

# SNaP: One-Step Posterior Sampling for Noisy Inverse Problems

Shirin Shoushtari<sup>1\*</sup> Edward P. Chandler<sup>2\*</sup> Xiao Shi<sup>2</sup> Ulugbek S. Kamilov<sup>2</sup>

<sup>1</sup>WashU, <sup>2</sup>UW-Madison

s.shirin@wustl.edu, epchandler@wisc.edu, xiao.shi@wisc.edu, kamilov@wisc.edu

\*Equal contribution

## Abstract

Diffusion and flow-matching models can produce high-quality posterior samples for inverse problems, but typically require tens to thousands of network evaluations per draw. MeanFlow enables one-step generation, yet applying it to inverse problems leaves no intermediate steps at which to enforce measurement consistency. We introduce SNaP, a *one-step* MeanFlow posterior sampler for linear inverse problems with Gaussian noise. Its central innovation is a measurement-adapted source: a Gaussian distribution whose mean and anisotropic covariance are determined by the measurement operator, observation, and noise level. The source anchors well-measured directions while preserving variation where the measurements are weak or uninformative. We show that the exact conditional flow transports this source to the true posterior. Across natural-image restoration and multi-coil MRI, SNaP produces diverse, high-quality samples with one network evaluation per draw, 30 to 2250 $\times$  faster than iterative samplers.

## 1 Introduction

Many imaging tasks require recovering an unknown image  $\mathbf{x}_1$  from noisy linear measurements  $\mathbf{y} = \mathbf{A}\mathbf{x}_1 + \mathbf{n}$ , where  $\mathbf{A}$  models the imaging system and  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$  denotes additive Gaussian noise. This formulation encompasses magnetic resonance imaging (MRI), computed tomography (CT), microscopy, astronomical imaging, and remote sensing [1, 2, 3]. Because  $\mathbf{A}$  is often ill-conditioned or rank-deficient, the measurements do not uniquely determine  $\mathbf{x}_1$ . Therefore, we need prior information on images for a reliable recovery. Classical methods impose handcrafted regularizers such as total variation [4]. More recently, deep generative models, including diffusion models [5, 6] and flow-matching models [7, 8, 9], have learned rich image distributions that can serve as priors for inverse problems [10].

Most generative inverse solvers use an unconditional generative model as an image prior and impose measurement consistency repeatedly during an iterative sampling trajectory. One family modifies the generative dynamics using data-fidelity gradients or pseudoinverse corrections [11, 12, 13, 14, 15, 16]. Another alternates generative updates with explicit data-consistency steps, following the plug-and-play paradigm [17, 18, 19, 20, 21]. Although these corrections improve agreement with the measurement  $\mathbf{y}$ , they require many sequential network evaluations and therefore make posterior sampling slow. Depending on the method, generating a single sample may require tens to thousands of network evaluations (NFEs). In our CelebA Gaussian-deblurring experiment, for example, DPS [11] requires 1000 NFEs and 27 seconds to produce one (128  $\times$  128) reconstruction (Figure 1).

Consistency models [22, 23], Shortcut Models [24], and MeanFlow[25] enable generation with one network evaluation, avoiding the cost of integrating an iterative sampling trajectory. This efficiency is particularly attractive for posterior sampling, where generating a single reconstruction can require tens to thousands of network evaluations. However, collapsing the trajectory into one step also eliminates the intermediate states for enforcing measurement consistency. A direct extension is to condition MeanFlow on the measurement while retaining the standard isotropic source. Recent studies of




Figure 1: *Left*: LPIPS versus wall-clock time per sample for CelebA Gaussian deblurring. SNaP achieves lowest LPIPS at a fraction of the cost of iterative samplers. *Right*: 4 $\times$  super-resolution on AFHQ-Cat. Columns show one SNaP draw, averages of  $M = 4, 16$ , and 100 draws, and the ground truth; zoomed crops, error maps, and PSNR are shown below. Each draw requires one network evaluation. Averaging reduces pixel error and improves PSNR while smoothing fine detail.

one-step generation have reported mean-seeking bias in MeanFlow and averaging-induced diversity degradation in class-conditional flow models [26, 27]. We observe the same failure mode in inverse problems: despite conditioning on the measurement, a MeanFlow with an isotropic source produces nearly identical reconstructions across draws, so averaging offers no improvement (Table 5). This motivates shaping the one-step transport to reflect the geometry of the forward model.

We propose SNaP, a one-step sampler for *noisy* linear inverse problems. Its key innovation is a *measurement-adapted* Gaussian source that builds the forward model and noise level into the starting point of the transport. The source anchors well-measured directions while retaining stochastic variation in weakly measured and null-space directions. Samples from this source can be drawn efficiently, and a single network evaluation maps each draw to a reconstruction.

Our contributions are: **(a)** We extend one-step MeanFlow sampling to noisy linear inverse problems by constructing a measurement-adapted transport aligned with the geometry of the likelihood. **(b)** We derive the source center by minimizing its expected squared displacement to posterior samples among linear estimators. The resulting construction uses a Tikhonov estimate to anchor the transport. **(c)** Across natural-image restoration and multi-coil MRI, SNaP produces diverse, high-quality reconstructions with one network evaluation per draw. Its samples achieve near-unit calibration ratios on the tested CelebA tasks, while sampling 30 to 2250 $\times$  faster than iterative posterior samplers.

