, Ensemble Kalman methods propagate information through state–observation cross-covariances (Evensen, 2003), while flow-reconstruction studies relate recoverability to dynamical sensitivity and observation coverage (Zaki & Wang, 2021; Du et al., 2023). Recent approaches generate analysis ensembles through training-free diffusion-based filtering (Bao et al., 2025), or involving dynamivs into latent space (Xiao et al., 2026). PRBridge uses long-window forecast–analysis ensembles as detached distributional training targets, supplementing local KL optimization with explicit dynamical information.

Diffusion bridge create controlled between two endpoints (Zhou et al., 2024; Zheng et al., 2025). Guided diffusion bridges approximate stochastic dynamics conditioned on observations (Mider et al., 2021). NGDB learns neural drift corrections through path-space variational optimization, differentiating through numerical trajectory simulation (Yang et al., 2025). PRBridge combines a temporally correlated Gaussian reference with a local invertible transport. After reference construction, joint reference windows permit unbiased local path-KL gradient estimation, while retaining temporal dependence through the reference.

B From physical bridges to local variational training

Throughout this section, fix the conditioning context c. All densities are defined on the same free-state coordinates, with the deterministic initial state x_0 excluded from the density variables. Assume positive physical transition densities, finite positive evidence, a nonsingular Gaussian reference, and finite expectations in the expressions below. The reference statistics are independent of ϕ. For gradient identities, assume that differentiation can be interchanged with expectation.

B.1 Derivation of the path-space variational objective

Fix a context c and the known initial state x_0. The normalized form of Equation equation 5 is

$$\pi(x_{1:N} \mid c) = \frac{1}{\mathcal{Z}(c)} \prod_{n=1}^{N} p_n(x_n \mid x_{n-1})\, g_n(x_n; c), \tag{25}$$

where

$$\mathcal{Z}(c) = \int \prod_{n=1}^{N} p_n(x_n \mid x_{n-1})\, g_n(x_n; c)\, \mathrm{d}x_{1:N}. \tag{26}$$

We assume $0 < \mathcal{Z}(c) < \infty$ and that the expectations below are finite.

Taking logarithms gives

$$\log \pi(x_{1:N} \mid c) = \sum_{n=1}^{N} \log(p_n(x_n \mid x_{n-1}) g_n(x_n; c)) - \log \mathcal{Z}(c). \tag{27}$$

By definition, the reverse KL is

$$\mathrm{KL}(q_\phi(x_{1:N} \mid c) \,\|\, \pi(x_{1:N} \mid c)) = \mathbb{E}_{x_{1:N} \sim q_\phi(\cdot|c)} \left[\log \frac{q_\phi(x_{1:N} \mid c)}{\pi(x_{1:N} \mid c)} \right]. \tag{28}$$

Substituting Equation equation 27 into this expression yields

$$\mathrm{KL}(q_\phi(x_{1:N} \mid c) \,\|\, \pi(x_{1:N} \mid c)) = \mathbb{E}_{q_\phi} \left[\log q_\phi(x_{1:N} \mid c) - \sum_{n=1}^{N} \log(p_n(x_n \mid x_{n-1}) g_n(x_n; c)) \right] + \log \mathcal{Z}(c), \tag{29}$$

where $\mathbb{E}_{q_\phi}$ denotes expectation over $x_{1:N} \sim q_\phi(\cdot \mid c)$.

Averaging over the fixed distribution of training contexts gives Equation equation 7, with

$$C_{\mathrm{phys}} = \mathbb{E}_c[\log \mathcal{Z}(c)]. \tag{30}$$

The physical model and observation likelihoods are independent of ϕ, so this constant has zero ϕ-gradient.

Equivalently, the context-specific ELBO is

$$\mathrm{ELBO}_\phi(c) = \mathbb{E}_{q_\phi} \left[\sum_{n=1}^{N} \log(p_n(x_n \mid x_{n-1}) g_n(x_n; c)) - \log q_\phi(x_{1:N} \mid c) \right], \tag{31}$$

and therefore

$$\mathcal{L}_{\mathrm{KL}}(\phi) = -\mathbb{E}_c[\mathrm{ELBO}_\phi(c)] + C_{\mathrm{phys}}. \tag{32}$$

After introducing the reference-based sampler $z_{1:N} \sim q_g(\cdot \mid c)$ and $x_{1:N} = T_\phi(z_{1:N}; c)$, the same objective can be evaluated as

$$\mathcal{L}_{\mathrm{KL}}(\phi) = \mathbb{E}_{c,\, z_{1:N} \sim q_g(\cdot|c)} \left[\log q_\phi(x_{1:N} \mid c) - \sum_{n=1}^{N} \log(p_n(x_n \mid x_{n-1}) g_n(x_n; c)) \right] + C_{\mathrm{phys}}, \tag{33}$$

where $x_{1:N} = T_\phi(z_{1:N}; c)$ inside the expectation. This change of sampling variables follows from the definition of q_ϕ as the pushforward of q_g.

