

FLAPS

Flow Annealing Posterior Sampling for Function-Space Regression and Inverse Problems

Yaozhong Shi

California Institute of Technology

Joint work with Zachary E. Ross and Yisong Yue

Group meeting

[arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2)

github.com/yzshi5/FLAPS

NeurIPS 2026

Bayesian inference from noisy measurements

Unknown function or field u , observations y , and a measurement model:

$$y = \mathcal{A}(u) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Gamma).$$

Prior and likelihood


$\mu(du)$ describes plausible functions before observing y .

$p(y | u)$ measures agreement with the data under the noise model.

$$\pi^y(du) = \frac{1}{Z_y} p(y | u) \mu(du).$$

Posterior samples quantify uncertainty.

Bayesian formulation: [FLAPS](#), [arXiv:2606.22346v2](#), Sec. 3.2. The scalar plot is an explanatory example.



Scalar illustration: $y = u + \varepsilon$.

Linear and nonlinear inverse problems

Linear Gaussian setting ($u \in \mathbb{R}^n$)

$$u \sim \mathcal{N}(m_0, C_0), \quad y = Au + \varepsilon.$$

The posterior remains Gaussian:

$$u \mid y \sim \mathcal{N}(m_{\text{post}}, C_{\text{post}}),$$
$$C_{\text{post}} = (C_0^{-1} + A^\top \Gamma^{-1} A)^{-1}.$$


Closed-form conditioning gives the full posterior.

Linearity refers to the map $u \mapsto y$.

Linear Gaussian conditioning and the nonlinear plot are explanatory examples.

Nonlinear setting

$$u \sim \mathcal{N}(0, 1), \quad y = u^2 + \varepsilon.$$



$$y = 1, \text{ noise sd } 0.15.$$

Both signs can explain the observation.

Image restoration: a linear forward model

$$y = Au + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma_y^2 I).$$

Different tasks change the operator


Task	Operator A
Denoising	Identity I
Inpainting	Pixel selection P_Ω
Deblurring	Known convolution H
Super-resolution	Blur and downsampling DH

$$\nabla_u \log p(y | u) = \sigma_y^{-2} A^\top (y - Au).$$

A linear map can still yield a non-Gaussian image posterior.

Pixel-mask example: [FLAPS](#), [arXiv:2606.22346v2](#), Fig. 3(c). Other degradation operators are illustrative.

Partial pixel observations



Black-hole images in FLAPS: 256 observed pixels on a 64×64 grid.

Seismic imaging: a nonlinear forward model

Infer a subsurface velocity field $c(x)$ from waveforms recorded at receivers.

$$m(x) = c(x)^{-2}, \quad m(x) \partial_{tt} w_s(x, t) - \Delta w_s(x, t) = f_s(x, t).$$

Wave propagation

For fixed sources and boundary conditions,

$$w_s = \mathcal{G}_s(m), \quad \mathcal{A}(m) = \{P_R \mathcal{G}_s(m)\}_s.$$

P_R extracts receiver waveforms, and s indexes sources.

Bayesian imaging with FLAPS

$$\pi^y(dm) \propto \exp \left[-\frac{\|y - \mathcal{A}(m)\|_2^2}{2\sigma_y^2} \right] \mu_\theta(dm).$$

Learn a prior on subsurface fields.

Condition it using a differentiable wave solver or forward surrogate.

Posterior samples quantify imaging uncertainty.

The PDE is linear in w_s , but $m \mapsto \mathcal{A}(m)$ is nonlinear.

Acoustic imaging illustration with fixed sources and density. The FLAPS paper does not report this experiment.

Gaussian process regression

A Gaussian process specifies a distribution over functions:

$$f \sim \text{GP}(m, k), \quad y_i = f(x_i) + \varepsilon_i.$$

Conditioning at query points Q

$$f_Q \mid y \sim \mathcal{N}(m_Q^y, K_Q^y),$$

$$m_Q^y = m_Q + K_{QO}B^{-1}(y - m_O),$$

$$K_Q^y = K_{QQ} - K_{QO}B^{-1}K_{OQ}, \quad B = K_{OO} + \sigma_y^2 I.$$


Here O denotes observed locations.

Linear Gaussian observations give

closed-form posteriors.

Rasmussen & Williams, Gaussian Processes for Machine Learning (2006). See also FLAPS, App. B.2.

Yaozhong Shi



Mean, posterior samples, and ± 2 sd.
Illustrative RBF GP with seven observations.

General stochastic-process regression

Scientific priors can be non-Gaussian


Interfaces, coherent flow structures, and image morphology can require richer joint distributions.

Sparse measurements may leave several plausible functions.

A posterior mean can blur or cancel those alternatives. **Goal: condition an**

expressive distribution over functions on new data.

An ambiguous functional posterior



Equal mass on teal functions $u_{\pm} = \pm \sin(\pi x)$.
Dashed: posterior mean. Dots: observations.

Motivation: [FLAPS](#), [arXiv:2606.22346v2](#), Introduction. The two-function mixture is an explanatory example.