## 2 Background

**Inverse Problems.** We consider a linear measurement operator  $\mathbf{A} \in \mathbb{R}^{m \times n}$  with additive white Gaussian noise,

$$\mathbf{y} = \mathbf{A}\mathbf{x}_1 + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}). \quad (1)$$

The goal is to infer the unknown image  $\mathbf{x}_1 \in \mathbb{R}^n$  from the measurement  $\mathbf{y} \in \mathbb{R}^m$ ; in posterior sampling, this means drawing samples from  $p(\mathbf{x}_1 | \mathbf{y})$ . We write the singular value decomposition of  $\mathbf{A}$  as  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  and let  $\rho := \text{rank}(\mathbf{A})$ . We order the singular values so that  $s_1 \geq \dots \geq s_\rho > 0$  and set  $s_i = 0$  for  $\rho < i \leq n$ . The corresponding right singular vectors are the columns of  $\mathbf{V}$ . We denote the pseudoinverse by  $\mathbf{A}^\dagger$  and the orthogonal projection onto  $\text{null}(\mathbf{A})$  by  $\mathbf{P} := \mathbf{I} - \mathbf{A}^\dagger \mathbf{A}$ .

**Flow Matching.** Flow matching [7, 8, 9] learns a velocity field that transports a source distribution  $p_0$  to a data distribution  $p_1$ . We index time so that  $t = 0$  corresponds to the source and  $t = 1$  to the data. The flow  $\psi_t$  is defined by the ODE  $\frac{d}{dt}\psi_t(\mathbf{x}) = \mathbf{v}_t(\psi_t(\mathbf{x}))$  for  $t \in (0, 1)$ , and it maps samples from  $p_0$  to samples from  $p_1$ . For the exact marginal velocity field, the terminal map satisfies  $(\psi_1)_\# p_0 = p_1$ . This marginal velocity is generally not available in closed form, so flow matching instead regresses a sample-conditional velocity. Given a coupling of endpoints  $(\mathbf{x}_0, \mathbf{x}_1)$  with  $\mathbf{x}_0 \sim p_0$  and  $\mathbf{x}_1 \sim p_1$ , the straight-line path

$$\mathbf{z}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1, \quad \mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0, \quad (2)$$

has constant conditional velocity  $\mathbf{v}$ . A network  $\mathbf{v}^\theta$  is trained by minimizing  $\mathbb{E} \|\mathbf{v}^\theta(\mathbf{z}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|_2^2$ . The expected parameter gradient of this conditional regression objective coincides with that of the marginal flow-matching objective [7]. The source is commonly chosen as  $\mathcal{N}(\mathbf{0}, \mathbf{I})$  because it



is easy to sample; its random initial conditions provide the variability propagated by the deterministic ODE.

**Mean Flows.** Sampling a flow-matching model requires integrating the velocity  $\mathbf{v}$ , which takes many network evaluations. MeanFlow [25, 23] avoids this integration by modeling the average velocity  $\mathbf{u}$  over an interval  $[r, t]$ , defined as

$$\mathbf{u}(\mathbf{z}_r, r, t) := \frac{1}{t-r} \int_r^t \mathbf{v}(\mathbf{z}_s, s) ds. \quad (3)$$

Once  $\mathbf{u}$  is learned, the entire flow reduces to a single evaluation,  $\mathbf{z}_1 = \mathbf{z}_0 + \mathbf{u}(\mathbf{z}_0, 0, 1)$ . Learning  $\mathbf{u}$  directly from equation 3 is intractable, since it requires the integral of  $\mathbf{v}$ . Instead, rewriting equation 3 as  $(t-r)\mathbf{u} = \int_r^t \mathbf{v} ds$  and differentiating with respect to  $r$ , with  $t$  fixed, gives the MeanFlow identity

$$\mathbf{u}(\mathbf{z}_r, r, t) = \mathbf{v}(\mathbf{z}_r, r) + (t-r) \frac{d}{dr} \mathbf{u}(\mathbf{z}_r, r, t). \quad (4)$$

This differential identity provides the training relation used by MeanFlow. It also connects MeanFlow to consistency and shortcut models, which learn finite-time flow maps through related consistency relations [22, 23, 24].

**Measurement-dependent sources.** Standard MeanFlow formulations use a fixed, measurement-independent source. Applying MeanFlow to inverse problems therefore requires incorporating the measurement model into the one-step transport. NullFlow [28] is a one-step posterior sampler for noiseless linear inverse problems that uses a measurement-dependent source:

$$\mathbf{x}_0 = \mathbf{A}^\dagger \mathbf{y} + \mathbf{P}\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (5)$$

For a noiseless measurement, the source satisfies  $\mathbf{A}\mathbf{x}_0 = \mathbf{y}$ , and restricting the velocity to  $\text{null}(\mathbf{A})$  keeps the entire trajectory in the measurement-consistent affine set  $\{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{y}\}$ ; however, with noisy measurements, the posterior is no longer supported on this set. Moreover, the anchor  $\mathbf{A}^\dagger \mathbf{y}$  contains the row-space noise term  $\mathbf{A}^\dagger \mathbf{n}$ , which a velocity restricted to  $\text{null}(\mathbf{A})$  cannot alter.

### 3 Method

SNaP builds the geometry of the forward model into a one-step posterior transport. Its measurement-dependent source is designed to keep the path to the posterior short while preserving the randomness needed for sampling. It concentrates draws in well-measured directions and retains variation in weakly measured and null-space directions. A conditional MeanFlow then learns to transport these draws toward  $p(\mathbf{x}_1 | \mathbf{y})$ , producing an approximate posterior sample from one source draw and one network evaluation. We now derive the source, characterize the exact transport, and describe how the model is trained and sampled.

Along the straight-line path between a source sample  $\mathbf{x}_0$  and a posterior sample  $\mathbf{x}_1$ , the sample-conditional velocity is  $\mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0$ , and therefore  $\|\mathbf{v}\|_2^2 = \|\mathbf{x}_1 - \mathbf{x}_0\|_2^2$  [8]. We keep the path and the conditional-independence structure of the endpoint coupling fixed and use the source distribution as our design variable. This differs from coupling-based methods that modify the endpoint pairing [29, 30, 31]. Unlike approaches that prescribe or learn data-dependent sources [32, 33], we derive the source from the measurement model using an explicit design criterion.