B.2 Proof of Proposition 3.3

Proof. Fix a context c and define $\mathcal{I}_n = \{\max(1, n-2), \ldots, n\}$. By the local transport construction, $F_{\phi,n}$ depends only on $z_{\mathcal{I}_n}$ and c, with $z_0 = x_0$ fixed. Its expectation under the full reference therefore equals its expectation under the corresponding joint window marginal.

Sample K uniformly from $\{1, \ldots, N\}$ and, conditional on $K = n$, draw $z_{\mathcal{I}_n} \sim q_g(z_{\mathcal{I}_n} \mid c)$. For $\widehat{\mathcal{L}}_c = N F_{\phi,K}$,

$$\mathbb{E}[\widehat{\mathcal{L}}_c \mid c] = \sum_{n=1}^{N} \frac{1}{N} \mathbb{E}_{q_g(z_{\mathcal{I}_n}|c)}[N F_{\phi,n}] = \mathbb{E}_{z_{1:N} \sim q_g(\cdot|c)} \left[\sum_{n=1}^{N} F_{\phi,n} \right]. \tag{34}$$

By the transport density formula and the posterior factorization, the final expression differs from $\mathrm{KL}(q_\phi(x_{1:N} \mid c) \| \pi(x_{1:N} \mid c))$ only by a constant independent of ϕ.

The reference distribution and index probabilities are independent of ϕ. Interchanging differentiation and expectation under the assumed regularity conditions, then averaging over the training-context distribution, gives

$$\mathbb{E}_c \mathbb{E}[\nabla_\phi \widehat{\mathcal{L}}_c \mid c] = \mathbb{E}_c \left[\nabla_\phi \mathrm{KL}(q_\phi(x_{1:N} \mid c) \,\|\, \pi(x_{1:N} \mid c)) \right] = \nabla_\phi \mathcal{L}_{\mathrm{KL}}(\phi). \tag{35}$$

This proves the result. □

The gradient of $F_{\phi,n}$ includes differentiation through both transported states x_n and x_{n-1} when $n \geq 2$. Holding the reference statistics fixed does not justify detaching the transported predecessor x_{n-1}.

B.3 Stage 1 Settings

Terminal KL is bounded by path KL. Factor both path laws into a terminal marginal and a conditional distribution over the preceding states. The KL chain rule gives

$$\begin{aligned}&\mathrm{KL}(q_\phi(\cdot \mid c) \,\|\, \pi(\cdot \mid c))\\ &= \mathrm{KL}(q_{\phi,N}(\cdot \mid c) \,\|\, \pi_N(\cdot \mid c))\\ &\quad + \mathbb{E}_{x_N \sim q_{\phi,N}(\cdot|c)} \left[\mathrm{KL}(q_\phi(x_{1:N-1} \mid x_N, c) \,\|\, \pi(x_{1:N-1} \mid x_N, c))\right].\end{aligned} \tag{36}$$

Nonnegativity of the second term proves KL divergence after averaging over c. Equality holds when the two conditional history laws agree for $q_{\phi,N}$-almost every terminal state. The bound does not imply that every optimization step reducing path KL must also reduce terminal KL.

The numerical Doob transform. Define the discrete backward information function by

$$h_n(x; c) = \mathbb{E}_{\mathrm{phys}}\left[\prod_{k=n+1}^{N} g_k(x_k; c) \,\middle|\, x_n = x\right], \qquad h_N(x; c) = 1. \tag{37}$$

The Markov property of the numerical physical model yields

$$h_{n-1}(x; c) = \int p_n(x' \mid x)\, g_n(x'; c)\, h_n(x'; c)\, \mathrm{d}x'. \tag{38}$$

Consequently, Equation posteroir transition defines a normalized transition density. Multiplication along a path gives

$$\begin{aligned}\prod_{n=1}^{N} k_n^\star(x_n \mid x_{n-1}, c) &= \left[\prod_{n=1}^{N} p_n(x_n \mid x_{n-1})\, g_n(x_n; c)\right] \frac{h_N(x_N; c)}{h_0(x_0; c)}\\ &= \pi(x_{1:N} \mid c),\end{aligned} \tag{39}$$

because $h_N = 1$ and $h_0(x_0; c) = \mathcal{Z}(c)$. This is an exact identity for the declared numerical model. It does not assume that conditioning and numerical discretization commute at a finite step size.

Conditional transition KL decomposition. An arbitrary variational path law factorizes as

$$q_\phi(x_{1:N} \mid c) = \prod_{n=1}^{N} q_\phi(x_n \mid x_{<n}, c). \tag{40}$$

Combining this identity with Equation equation 39, and taking conditional expectations at each time, proves KL desomposition.