Regression and PDE inversion share one posterior

$$\pi^y(du) \propto \exp\left[-\frac{\|y - \mathcal{A}(u)\|_2^2}{2\sigma_y^2}\right] \mu(du).$$

Task	Unknown	Observation operator
Functional regression	Function values $u(x)$	$\mathcal{A}(u) = P_\Omega u$
PDE inverse problem	Input field $u(x)$	$\mathcal{A}(u) = P_\Omega \mathcal{G}_\phi(u)$

Example: infer Darcy permeability from sparse solution measurements

$$-\nabla \cdot (u(x) \nabla v(x)) = f(x), \quad v|_{\partial D} = 0, \quad v = \mathcal{G}_\phi(u).$$

The PDE is linear in v , while the coefficient-to-solution map $u \mapsto v$ is nonlinear.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Sec. 3.2, Eqs. (7)–(10). P_Ω extracts measurements at observed locations.

Learning a reusable flow prior

Operator Flow Matching learns a distribution over functions

$$u_0 \sim \gamma = \mathcal{N}(0, C_0), \quad u_1 \sim \mu_{\text{data}}, \quad u_t = tu_1 + (1 - t)u_0.$$

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t, u_0, u_1} \left[\|v_\theta(t, u_t) - (u_1 - u_0)\|_{\mathcal{U}}^2 \right].$$


$$\frac{du_t}{dt} = v_\theta(t, u_t), \quad \mu_\theta = (\Phi_{0 \rightarrow 1}^\theta)_\# \gamma.$$

- Train the prior once on unconditional function samples.
- Evaluate a process prior on different query sets $X = \{x_i\}$.
- Freeze the prior when conditioning on new measurements.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Sec. 3.1. The displayed path omits the small residual noise used to regularize training. FLAPS also supports conventional fixed-grid flow priors.

FLAPS posterior sampling

$$0 = t_0 < t_1 < \dots < t_K = 1, \quad u_{t_0} \sim \gamma.$$



Re-bridging uses $u_{t_{k+1}} = t_{k+1}u_k^{\text{clean}} + (1 - t_{k+1})\xi$, $\xi \sim \gamma$.

FLAPS, arXiv:2606.22346v2, Fig. 1 (posterior-sampling panel), Algorithm 1 and App. A.1. At the last level, return u_k^{clean} .

One FLAPS annealing level

At level t_k , the current state mixes a candidate field with reference noise.

1. Flow transport supplies a plausible endpoint

$$\hat{u}_k = \Phi_{t_k \rightarrow 1}^\theta(u_{t_k}).$$

The frozen prior gives an anchor for the clean field.

2. Likelihood correction incorporates the measurements

$$u_k^{\text{clean}} \sim \tilde{\pi}_k, \quad \tilde{\pi}_k(u) \propto p(y \mid u) \mathcal{N}(u; \hat{u}_k, \lambda_k^2 C_X).$$

Langevin updates balance data agreement with the local prior approximation.

3. Re-bridging prepares the next level

$$u_{t_{k+1}} = t_{k+1} u_k^{\text{clean}} + (1 - t_{k+1}) \xi, \quad \xi \sim \gamma.$$

Fresh noise prepares another refinement at a lower noise level.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Sec. 3.3 and Algorithm 1. Return the final corrected clean sample at $t_K = 1$.

Likelihood-guided local correction

The flow predicts an endpoint anchor:

$$\hat{u}_k = \Phi_{t_k \rightarrow 1}^\theta(u_{t_k}).$$

FLAPS approximates the clean conditional around that anchor:

$$\tilde{\pi}_k(u) \propto p(y | u) \mathcal{N}(u; \hat{u}_k, \lambda_k^2 C_X).$$

$$u^{\ell+1} = u^\ell + \eta \left[-\frac{u^\ell - \hat{u}_k}{\lambda_k^2} + C_X \nabla_u \log p(y | u^\ell) \right] + \sqrt{2\eta} \zeta^\ell, \quad \zeta^\ell \sim \mathcal{N}(0, C_X).$$

$$\nabla_u \log p(y | u) = \frac{1}{\sigma_y^2} J_{\mathcal{A}}(u)^*(y - \mathcal{A}(u)).$$

For PDE inversion, differentiate through the frozen forward model $\mathcal{G}_\phi$.

The correction avoids evaluating the learned prior density.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Sec. 3.3, Eqs. (12)–(18). The local Gaussian conditional and finite Langevin steps are approximations.

Covariance preconditioning

Estimate C_X offline from clean samples of the learned prior:

$$C_X \approx Q_r \Lambda_r Q_r^\top + \sigma_{\text{res}}^2 I.$$

Correlated updates from sparse data




Illustration using a correlated covariance.

Matérn GP ablation

Preconditioner	SWD ↓	MMD ↓
Identity	0.663	0.508
Rank 4	0.0626	0.0420
Rank 8	0.0470	0.0398
Rank 32	0.0356	0.0172
Rank 64	0.0450	0.0302

512 queries. 128 posterior samples.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Sec. 3.4 and Table 11. C_X estimates covariance of the target process.