Although the posterior is unavailable in closed form, the likelihood provides a tractable geometry for images compatible with the measurement. Under additive Gaussian noise,

$$p(\mathbf{y} | \mathbf{x}) \propto \exp\left(-\frac{\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2}{2\sigma_n^2}\right),$$

and is maximized on the least-squares solution set

$$\mathcal{S}_\mathbf{y} := \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2.$$

Its superlevel sets are convex tubes around  $\mathcal{S}_\mathbf{y}$ . Along a measured right singular vector  $\mathbf{v}_i$ , their width scales as  $\sigma_n/s_i$ : well-measured directions are tightly constrained, weakly measured directions are less constrained, and null-space directions are unconstrained (Fig. 2a). The noise level sets the scale




Figure 2: Geometry-aware one-step transport (schematic). **Left:** Likelihood superlevel sets form tubes around  $\mathcal{S}_y$ : narrow in well-measured directions, wider in weakly measured directions, and unbounded along  $\text{null}(\mathbf{A})$ . Gray points represent posterior samples. **Middle:** An isotropic source spreads samples across directions regardless of measurement strength, leading to large source-to-target displacements. **Right:** The SNaP source concentrates samples in well-measured directions while retaining variation in weakly measured and null-space directions, shortening displacements.

because  $\mathbf{A}\mathbf{x}_1 - \mathbf{y} = -\mathbf{n}$ . Since these tubes are convex, a straight path between two points in the same tube stays inside it. This geometry motivates a source with little variation in well-measured directions and greater variation in weakly measured and null-space directions (Fig. 2c), in contrast to an isotropic source (Fig. 2b). We construct this source from the observed  $\mathbf{y}$ , known  $\mathbf{A}$  and  $\sigma_n$ , and auxiliary randomness. We consider sources with an explicit, linear measurement-dependent center:

$$\mathbf{x}_0 = \mathbf{K}\mathbf{y} + \boldsymbol{\xi}, \quad (6)$$

where  $\mathbf{K}$  is fixed for a given measurement model and  $\boldsymbol{\xi}$  has mean  $\mathbf{0}$  and covariance  $\boldsymbol{\Sigma}$ , independently of  $(\mathbf{x}_1, \mathbf{n})$ . This choice admits an exact analysis and avoids training a separate source model [33]. The family contains the standard isotropic source when  $\mathbf{K} = \mathbf{0}$  and  $\boldsymbol{\Sigma} = \mathbf{I}$ , and sources centered at the least-squares solution when  $\mathbf{K} = \mathbf{A}^\top$ .

Because the one-step map is deterministic given  $(\mathbf{x}_0, \mathbf{y})$ , all sampling randomness originates from the source. For the family in equation 6, its expected squared displacement decomposes as

$$\mathbb{E}[\|\mathbf{x}_1 - \mathbf{x}_0\|_2^2] = \text{tr}(\mathbf{R}(\mathbf{K})) + \text{tr}(\boldsymbol{\Sigma}), \quad \mathbf{R}(\mathbf{K}) := \mathbb{E}[(\mathbf{x}_1 - \mathbf{K}\mathbf{y})(\mathbf{x}_1 - \mathbf{K}\mathbf{y})^\top]. \quad (7)$$

For fixed  $\boldsymbol{\Sigma}$ , minimizing the displacement reduces to minimizing the center error  $\text{tr}(\mathbf{R}(\mathbf{K}))$  over  $\mathbf{K}$ . Minimizing over  $\boldsymbol{\Sigma}$  would instead give  $\boldsymbol{\Sigma} = \mathbf{0}$  and eliminate source variability, so we optimize the center and choose the spread separately.

**Proposition 1.** *Let  $\mathbf{x}_1$  have mean  $\mathbf{0}$  and covariance  $\mathbf{C}$ , not necessarily Gaussian, and let  $\mathbf{y} = \mathbf{A}\mathbf{x}_1 + \mathbf{n}$  with  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$  and  $\sigma_n > 0$ . For any fixed  $\boldsymbol{\Sigma} \succeq \mathbf{0}$ , the expected squared displacement of the source family in equation 6 is uniquely minimized over  $\mathbf{K}$  by*

$$\mathbf{K}_C := \mathbf{C}\mathbf{A}^\top (\mathbf{A}\mathbf{C}\mathbf{A}^\top + \sigma_n^2 \mathbf{I})^{-1}. \quad (8)$$

$\mathbf{K}_C \mathbf{y}$  is thus the linear minimum-mean-squared-error estimator of  $\mathbf{x}_1$  from  $\mathbf{y}$ , with error covariance

$$\boldsymbol{\Sigma}_C := \mathbf{R}(\mathbf{K}_C) = \mathbf{C} - \mathbf{K}_C \mathbf{A} \mathbf{C}.$$

For the isotropic working covariance  $\mathbf{C} = \tau^2 \mathbf{I}$ , with  $\tau > 0$ , the optimal center is the Tikhonov estimate  $\mathbf{m}_y = \mathbf{A}^\top (\mathbf{A}\mathbf{A}^\top + \lambda \mathbf{I})^{-1} \mathbf{y}$ , with  $\lambda = \sigma_n^2 / \tau^2$ , and its error covariance is  $\tau^2 \mathbf{W}$ , where  $\mathbf{W} = \mathbf{I} - \mathbf{A}^\top (\mathbf{A}\mathbf{A}^\top + \lambda \mathbf{I})^{-1} \mathbf{A}$ .

The proof is given in App. A.1. Minimizing the displacement over  $\boldsymbol{\Sigma}$  would give  $\boldsymbol{\Sigma} = \mathbf{0}$  and collapse the source. For a chosen working covariance  $\mathbf{C}$ , we instead set  $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_C$ , matching the error covariance of the optimal linear estimator. Under the corresponding Gaussian working model, this is also the posterior covariance.