Furthermore, the backward-information terms cancel pathwise:

$$\sum_{n=1}^{N} [\log h_{n-1}(x_{n-1}; c) - \log h_n(x_n; c)] = \log \mathcal{Z}(c). \tag{41}$$

Hence

$$\begin{aligned}&\mathrm{KL}(q_\phi(\cdot \mid c) \,\|\, \pi(\cdot \mid c))\\ &= \mathbb{E}_{q_\phi}\left[\log q_\phi(x_{1:N} \mid c) - \sum_{n=1}^{N} \{\log p_n(x_n \mid x_{n-1}) + \log g_n(x_n; c)\}\right] + \log \mathcal{Z}(c).\end{aligned} \tag{42}$$

This removes the need to evaluate the exact bridge transitions during training.

Gaussian reference and forward sampling. The exact-smoother and Markov representations below describe a special case; in our implementation, as in the main text, we retain the surrogate smoothing mean and prescribe a temporally correlated covariance for q_g, whose Markov factorization is not required for the local-gradient result. Consider a fixed linear-Gaussian surrogate

$$\widetilde{p}_n(z' \mid z) = \mathcal{N}\left(z'; A_n^g z + \beta_n^g, Q_n^g\right), \qquad Z_0 = x_0. \tag{43}$$

Conditioning this surrogate on the linear Gaussian measurements gives

$$\begin{aligned}q_g(Z_{1:N} \mid c) &= \frac{1}{\mathcal{Z}_g(c)} \prod_{n=1}^{N} \widetilde{p}_n(Z_n \mid Z_{n-1})\, g_n(Z_n; c)\\ &= \mathcal{N}(m_g, \Sigma_g).\end{aligned} \tag{44}$$

Let h_n^g denote the surrogate's backward information function. Its conditioned transition is

$$r_n^g(z' \mid z, c) = \frac{\widetilde{p}_n(z' \mid z)\, g_n(z'; c)\, h_n^g(z'; c)}{h_{n-1}^g(z; c)}. \tag{45}$$

The linear Gaussian construction makes this transition Gaussian, so it admits a forward sampling representation

$$Z_n = \overline{A}_n^g Z_{n-1} + \overline{\beta}_n^g + \overline{L}_n^g \xi_n, \qquad \xi_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I), \tag{46}$$

where the coefficients depend on the full context c. Thus, forward sampling proceeds in physical time while retaining information from future observations.

Once the reference statistics have been prepared, any joint window $\mathcal{W}$ can also be sampled directly:

$$\begin{aligned} Z_{\mathcal{W}} &= m_{g,\mathcal{W}} + L_{g,\mathcal{W}} \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I), \\ L_{g,\mathcal{W}} L_{g,\mathcal{W}}^\top &= \Sigma_{g,\mathcal{W},\mathcal{W}}. \end{aligned} \tag{47}$$

The covariance blocks must come from the same joint reference law. Sampling each time marginal independently would generally produce a different training objective. Reference preparation remains a separate computational cost.

Triangular transport and generated conditionals. Write equation as

$$\begin{aligned} x_n &= D_{\phi,n}(Z_{n-1}, c) Z_n + u_{\phi,n}(Z_{n-1}, c), \\ D_{\phi,n}(z, c) &:= \operatorname{diag}\left(\exp s_{\phi,n}(z, c)\right). \end{aligned} \tag{48}$$

Starting from $Z_0 = x_0$, the inverse is obtained recursively:

$$Z_n = D_{\phi,n}(Z_{n-1}, c)^{-1} \left[x_n - u_{\phi,n}(Z_{n-1}, c)\right]. \tag{49}$$

The Jacobian of the full map is block lower triangular, with diagonal blocks $D_{\phi,n}$. Therefore,

$$\begin{aligned} \log q_\phi(x_{1:N} \mid c) = \log q_g(Z_{1:N} \mid c) \\ - \sum_{n=1}^{N} \sum_{i=1}^{d} s_{\phi,n,i}(Z_{n-1}, c). \end{aligned} \tag{50}$$

For the conditioned Markov Gaussian reference, the same argument gives the history-conditioned density

$$\begin{aligned} q_\phi(x_n \mid x_{<n}, c) = r_n^g(Z_n \mid Z_{n-1}, c) \\ \times \exp\left[-\sum_{i=1}^{d} s_{\phi,n,i}(Z_{n-1}, c)\right], \end{aligned} \tag{51}$$

evaluated at the recursively recovered latent states. Although the reference transition depends only on Z_{n-1}, recovering that latent state generally requires the generated history. The corrected process therefore need not be Markov in x_n alone.

Local objective and unbiased gradients. Substituting Equation equation 50 into Equation equation 42 yields

$$\begin{aligned} &\mathrm{KL}(q_\phi(\cdot \mid c) \,\|\, \pi(\cdot \mid c)) \\ &= \mathbb{E}_{Z \sim q_g(\cdot \mid c)} \left[\sum_{n=1}^{N} F_{\phi,n}\right] + C(c), \\ C(c) &:= \mathbb{E}_{q_g} \log q_g(Z_{1:N} \mid c) + \log \mathcal{Z}(c). \end{aligned} \tag{52}$$

Because q_g is fixed with respect to ϕ, $C(c)$ has zero ϕ-gradient.