Annealing and the ideal posterior identity

The noisy posterior marginal is a bridge mixture of the clean posterior:

$$\pi_t^y(du_t) = \int q_t(du_t \mid u_1) \pi_1^y(du_1), \quad q_t(\cdot \mid u_1) = \mathcal{N}(tu_1, (1-t)^2 \Sigma_0^X).$$

Exact conditional sampling preserves the intended marginals

If $u_{t_k} \sim \pi_{t_k}^y$, perform

$$u_1 \sim \pi_1^y(\cdot \mid u_{t_k}), \quad u_{t_{k+1}} \sim q_{t_{k+1}}(\cdot \mid u_1).$$

Then $u_{t_{k+1}} \sim \pi_{t_{k+1}}^y$, and the endpoint is the posterior.

Practical FLAPS approximates the clean conditional.


Multiple levels allow repeated local correction as the bridge noise decreases. The identity alone does not prove exactness of the practical sampler.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Proposition 1 and App. A.1–A.3. Posterior sampling uses the residual-noise-free bridge $s(t) = 1 - t$.

GP regression with an exact posterior reference

Same seven observations, with inference at different query resolutions.

FLAPS at 1024 query locations



Orange: FLAPS samples and uncertainty.

Blue: exact GP posterior mean and uncertainty.


FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Fig. 2 (FLAPS, 1024-query panel) and Table 1. SWD compares generated and exact-reference posterior samples.

Posterior SWD ↓

Query size	128	512
CNP	0.202	0.201
ANP	0.204	0.203
NDP	0.265	0.465
FlowNP	0.224	0.266
OFM posterior	0.217	0.292
FLAPS	0.142	0.147

Table 1, selected test-set means.

Regression on non-Gaussian scientific fields



Navier–Stokes vorticity from sparse, noisy direct observations.


Dataset	Best NP CRPS ↓	FLAPS CRPS ↓	FLAPS SSR
Navier–Stokes	0.0737	0.0279	1.01
Black hole	0.0149	0.0126	0.727
Global climate	0.0598	0.0228	0.890

CRPS ↓: predictive quality. SSR: spread/error, best near 1.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Fig. 3(a), Table 2. Best NP: minimum CRPS across CNP, ANP, NDP, and FlowNP.

PDE inverse problems and sampling cost

Darcy example: 128 noisy solution observations on a 128×128 grid (0.78%).



Input-field relative L^2 error ↓

PDE	Best baseline	FLAPS-FNO
Darcy	0.606	0.866
Poisson	0.562	0.515
Helmholtz	0.581	0.501
Navier–Stokes	0.330	0.411

Poisson runtime (s/case)

Method	32 samples
FLAPS-FNO	65.1
DiffusionPDE	112.5
FunDPS	100.7
DDIS	153.9

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Fig. 4(a), Table 3 and Table 10. Best baseline: lowest relative L^2 among DiffusionPDE, FunDPS, and DDIS.

Zero-shot super-resolution with a function prior

A process prior describes a function $u(x)$ across query resolutions.

$$y = \mathcal{A}(u) + \varepsilon, \quad \pi^y(du) \propto p(y | u) \mu_\theta(du).$$

Condition on coarse or sparse data

For direct reconstruction, choose

$$\mathcal{A}(u) = P_\Omega u \quad \text{or} \quad DHu.$$

For PDE inversion, choose

$$\mathcal{A}(u) = P_\Omega \mathcal{G}_\phi(u).$$

Evaluate the posterior on a finer query set X_{fine} using the same learned prior.

What zero-shot means here

Example	Training grid	Query grid
Matérn GP	128	512 / 1024
Darcy inverse	128 ²	160 ²


No retraining for the finer query grid.

Cross-resolution inference uses the function-space FLAPS-FNO prior. A fixed-grid UNet prior does not provide it automatically.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Fig. 2, App. B.1, B.5 and G. The covariance and forward discretization must match the inference grid.

Zero-shot Darcy inversion on a finer grid

Train prior and forward surrogate at 128×128 . Infer at 160×160 without retraining.



128 noisy solution measurements: 0.5% of the 160×160 grid.

- Posterior samples retain sharp coefficient interfaces.
- Uncertainty grows near interfaces and other ambiguous regions.
- Posterior predictive solutions agree with the observed solution values.

FLAPS, [arXiv:2606.22346v2](https://arxiv.org/abs/2606.22346v2), Fig. 9 and App. G. Qualitative demonstration of inference on an unseen query resolution.

Reusable priors for Bayesian function inference

$$\pi^y(du) \propto p(y | u) \mu_\theta(du).$$

- **One Bayesian formulation**

Regression and PDE inversion differ through the observation operator.

- **A reusable learned prior**

Reuse the function prior for zero-shot inference on finer query grids.

- **Approximate posterior sampling**

Alternate frozen flow transport, likelihood correction, and re-bridging.

Open questions: local conditional accuracy, mixing across modes, and forward-model error.

Paper: [arXiv:2606.22346](https://arxiv.org/abs/2606.22346)

Code: github.com/yzshi5/FLAPS