Because the true image covariance is unknown, we instantiate Proposition 1 with the isotropic working covariance  $\mathbf{C} = \tau^2 \mathbf{I}$ , where  $\tau$  controls the source scale. Under this model,  $\mathbf{m}_y$  is the displacement-minimizing linear center and  $\tau^2 \mathbf{W}$  is its error covariance. The resulting Tikhonov center regularizes weakly measured directions and avoids amplifying their measurement noise, unlike the least-squares estimate  $\mathbf{A}^\dagger \mathbf{y}$ . The working model is used only to construct the source; it does not assume that the true image distribution is isotropic or Gaussian. Using this mean and covariance, we define

$$q_\tau(\mathbf{x} | \mathbf{y}) = \mathcal{N}(\mathbf{x}; \mathbf{m}_y, \tau^2 \mathbf{W}). \quad (9)$$

In the right singular basis of  $\mathbf{A}$ ,

$$\mathbf{W} = \mathbf{V} \text{diag}(w_i) \mathbf{V}^\top, \quad w_i = \frac{\sigma_n^2}{\sigma_n^2 + \tau^2 s_i^2}. \quad (10)$$

The source variance along  $\mathbf{v}_i$  is therefore  $\tau^2 w_i$ . It is small in well-measured directions, approaches  $\sigma_n^2/s_i^2$  when  $\tau^2 s_i^2 \gg \sigma_n^2$ , and equals  $\tau^2$  on  $\text{null}(\mathbf{A})$ . Thus, the source suppresses variation in well-measured directions while retaining it where the measurement is weak or uninformative. As  $\sigma_n \rightarrow 0$ , the isotropic working source converges to  $q_\tau(\cdot | \mathbf{y}) \rightarrow \mathcal{N}(\mathbf{A}^\dagger \mathbf{y}, \tau^2 \mathbf{P})$ , which recovers the NullFlow source when  $\tau = 1$ .

Replacing the isotropic source by the working posterior changes the transport, but not its target.

**Proposition 2.** *Let  $\sigma_n > 0$ , and let  $\mathbf{x}_0 \sim q_\tau(\cdot | \mathbf{y})$  and  $\mathbf{x}_1 \sim p(\cdot | \mathbf{y})$  be conditionally independent given  $\mathbf{y}$  and joined by the straight-line path in equation 2. Assume that  $p(\cdot | \mathbf{y})$  has finite second moments, and that the flow of the marginal velocity field has a terminal limit: writing  $\psi_{0 \rightarrow t}$  for its flow map from time 0 to time  $t$ , the limit  $\psi_{0 \rightarrow 1} := \lim_{t \rightarrow 1} \psi_{0 \rightarrow t}$  exists  $q_\tau(\cdot | \mathbf{y})$ -almost everywhere. Then the one-step map*

$$T_{\mathbf{y}}(\mathbf{x}) := \mathbf{x} + \mathbf{u}(\mathbf{x}, 1 | \mathbf{y}) \quad (11)$$

satisfies

$$(T_{\mathbf{y}})_\# q_\tau(\cdot | \mathbf{y}) = p(\cdot | \mathbf{y}) \quad \text{for every } \tau > 0. \quad (12)$$

Here  $\mathbf{u}$  is the exact average velocity of equation 3; the trained network  $\mathbf{u}^\theta$  approximates it, so the guarantee holds for SNaP samples only up to this approximation error. The exact-map result identifies the posterior target of SNaP; the effect of the source choice on the learned one-step approximation is evaluated through the controlled comparison with isotropic sources in Table 5.

### 3.1 Training and sampling

A direct draw from  $q_\tau(\cdot | \mathbf{y})$  would require applying the matrix square root  $\mathbf{W}^{1/2}$ , which is generally impractical for large imaging operators. We instead use perturb-and-solve [34, 35]:

$$\mathbf{x}_0 = \boldsymbol{\epsilon} + \mathbf{A}^\top (\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I})^{-1} (\mathbf{y} - \mathbf{A} \boldsymbol{\epsilon} - \mathbf{n}'), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I}), \quad \mathbf{n}' \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}). \quad (13)$$

Here,  $\boldsymbol{\epsilon}$  and  $\mathbf{n}'$  are independent of each other and of  $(\mathbf{x}_1, \mathbf{n})$ . Lemma 1 shows that  $\mathbf{x}_0 | \mathbf{y} \sim q_\tau(\cdot | \mathbf{y})$  exactly and that  $\mathbf{x}_0$  and  $\mathbf{x}_1$  are conditionally independent given  $\mathbf{y}$ . Each source draw requires one solve with  $\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I}$ ; operator-specific implementations are given in App. C.2.

**Training.** The network parameterizes the conditional average velocity

$$\mathbf{u}_{r,t}^\theta = \mathbf{u}^\theta(\mathbf{z}_r, r, t | \mathbf{y}, \sigma_n).$$

Because  $\mathbf{y}$  and  $\sigma_n$  remain fixed along the path, the conditional MeanFlow identity gives the target

$$\mathbf{u}_{\text{tgt}} = \mathbf{v} + (t - r) \text{sg}[(\partial_{\mathbf{z}} \mathbf{u}_{r,t}^\theta) \mathbf{v} + \partial_r \mathbf{u}_{r,t}^\theta], \quad \mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0. \quad (14)$$

The derivative term is evaluated using one Jacobian–vector product, and  $\text{sg}[\cdot]$  prevents gradients from propagating through the target. We minimize

$$\mathcal{L}(\boldsymbol{\theta}) = \mathbb{E}[\text{sg}[\omega(\ell)] \ell], \quad \ell := \frac{1}{n} \|\mathbf{u}_{r,t}^\theta - \mathbf{u}_{\text{tgt}}\|_2^2, \quad (15)$$

using the adaptive weight  $\omega(\ell) = (\ell + c)^{-p}$  from [25]. Implementation details and Algorithm 1 are given in App. C.3.