For $n \geq 2$, x_n depends on (Z_{n-1}, Z_n) and x_{n-1} depends on (Z_{n-2}, Z_{n-1}). Thus $F_{\phi,n}$ is measurable with respect to a three-state reference window. The boundary cases use the deterministic $Z_0 = x_0$.

For $\Pr(K = n) = \rho_n > 0$, sampling a valid joint window gives

$$\begin{aligned}\mathbb{E}_{K,Z}\left[\frac{F_{\phi,K}}{\rho_K}\right] &= \sum_{n=1}^{N} \rho_n \, \mathbb{E}_{q_g}\left[\frac{F_{\phi,n}}{\rho_n}\right] \\ &= \mathbb{E}_{q_g}\left[\sum_{n=1}^{N} F_{\phi,n}\right].\end{aligned} \tag{53}$$

Consequently,

$$\mathbb{E}_{c,K,Z}\left[\nabla_\phi \frac{F_{\phi,K}}{\rho_K}\right] = \nabla_\phi \mathcal{L}_{\mathrm{KL}}(\phi). \tag{54}$$

For uniform time sampling, the estimator is simply $NF_{\phi,K}$. The local contributions need not be nonnegative; their expected sum equals the path KL up to a parameter-independent constant.

Terminal sampling. The terminal coordinate depends only on the final reference pair. Therefore, for any measurable terminal-state set A,

$$q_{\phi,N}(A \mid c) = \mathbb{P}_{(z_{N-1},z_N)\sim q_g(\cdot|c)}\left(T_{\phi,N}(z_{N-1}, z_N; c) \in A\right). \tag{55}$$

Thus the final joint reference pair suffices to sample the terminal marginal of the learned path law after reference preparation. This is exact sampling from $q_{\phi,N}$; approximation to the target π_N depends on the learned transport.

Why pairwise transition KLs are insufficient in general. If the generated process is not Markov, replacing its history-conditioned transitions by $q_\phi(x_n \mid x_{n-1}, c)$ omits a nonnegative term:

$$\begin{aligned}&\mathrm{KL}(q_\phi(\cdot \mid c) \,\|\, \pi(\cdot \mid c)) \\ &= \sum_{n=1}^{N} \mathbb{E}_{q_\phi(x_{n-1}|c)}\left[\mathrm{KL}(q_\phi(\cdot \mid x_{n-1}, c) \,\|\, k_n^\star(\cdot \mid x_{n-1}, c))\right] \\ &\quad + \sum_{n=3}^{N} I_{q_\phi}\left(x_n; x_{1:n-2} \mid x_{n-1}, c\right).\end{aligned} \tag{56}$$

The additional terms are conditional mutual informations. This explains why the exact conditional decomposition retains the generated history, whereas computational locality is obtained through the reference transport and the additive objective.

C Proof of the sensor-dependent reconstruction bound

Physical-noise coordinates. Consider the additive numerical model on its free state coordinates,

$$x_n = f_n(x_{n-1}) + \eta_n, \qquad n = 1, \ldots, N, \tag{57}$$

where x_0 and the physical parameters are fixed. The noise history $\eta = (\eta_1, \ldots, \eta_N)$ determines the numerical trajectory, and a compatible path determines its increments through $\eta_n(x) = x_n - f_n(x_{n-1})$.

For this analysis, define

$$G(\eta) = \mathrm{col}_{j=1}^{J}\left(Hx_{n_j}(\eta)\right), \qquad R(\eta) = P_B x_N(\eta). \tag{58}$$

Unlike the state-perturbation maps G_a, R_a in Definition 3.4, these maps vary the complete physical-noise history. Accordingly, throughout this proof, $J_G = DG(\eta^\star)$ and $J_R = DR(\eta^\star)$ denote derivatives with respect to the admissible noise coordinates.

Second-order regularity. The constants in the theorem satisfy

$$\sup_{\eta\in\mathcal{B}_\rho(\eta^\star)} \|D^2 G(\eta)\|_{\mathrm{op}} \le M_G, \qquad \sup_{\eta\in\mathcal{B}_\rho(\eta^\star)} \|D^2 R(\eta)\|_{\mathrm{op}} \le M_R. \tag{59}$$

For a vector-valued map F, the second-derivative operator norm is the bilinear norm

$$\|D^2 F(\eta)\|_{\mathrm{op}} := \sup_{\|u\|_2=\|v\|_2=1} \|D^2 F(\eta)[u, v]\|_2, \tag{60}$$

with directions restricted to the admissible noise coordinates.

Proof of Theorem 3.5.

Proof. Fix the condition c and generating noise history $\eta^\star$. Write the stacked observations as $y = G(\eta^\star) + \varepsilon$, with $\|\varepsilon\|_2 \le \delta$.