**Sampling.** At test time, SNaP draws  $\mathbf{x}_0$  using equation 13 and returns  $\hat{\mathbf{x}}_1 = \mathbf{x}_0 + \mathbf{u}^\theta(\mathbf{x}_0, 0, 1 | \mathbf{y}, \sigma_n)$ . Each sample requires one source solve and one network evaluation. Redrawing  $(\boldsymbol{\epsilon}, \mathbf{n}')$  with  $\mathbf{y}$  fixed produces conditionally independent samples from the learned model.

## 4 Experiments

**Datasets.** We evaluate on CelebA (128×128), AFHQ-Cat (256×256), and multi-coil fastMRI brain AXT2 (320×320), averaging all metrics over 100 test measurements. Details are in App. C.1.

Table 1: CelebA, 100 test images. We report single draw ( $M=1$ ) and the average of  $M$  draws for SNaP. Note that a single SNaP draw is competitive in LPIPS at one network evaluation, while averaging draws attains the best SSIM on every task and the best PSNR on three of the four. **Best** and **second best** are shown in blue and red.

| Method           | NFE  | Deblurring      |                 |                    | Super-resolution |                 |                    | Random inpainting |                 |                    | Box inpainting  |                 |                    |
|------------------|------|-----------------|-----------------|--------------------|------------------|-----------------|--------------------|-------------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                  |      | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$  | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$   | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| Degraded         | –    | 27.26           | 0.838           | 0.214              | 11.67            | 0.182           | 0.859              | 11.98             | 0.198           | 1.070              | 22.33           | 0.754           | 0.217              |
| MSE regressor    | 1    | <b>35.74</b>    | 0.953           | 0.029              | <b>34.17</b>     | <b>0.945</b>    | 0.033              | 34.42             | 0.955           | 0.022              | –               | –               | –                  |
| PnP-GS           | 23   | 33.97           | 0.924           | 0.041              | 31.23            | 0.890           | 0.065              | 29.17             | 0.874           | 0.066              | –               | –               | –                  |
| DDRM             | 20   | 35.02           | 0.946           | 0.027              | 32.24            | 0.927           | <b>0.027</b>       | 32.40             | 0.942           | 0.029              | –               | –               | –                  |
| DiffPIR          | 100  | 34.83           | 0.938           | 0.027              | 31.87            | 0.892           | 0.030              | 32.45             | 0.924           | 0.024              | 30.48           | 0.918           | 0.028              |
| OT-ODE           | 180  | 32.96           | 0.920           | 0.029              | 31.34            | 0.903           | <b>0.027</b>       | 28.69             | 0.871           | 0.049              | 29.37           | 0.919           | 0.038              |
| PnP-Flow5        | 500  | 34.80           | 0.941           | 0.046              | 31.44            | 0.905           | 0.056              | 33.98             | 0.953           | 0.021              | 31.06           | 0.939           | 0.043              |
| Flower5-OT       | 500  | 35.65           | <b>0.954</b>    | 0.031              | 33.03            | 0.931           | 0.039              | 33.97             | 0.952           | 0.019              | 31.85           | 0.952           | 0.023              |
| DPS              | 1000 | 34.27           | 0.936           | <b>0.019</b>       | 32.23            | 0.917           | <b>0.023</b>       | 33.43             | 0.951           | <b>0.013</b>       | 30.26           | 0.942           | 0.022              |
| SNaP ( $M=1$ )   | 1    | 32.90           | 0.919           | <b>0.018</b>       | 31.27            | 0.903           | <b>0.023</b>       | 31.97             | 0.927           | 0.018              | 30.73           | 0.934           | 0.021              |
| SNaP ( $M=4$ )   | 4    | 34.98           | 0.947           | 0.022              | 33.26            | 0.934           | <b>0.023</b>       | 33.97             | 0.951           | <b>0.015</b>       | 32.78           | 0.954           | 0.018              |
| SNaP ( $M=16$ )  | 16   | 35.68           | <b>0.954</b>    | 0.028              | 33.95            | 0.943           | 0.029              | <b>34.64</b>      | <b>0.957</b>    | 0.018              | <b>33.43</b>    | <b>0.959</b>    | <b>0.020</b>       |
| SNaP ( $M=100$ ) | 100  | <b>35.91</b>    | 0.956           | 0.030              | <b>34.16</b>     | 0.946           | 0.031              | 34.84             | 0.959           | 0.019              | <b>33.69</b>    | <b>0.961</b>    | 0.022              |

Table 2: Distributional and calibration metrics for CelebA.  $\Delta_{\text{cal}}$  is the calibration ratio of App. C.4, for which 1 is ideal. Note that SNaP attains the best or second-best FID among the posterior-sampling baselines in one network evaluation. **Best** and **second best** are shown in blue and red.

| Method        | NFE  | Gaussian deblurring |                                       | Super-resolution |                                       | Box inpainting   |                                       |
|---------------|------|---------------------|---------------------------------------|------------------|---------------------------------------|------------------|---------------------------------------|
|               |      | FID $\downarrow$    | $\Delta_{\text{cal}} (\rightarrow 1)$ | FID $\downarrow$ | $\Delta_{\text{cal}} (\rightarrow 1)$ | FID $\downarrow$ | $\Delta_{\text{cal}} (\rightarrow 1)$ |
| MSE regressor | 1    | 22.06               | –                                     | 24.76            | –                                     | –                | –                                     |
| DDRM          | 20   | 20.26               | 0.31                                  | <b>18.94</b>     | 0.19                                  | –                | –                                     |
| Flower1-OT    | 100  | <b>17.86</b>        | 0.23                                  | 23.79            | 0.24                                  | 18.72            | <b>0.24</b>                           |
| DPS           | 1000 | 18.25               | <b>0.64</b>                           | 19.94            | <b>0.67</b>                           | <b>15.42</b>     | 2.15                                  |
| SNaP (ours)   | 1    | <b>15.80</b>        | 1.03                                  | 16.80            | 0.99                                  | <b>15.53</b>     | <b>1.08</b>                           |