First consider any path $x_{1:N} \in \mathcal{A}_{\rho,\tau}$, and set $v = \eta(x) - \eta^\star$. By the definition of $\mathcal{A}_{\rho,\tau}$,

$$\|v\|_2 \le \rho, \qquad \|G(\eta^\star + v) - y\|_2 \le \tau. \tag{61}$$

For either $F = G$ or $F = R$, Taylor's theorem gives

$$F(\eta^\star + v) - F(\eta^\star) = DF(\eta^\star)v + r_F(v), \tag{62}$$

where

$$r_F(v) = \int_0^1 (1-s)D^2F(\eta^\star + sv)[v,v]\,\mathrm{d}s, \qquad \|r_F(v)\|_2 \le \frac{1}{2}M_F\|v\|_2^2. \tag{63}$$

Consequently, the linearized observation response satisfies

$$\|J_G v\|_2 \le \|G(\eta^\star + v) - y\|_2 + \|y - G(\eta^\star)\|_2 + \|r_G(v)\|_2 \le \tau + \delta + \frac{1}{2}M_G\rho^2. \tag{64}$$

Next, decompose the terminal response as

$$J_R v = J_R J_G^\dagger J_G v + J_R\left(I - J_G^\dagger J_G\right)v. \tag{65}$$

By the definitions of κ and β,

$$\|J_R v\|_2 \le \kappa\|J_G v\|_2 + \beta\|v\|_2. \tag{66}$$

Combining this inequality with the target Taylor remainder yields

$$\|R(\eta^\star + v) - R(\eta^\star)\|_2 \le \kappa(\tau + \delta) + \beta\rho + \frac{1}{2}(M_R + \kappa M_G)\rho^2 = r_{\rho,\tau}. \tag{67}$$

Since $R(\eta(x)) = P_B x_N$, the preceding inequality holds for every path in $\mathcal{A}_{\rho,\tau}$. Therefore, the event

$$\mathcal{E} := \{x_{1:N} : \|P_B x_N - R(\eta^\star)\|_2 > r_{\rho,\tau}\} \tag{68}$$

is contained in $\mathcal{A}_{\rho,\tau}^c$. The assumed posterior concentration thus implies $\pi(\mathcal{E} \mid c) \le \gamma$.

Finally, using the convention $\mathrm{TV}(P,Q) = \sup_A |P(A) - Q(A)|$, Pinsker's inequality van Erven & Harremos (2014) gives

$$q_\phi(\mathcal{E} \mid c) \le \pi(\mathcal{E} \mid c) + \mathrm{TV}\left(q_\phi(\cdot \mid c), \pi(\cdot \mid c)\right) \le \gamma + \sqrt{\frac{1}{2}\mathrm{KL}(q_\phi(\cdot \mid c)\|\pi(\cdot \mid c))}. \tag{69}$$

This proves the claimed reconstruction bound. □

D Detailed experimental setup

D.1 Synthetic data under benckmark protocals

Synthetic stochastic dynamics. After spatial discretization, the synthetic systems take the form

$$\mathrm{d}x_t = b_\theta(x_t, t)\,\mathrm{d}t + \sigma Q^{1/2}\,\mathrm{d}W_t, \tag{70}$$

where W_t is a standard Wiener process and Q specifies the spatial covariance of the forcing. For FHN, the noise term is applied separately to the two state channels with independent Wiener drivers. All three systems use periodic boundary conditions and an implicit–explicit (IMEX) Euler integrator. Our datasets use independently specified equation parameters, initial conditions, and noise configurations, with simulation implementations informed by SPDEBench Zhu et al. (2025) and its `SPDE_hackathon` repository.

Stochastic Burgers. We simulate the one-dimensional Burgers equation on $[0, 1)$ using 512 spatial grid points. The viscosity is $\nu = 0.01$, the noise amplitude is $\sigma = 0.2$, and the terminal time is $T = 1$. Each trajectory uses 2,000 integration steps, corresponding to a time step of 5×10^{-4}.

Stochastic FitzHugh–Nagumo. The FHN system comprises two reaction–diffusion fields (u, v) on $[-1, 1)^2$, discretized on a 128×128 grid. We use diffusion coefficients $D_u = 10^{-3}$ and $D_v = 5 \times 10^{-3}$, reaction parameter $k = 5 \times 10^{-3}$, and terminal time $T = 5$. Each trajectory uses 1,000

integration steps. Both channels have noise amplitude 0.1, with independent process noise across channels.

Stochastic Navier–Stokes. We simulate the two-dimensional incompressible Navier–Stokes equation in vorticity form on $[0,1)^2$ using a 128×128 grid. The viscosity is $\nu = 10^{-4}$, the noise amplitude is $\sigma = 0.3$, and the deterministic forcing is

$$f(\boldsymbol{r}) = 0.1\left[\sin\big(2\pi(r_1+r_2)\big) + \cos\big(2\pi(r_1+r_2)\big)\right], \tag{71}$$

where $\boldsymbol{r} = (r_1, r_2)$ denotes spatial position. We integrate to $T = 5$ with 10,000 steps of size 5×10^{-4}. Before each reconstruction trajectory, the initial vorticity is obtained through a 10-time-unit stochastic spin-up under the same dynamics. All reported synthetic NS results use this configuration, referred to as NS v2 in our implementation.

Spatial covariance of the process noise. The covariance is diagonal in the Fourier basis, with a Whittle–Matérn-type spectrum

$$q_{\boldsymbol{k}} \propto \left(1 + \ell^2 \left\| \frac{2\pi \boldsymbol{k}}{L} \right\|_2^2 \right)^{-2}, \tag{72}$$

where L is the domain side length. The discrete covariance is normalized so that $\mathrm{diag}(Q) = \mathbf{1}$. Consequently, σ controls the pointwise standard deviation of the noise increment per square root of time. The correlation lengths are $\ell = 0.02$ for Burgers, 0.05 for FHN, and 0.04 for NS. For NS, the zero Fourier mode is removed before normalization to preserve zero-mean vorticity. The process-noise covariance Q and the measurement-noise covariance Γ_j are treated as distinct quantities.

Initial conditions and dataset construction. For each equation, we generate 1,000 initial conditions and 10 independent process-noise realizations per initial condition, yielding 10,000 trajectories. Each trajectory contains 101 saved frames, including the initial and terminal fields. The initial fields are sampled from band-limited Gaussian random fields with Fourier modes $0 < \|\boldsymbol{k}\|_2 \le 4$ and dataset-specific RMS normalization. For NS, these fields are subsequently evolved through the spin-up procedure described above. We assign 800, 100, and 100 initial conditions to training, validation, and testing, respectively, yielding 8,000/1,000/1,000 trajectories. All process-noise realizations associated with an initial condition remain in the same split.

Synthetic sensor layouts and measurements. For each experimental setting, the sensor operator H is generated once and remains fixed across all trajectories, splits, and observation times. Uniform layouts are origin-aligned lattices with $n_s \in \{8, 16, 32\}$ sensors per spatial dimension. This gives n_s point sensors for Burgers and n_s^2 sensor locations for FHN and NS. The two-dimensional budgets are therefore 64, 256, and 1,024 locations, corresponding to approximately 0.39%, 1.56%, and 6.25% of the spatial grid. The three uniform layouts are nested.

Observations are collected at all 100 non-initial saved frames:

$$y_j = Hx_{s_j} + \varepsilon_j, \qquad \varepsilon_j \sim \mathcal{N}(0, 0.05^2 I), \qquad j = 1, \ldots, 100. \tag{73}$$

Measurement noise is generated once using fixed random seeds and reused across methods.

Permanently unobserved FHN regions. The nonuniform FHN settings exclude all sensors from a central 64×64 box, which occupies one quarter of the spatial domain. Either 64 or 256 sensor locations are sampled outside this box using a fixed random seed and then held fixed. The reported non-sensor metrics are evaluated over all non-sensor locations, including the central box. They are not restricted to the box itself.

D.2 Real-world Cylinder benchmark

Recordings and nondimensionalization. We use the real PIV velocity measurements from the Cylinder scenario of RealPDEBench Hu et al. (2026). Our reconstruction protocol retains the 53 recordings with $Re \ge 6000$ from the 92 available recordings. The lower-Re recordings are excluded because our kinematic consistency checks identify discrepancies between the recorded velocities and the associated spatial and temporal scales.

Spatial coordinates are normalized by the cylinder diameter $D = 30\,\mathrm{mm}$ and velocities by the recording-specific free-stream velocity U. The resulting velocity fields (u, v) are represented on a common 64×112 grid, with approximately 11.8 grid cells per cylinder diameter. The parameters $\theta = (Re, \Delta t^*)$ remain constant within each recording.

Temporal windows and data splits. Each window targets a duration of 10 convective time units and contains 101 frames selected using an integer temporal stride. Only recorded frames are used; no temporal interpolation is applied. Training windows may overlap in time, but recordings are assigned to disjoint data splits.

The training split contains 31 recordings and 2,145 windows, and the validation split contains 8 recordings and 140 windows. The in-distribution test split contains 7 recordings and 121 windows. The out-of-distribution test split contains 7 recordings and 133 windows, comprising the highest-Reynolds-number block with $Re \in [11250, 11812]$. This split evaluates extrapolation to Reynolds numbers beyond the training range.

Masks, sensors, and observation covariance. The cylinder body and an identified PIV-artifact region are excluded from both sensing and evaluation. Uniform lattices with spacings of 16, 8, or 4 grid cells yield 28, 107, or 432 sensor locations after masking, approximately 0.4%, 1.5%, or 6.2% of the valid evaluation cells. Each sensor measures both velocity components at every selected observation frame. No synthetic measurement noise is added. The observation covariance is specified using the recording-specific temporal variability of the free stream.

Reference selection. For Cylinder, the Gaussian-reference hyperparameters are selected using dense validation fields. This introduces full-field information during reference selection, even for variants whose neural training objectives use only sparse post-initial observations. We therefore disclose reference selection separately from the supervision used to train the learned sampler.

E EVALUATION PROTOCOL

Evaluation coordinates and ensemble predictions. Let $\mathcal{E}$ denote the set of valid field coordinates and $\mathcal{U} \subseteq \mathcal{E}$ the coordinates not directly measured by sensors. For Cylinder, $\mathcal{E}$ excludes the cylinder body and the PIV-artifact mask. For the FHN blind-box settings, $\mathcal{U}$ includes both the central box and all other non-sensor coordinates.