**Forward operators.** We adopt the degradation settings of [20] and [21] for natural images, with  $\sigma_n = 0.05$  everywhere except random inpainting, where  $\sigma_n = 0.01$ . Deblurring uses a  $61 \times 61$  Gaussian kernel with  $\sigma_b = 1.0$  on CelebA and 3.0 on AFHQ-Cat; super-resolution uses stride subsampling at  $2\times$  on CelebA and  $4\times$  on AFHQ-Cat; random inpainting masks 70% of pixels; box inpainting uses a centered  $40 \times 40$  mask on CelebA and  $80 \times 80$  on AFHQ-Cat. For MRI,  $\mathbf{A} = \mathbf{MFS}$  with sampling mask  $\mathbf{M}$ , Fourier transform  $\mathbf{F}$  and per-slice ESPIRiT sensitivities  $\mathbf{S}$ , at  $\times 4$  and  $\times 8$  Cartesian acceleration, with noise scaled to 20 and 30 dB input SNR. Closed forms of equation 13 for each operator are given in App. C.2.

**Baselines.** For natural images, we compare against a supervised MSE regressor trained with the SNaP backbone, the plug-and-play solvers PnP-GS [36] and PnP-Flow [20], and the iterative posterior samplers DDRM [13], DiffPIR [17], OT-ODE [15], Flower [21] and DPS [11]. Flower and PnP-Flow are run at the recommended setting, where the exact budget of every method is listed in App. C.3. For MRI, we compare against zero-filled, wavelet- $\ell_1$  and TV reconstructions, the MSE regressor, and the posterior samplers CSGM [37], DiffPIR, DPS, DAPS [19] and PnP-DM [38, 39]. Variational Flow Maps (VFM) [33] also place the measurement in the source, but learn it with an adapter rather than constructing it in closed form, and report results for latent-space flow models, whereas SNaP operates in pixel space; we compare the two source constructions in App. B.4. NullFlow [28] is not defined for  $\sigma_n > 0$  (Sec. 2), so we compare against it from  $\sigma_n = 0$  to 0.05 in App. B.5.

**Metrics.** We report PSNR, SSIM [40], LPIPS [41] and FID [42]. None of them says whether the spread of the draws is of the right size, so we also report the calibration ratio  $\Delta_{\text{cal}} = V/B$ , which divides the variance  $V$  of the draws by the squared error  $B$  of their mean. For an exact posterior sampler the two are equal, so  $\Delta_{\text{cal}} = 1$  is ideal: values below 1 indicate draws collapsed toward a point estimate, and values above 1 indicate over-dispersed draws. It is a second-moment summary and does not certify conditional coverage. We give the definition and estimators in App. C.4.

Table 3: fastMRI brain, multi-coil. PSNR and SSIM over 100 test slices, at  $\times 4$  and  $\times 8$  acceleration and two input SNRs. For SNaP, we report results for a single draw  $M = 1$  and average of  $M$  draws. Note that among posterior samplers SNaP is very competitive even using one NFE. **Best** and **second best** are shown in blue and red.

| Method            | NFE  | $\times 4$      |                 |                 |                 | $\times 8$      |                 |                 |                 |
|-------------------|------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
|                   |      | 20 dB           |                 | 30 dB           |                 | 20 dB           |                 | 30 dB           |                 |
|                   |      | PSNR $\uparrow$ | SSIM $\uparrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ |
| Zero-filled       | —    | 25.35           | 0.791           | 25.37           | 0.793           | 21.98           | 0.694           | 21.99           | 0.694           |
| Wavelet+ $\ell_1$ | —    | 26.57           | 0.672           | 27.75           | 0.768           | 23.02           | 0.551           | 23.81           | 0.643           |
| TV                | —    | 26.19           | 0.798           | 26.27           | 0.799           | 22.62           | 0.689           | 22.66           | 0.691           |
| MSE regressor     | 1    | 33.80           | 0.909           | 33.85           | 0.912           | 30.71           | 0.875           | 30.73           | 0.878           |
| DiffPIR           | 100  | 29.84           | 0.709           | 30.01           | 0.719           | 26.14           | 0.642           | 26.31           | 0.655           |
| CSGM              | 1000 | 27.74           | 0.734           | 32.39           | 0.802           | 27.74           | 0.733           | 28.33           | 0.754           |
| DPS               | 1000 | 28.81           | 0.646           | 29.83           | 0.669           | 24.89           | 0.575           | 25.17           | 0.574           |
| DAPS              | 1000 | 31.55           | 0.818           | 32.44           | 0.782           | 27.31           | 0.749           | 28.01           | 0.714           |
| PnP-DM            | 1000 | 32.16           | 0.778           | 33.39           | 0.786           | 27.18           | 0.701           | 29.00           | 0.712           |
| SNAP ( $M=1$ )    | 1    | 32.06           | 0.889           | 32.89           | 0.899           | 28.10           | 0.818           | 29.14           | 0.845           |
| SNAP ( $M=4$ )    | 4    | 32.68           | 0.904           | 33.80           | 0.919           | 29.06           | 0.843           | 29.88           | 0.864           |
| SNAP ( $M=16$ )   | 16   | 32.85           | 0.908           | 34.05           | 0.924           | 29.34           | 0.850           | 30.09           | 0.869           |

Table 4: **Computational cost.** NFE and 1-batch wall-clock time per posterior sample. **Best** in each row is in bold.