For each test instance, let x denote the terminal reference field and $\{X^{(m)}\}_{m=1}^{M}$ the predicted ensemble, with

$$\bar{X} = \frac{1}{M} \sum_{m=1}^{M} X^{(m)}. \tag{74}$$

We use $M = 32$ unless otherwise specified.

Point-estimation accuracy. Full-field RMSE is evaluated over all valid coordinates, including sensor locations:

$$\mathrm{RMSE} = \left[\frac{1}{|\mathcal{E}|} \sum_{i \in \mathcal{E}} (\bar{X}_i - x_i)^2 \right]^{1/2}. \tag{75}$$

Relative L_2 error is evaluated only at non-sensor coordinates:

$$\mathrm{RelL}_2 = \frac{\|\bar{X}_{\mathcal{U}} - x_{\mathcal{U}}\|_2}{\|x_{\mathcal{U}}\|_2}. \tag{76}$$

Continuous ranked probability score. For $M \geq 2$, we use the fair ensemble estimator of CRPS Ferro (2014):

$$\widehat{\mathrm{CRPS}} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \left[\frac{1}{M} \sum_{m=1}^{M} |X_i^{(m)} - x_i| - \frac{1}{2M(M-1)} \sum_{\substack{m,m'=1 \\ m \neq m'}}^{M} |X_i^{(m)} - X_i^{(m')}| \right]. \tag{77}$$

For conditionally independent ensemble draws, this estimator is unbiased for the CRPS of the underlying predictive marginal distribution. For a deterministic prediction, CRPS reduces to absolute error, averaged over non-sensor coordinates.

Interval coverage and width. Let $\widehat{q}_{\alpha,i}$ denote the empirical α-quantile of the ensemble at coordinate i. We report 90% interval coverage in percent:

$$\mathrm{Cov90} = \frac{100}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{1}\left\{ \widehat{q}_{0.05,i} \leq x_i \leq \widehat{q}_{0.95,i} \right\}. \tag{78}$$

The corresponding mean interval width is

$$\mathrm{W90} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \left(\widehat{q}_{0.95,i} - \widehat{q}_{0.05,i} \right). \tag{79}$$

The nominal coverage is 90%. Coverage is interpreted jointly with interval width: undercoverage may reflect predictive bias, insufficient dispersion, or other distributional mismatch, while excessively wide intervals can yield high coverage without informative uncertainty estimates.

Aggregation. Each metric is first computed for an individual test instance and then averaged over the evaluated subset. Posterior draws are used to estimate the predictive distribution and are not counted as separate test instances.

F Additional experimental results

F.1 Sensor-budget and layout comparisons

Tables 3–6 complement the intermediate budgets in Table 1; together they cover all 11 available synthetic-data settings. Unavailable method–setting pairs are shown as "-". All relative L_2 values use non-sensor locations.

Table 3: Burgers with 8 and 32 fixed uniform sensors. Metric definitions and within-group highlighting follow Table 1.

		8 sensors				32 sensors			
	Method	RMSE	Rel. L_2	CRPS	Cov90	RMSE	Rel. L_2	CRPS	Cov90
Sparse only	Gaussian smoother	0.0591	0.2843	0.0544	68.3	0.0386	0.1894	0.0502	67.1
	EnKF	0.0333	**0.1743**	**0.0210**	64.5	0.0211	0.1195	**0.0125**	87.3
	nMA	0.2082	0.9726	0.1236	77.3	0.2379	1.1699	0.1904	14.7
	SDE Matching	0.3747	1.8153	0.2880	16.4	0.1572	0.7653	0.1146	14.1
	NGDB	**0.0325**	0.1774	0.0467	65.2	0.0238	0.1367	0.0159	79.3
	PICSB	0.0700	0.3047	0.0432	39.1	0.0379	0.1687	0.0229	56.5
	PRBridge (Sparse)	0.0343	0.1839	0.0250	**79.6**	**0.0206**	**0.1171**	0.0140	**87.5**
Dense supervision	DiffusionPDE	0.0705	0.3462	0.0490	64.2	0.0690	0.2957	0.0321	**91.9**
	SDA	0.0938	0.3818	0.0699	**90.0**	0.0343	0.1538	0.0305	87.0
	PalSB	0.0671	0.3139	0.0428	43.1	0.0338	0.1421	0.0281	42.0
	PRBridge (Dense)	**0.0327**	**0.1784**	**0.0231**	77.5	**0.0186**	**0.1028**	**0.0194**	85.9

Table 4: FHN with 1024 fixed uniform sensors (32×32). Metric definitions and within-group highlighting follow Table 1.

		1024 sensors (32×32)			
	Method	RMSE	Rel. L_2	CRPS	Cov90
Sparse only	Gaussian smoother	0.0419	0.1319	0.0540	67.5
	EnKF	0.0265	0.0797	0.0136	85.6
	nMA	0.0465	0.1472	0.0325	21.6
	SDE Matching	0.2257	0.6161	0.1690	3.4
	NGDB	0.0277	0.0792	**0.0132**	**91.9**
	PICSB	0.0755	0.2403	0.0477	44.8
	PRBridge (Sparse)	**0.0245**	**0.0771**	0.0142	84.8
Dense supervision	DiffusionPDE	0.4154	0.0544	0.0781	**83.3**
	SDA	0.0268	0.0808	0.0373	81.3
	PalSB	0.0329	0.0971	0.0344	40.9
	PRBridge (Dense)	**0.0240**	**0.0753**	**0.0133**	83.0

Table 5: Navier–Stokes with 1024 fixed uniform sensors (32×32). Metric definitions and within-group highlighting follow Table 1.

		1024 sensors (32×32)			
	Method	RMSE	Rel. L_2	CRPS	Cov90
Sparse only	Gaussian smoother	0.0967	0.0687	0.1626	75.8
	EnKF	0.0866	0.0638	**0.0266**	82.1
	nMA	0.6235	0.4422	0.4841	2.1
	SDE Matching	0.1737	0.1401	0.1503	29.3
	NGDB	0.0876	0.0692	0.0617	71.4
	PICSB	0.1190	0.0830	0.0727	58.6
	PRBridge (Sparse)	**0.0780**	**0.0561**	0.0519	**82.6**
Dense supervision	DiffusionPDE	0.1657	0.1228	0.0722	78.3
	SDA	0.0912	0.0639	**0.0299**	96.9
	PalSB	0.0895	0.0656	0.0494	**86.3**
	PRBridge (Dense)	**0.0712**	**0.0505**	0.0392	85.9

Table 6: FHN with fixed random sensors outside a central 64×64 blind box. Metric definitions and within-group highlighting follow Table 1. The non-sensor metrics cover all non-sensor locations, not only the central box.

		64 sensors				**256 sensors**			
	Method	RMSE	Rel. L_2	CRPS	Cov90	RMSE	Rel. L_2	CRPS	Cov90
Sparse only	EnKF	0.1785	0.5638	0.0886	80.3	0.1397	0.4103	0.0727	78.0
	nMA	0.2700	0.8454	0.2146	13.0	0.1960	0.6164	0.1367	18.8
	SDE Matching	0.2510	0.7867	0.1856	16.3	0.2284	0.7167	0.1706	14.4
	NGDB	0.1770	0.5607	**0.0859**	84.2	0.1337	0.4056	**0.0564**	**82.1**
	PICSB	0.2910	0.9104	0.1962	13.6	0.2001	0.6288	0.1195	23.6
	PRBridge (Sparse)	**0.1736**	**0.5536**	0.1042	**85.3**	**0.1313**	**0.3811**	0.0672	63.6
Dense supervision	DiffusionPDE	0.2458	0.7707	0.1312	81.5	0.2382	0.7509	0.1277	82.1
	SDA	1.0448	2.3194	0.6039	70.3	0.1471	0.4781	0.1044	71.9
	PRBridge (Dense)	**0.1585**	**0.5394**	**0.0889**	**83.6**	**0.1206**	**0.3780**	**0.0539**	**84.3**

G LIMITATIONS

Model assumptions and recoverability. Our formulation assumes a known complete initial field, specified physical parameters, and prescribed process- and observation-noise models. Uncertain initial conditions or misspecified dynamics may affect both reconstruction accuracy and uncertainty estimates. Physical evolution also does not make every unobserved structure identifiable: perturbations that remain invisible to the observation history can still alter the terminal field. Our reconstruction bound is local and relies on smoothness and posterior-concentration assumptions, rather than guaranteeing recovery of arbitrary fine-scale structures.

Posterior approximation and calibration. The Gaussian reference uses a smoothing mean and a designed temporal covariance, while the transport has a restricted local parameterization. Long-window supervision additionally relies on finite ensembles and approximate Kalman analyses. The resulting targets provide auxiliary supervision rather than exact posterior samples, and their inclusion does not guarantee a decrease in the path-space KL.

Distillation and computational cost. The one-step student approximates the Stage 1 teacher distribution and can therefore inherit teacher errors while introducing additional approximation error. Our experiments show that similar relative L_2 errors can coexist with substantially reduced interval coverage after distillation. Although the student avoids reference preparation and ensemble analysis at inference, these computations remain necessary during training. Preserving distributional quality while reducing the overall computational cost remains an open direction.

Experimental scope and supervision. Our evaluation covers one- and two-dimensional systems with sensor layouts fixed within each setting. Generalization to unseen layouts, substantially sparser temporal observations, and large three-dimensional systems has not been established. For Cylinder, Gaussian-reference hyperparameters are selected using dense validation fields. Consequently, the sparse-trained variants avoid post-initial full-field targets in their neural training objectives, but the complete model-selection procedure still uses dense field information. Reference selection based entirely on sparse measurements would further broaden applicability.