| <i>CelebA, Gaussian deblurring</i>                     |               |              |              |              |              |
|--------------------------------------------------------|---------------|--------------|--------------|--------------|--------------|
|                                                        | DPS           | OT-ODE       | Flower       | DDRM         | SNaP         |
| NFE $\downarrow$                                       | 1000          | 180          | 100          | 20           | <b>1</b>     |
| Time / sample (s) $\downarrow$                         | 27.014        | 6.575        | 2.614        | 0.360        | <b>0.012</b> |
| Slowdown vs. ours $\downarrow$                         | 2251 $\times$ | 548 $\times$ | 218 $\times$ | 30 $\times$  | —            |
| <i>fastMRI brain, multi-coil <math>\times 8</math></i> |               |              |              |              |              |
|                                                        | CSGM          | DAPS         | PnP-DM       | DPS          | SNaP         |
| NFE $\downarrow$                                       | 1000          | 1000         | 1000         | 1000         | <b>1</b>     |
| Time / sample (s) $\downarrow$                         | 127.78        | 96.28        | 92.02        | 76.07        | <b>0.227</b> |
| Slowdown vs. ours $\downarrow$                         | 563 $\times$  | 424 $\times$ | 405 $\times$ | 335 $\times$ | —            |

#### 4.1 Main Results

On CelebA (Tab. 1), a single SNaP draw attains the best LPIPS on Gaussian deblurring, and averaging draws attains the best SSIM on all four operators and the best PSNR on three of the four. Results on AFHQ-Cat are reported in App. B.1. SNaP also attains the best FID on two of the three CelebA operators and the calibration ratio closest to 1 on all three (Tab. 2), so the spread of its draws matches the error they make. On multi-coil fastMRI (Tab. 3) it is competitive with posterior sampler at both accelerations and both noise levels with only **one** NFE. Figures 3 and 4 show the corresponding reconstructions on natural images and on brain MRI.

#### 4.2 Computational cost

Each SNaP sample costs one network evaluation and one source draw by equation 13, whose solve with  $(\mathbf{A}\mathbf{A}^\top + \lambda\mathbf{I})^{-1}$  is diagonal (or Fourier diagonal) for the image operators and reduces to a short conjugate-gradient solve for multi-coil MRI (App. C.2). As reported in Tab. 4, this makes SNaP 30 to 2250 $\times$  faster per sample than the iterative samplers on CelebA deblurring and 335 to 563 $\times$  faster on  $\times 8$  MRI.

#### 4.3 Ablations

**Source.** Table 5 isolates the effect of the source by comparing our proposed source to two isotropic sources: the first is the usual flow matching source  $\mathcal{N}(\mathbf{0}, \mathbf{I})$  and the other is  $\mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$  to control for the influence of  $\tau$ . We hold the other design choices constant and  $\mathbf{y}$  and  $\sigma_n$  are retained in the conditioning. Under the isotropic sources, the calibration ratio falls to 0.00: the draws are nearly identical, yielding a *point estimator* with higher single-draw PSNR and worse LPIPS. The measurement-adapted source keeps the draws dispersed, and its single-draw advantage is perceptual






Figure 3: Visual comparison of samples for three imaging tasks, CelebA super-resolution and AFHQ-Cat box inpainting and deblurring, with PSNR and LPIPS for the displayed image. Note that SNaP requires one NFE, while DiffPIR requires 100 and DPS and Flower require 1000. SNaP attains the best LPIPS on all three.




Figure 4: fastMRI brain, multi-coil,  $\times 8$  acceleration at 20 dB input SNR. Reconstructions with PSNR and SSIM, and error magnitude and a zoomed in region is displayed. Note that SNaP achieves competitive reconstruction using 1 NFE, while baselines use 1000.

rather than distortion-based. Figures 12 and 13 in App. B.3 show this collapse directly—the draws and their per-pixel spread—and trace both metrics against the averaging budget  $M$ .

**Additional results.** App. B.1 reports the full AFHQ-Cat tables and App. B.2 shows further reconstructions on natural images and brain MRI. App. B.5 compares SNaP with NullFlow across a range of measurement noise levels, and App. B.4 compares the two source constructions on the two-dimensional benchmark of [33]. App. B.7 verifies the method against a closed-form posterior, and App. B reports the remaining ablations, on the working-prior scale  $\tau$ , the  $(r, t)$  sampling ratio and few-step composition.

## 5 Related work

**Iterative posterior samplers.** Most generative solvers enforce consistency with the measurement repeatedly along the sampling trajectory, either by correcting the generative dynamics with a data-

Table 5: **Source distribution.** CelebA, Gaussian deblurring at  $\sigma_n = 0.05$ ; only the source differs across rows.  $M$  is the number of averaged draws, and  $\Delta_{\text{cal}}$  is the calibration ratio (App. C.4), for which 1 is ideal. An isotropic source collapses onto the posterior mean, where draws are near-identical and averaging brings no gain.

| Source                                                | $M = 1$         |                    | $M = 100$       |                    | $\Delta_{\text{cal}} (\rightarrow 1)$ |
|-------------------------------------------------------|-----------------|--------------------|-----------------|--------------------|---------------------------------------|
|                                                       | PSNR $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | LPIPS $\downarrow$ |                                       |
| $\mathcal{N}(\mathbf{0}, \mathbf{I})$                 | 35.85           | 0.030              | 35.86           | 0.030              | 0.00                                  |
| $\mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$          | 35.83           | 0.030              | 35.84           | 0.030              | 0.00                                  |
| $\mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$ (ours) | 32.90           | <b>0.018</b>       | <b>35.91</b>    | 0.030              | <b>1.03</b>                           |

fidelity gradient or a pseudoinverse update [11, 12, 13, 14, 15, 16, 43], or by alternating a generative step with an explicit data-consistency step [17, 18, 19, 20, 21]. These methods reach high sample quality, but a single posterior sample costs tens to thousands of network evaluations.

**One-step generative models.** Consistency models [22, 23], shortcut models [24] and MeanFlow [25] eliminate the trajectory and map a source sample to the data in a single network evaluation. Using them for inverse problems requires the measurement to enter through the network conditioning or through the source, since no intermediate state is available at which to enforce data consistency. For real-world super-resolution, MFSR [44] distills a text-to-image flow model into a MeanFlow in which the low-resolution image enters only through conditioning, without an explicit forward model or posterior sampling.

**Measurement-dependent sources.** Bridge models transport from a measurement-dependent source by starting from the degraded observation itself [45, 46, 47], but that source carries no per-direction covariance. NullFlow [28] restricts a MeanFlow to  $\text{null}(\mathbf{A})$  and VFM [33] learns a measurement-dependent adapter jointly with the flow map, which covers a broader class of problems but yields a source that is trained. MeanFlow has also been applied to PDE-based Bayesian inverse problems with a prior-aligned base measure that does not depend on the measurement [48].

## 6 Conclusion

We introduced SNaP, a one-step posterior sampler for linear inverse problems with Gaussian noise. Instead of correcting toward the measurement at every step, SNaP builds it into the transport. The transport starts from a distribution centered at the linear estimate that minimizes the expected transport distance, with a spread that matches the remaining uncertainty. We proved that the exact one-step map carries this distribution to the target posterior. On natural images and multi-coil MRI, SNaP draws a posterior sample in a single network evaluation, 30 to  $2250\times$  faster than iterative samplers. On CelebA, averaging its draws gives the best SSIM on all four operators and the best PSNR on three. Its draws also reach the best FID on two of the three operators tested and the calibration ratio closest to 1 on all three.

**Limitations.** The construction is tied to linear operators, since the working posterior is available in closed form only in that setting, and nonlinear forward models require further work. Sampling the source requires one solve with  $\mathbf{A}\mathbf{A}^\top + \lambda\mathbf{I}$ , which is inexpensive for the operators considered here but may be demanding for operators without exploitable structure. Finally, SNaP operates in pixel space, and extending the source construction to latent-space flow models remains open.

## Acknowledgments

This work was supported in part by the National Science Foundation under CAREER award CCF-2625643 and award CCF-2622128.

## References

- [1] Mario Bertero, Patrizia Boccacci, and Christine De Mol, *Introduction to inverse problems in imaging*, CRC Press, 2021.

- [2] Michael T McCann, Kyong Hwan Jin, and Michael Unser, “Convolutional Neural Networks for Inverse Problems in Imaging: A Review,” *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 85–95, 2017.
- [3] Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett, “Deep Learning Techniques for Inverse Problems in Imaging,” *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 39–56, 2020.
- [4] Leonid I Rudin, Stanley Osher, and Emad Fatemi, “Nonlinear t=Total Variation Based Noise Removal Algorithms,” *Physica D*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [5] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising Diffusion Probabilistic Models,” in *Proc. NeurIPS*, 2020.
- [6] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, “Score-Based Generative Modeling through Stochastic Differential Equations,” in *Proc. ICLR*, 2021.
- [7] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le, “Flow Matching for Generative Modeling,” in *Proc. ICLR*, 2023.
- [8] Xingchao Liu, Chengyue Gong, and Qiang Liu, “Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow,” in *Proc. ICLR*, 2023.
- [9] Michael Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden, “Stochastic Interpolants: A Unifying Framework for Flows and Diffusions,” *J. Mach. Learn. Res.*, vol. 26, no. 209, pp. 1–80, 2025.
- [10] Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, and Mauricio Delbracio, “A Survey on Diffusion Models for Inverse Problems,” *arXiv:2410.00083*, 2024.
- [11] Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye, “Diffusion Posterior Sampling for General Noisy Inverse Problems,” in *Proc. ICLR*, 2023.
- [12] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz, “Pseudoinverse-Guided Diffusion Models for Inverse Problems,” in *Proc. ICLR*, 2023.
- [13] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song, “Denoising Diffusion Restoration Models,” in *Proc. NeurIPS*, 2022.
- [14] Yinhuai Wang, Jiwen Yu, and Jian Zhang, “Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model,” in *Proc. ICLR*, 2023.
- [15] Ashwini Pople, Matthew J. Muckley, Ricky T. Q. Chen, and Brian Karrer, “Training-free Linear Image Inverses via Flows,” *TMLR*, 2024.
- [16] Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman, “D-Flow: Differentiating through Flows for Controlled Generation,” in *Proc. ICML*, 2024.
- [17] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool, “Denoising Diffusion Models for Plug-and-Play Image Restoration,” in *CVPR Workshops*, 2023, pp. 1219–1229.
- [18] Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye, “Decomposed Diffusion Sampler for Accelerating Large-Scale Inverse Problems,” in *Proc. ICLR*, 2024.
- [19] Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song, “Improving Diffusion Inverse Problem Solving with Decoupled Noise Annealing,” in *Proc. CVPR*, 2025.
- [20] Ségolène Tiffany Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl, “PnP-Flow: Plug-and-Play Image Restoration with Flow Matching,” in *Proc. ICLR*, 2025.
- [21] Mehrsa Pourya, Bassam El Rawas, and Michael Unser, “Flower: A Flow-Matching Solver for Inverse Problems,” in *Proc. ICLR*, 2026.
- [22] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever, “Consistency Models,” in *Proc. ICML*, 2023.
- [23] Nicholas Boffi, Michael Albergo, and Eric Vanden-Eijnden, “How to Build a Consistency Model: Learning Flow Maps via Self-Distillation,” in *Proc. NeurIPS*, 2025.

- [24] Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel, “One Step Diffusion via Shortcut Models,” in *Proc. ICLR*, 2025.
- [25] Zhengyang Geng, Mingyang Deng, Xingjian Bai, Zico Kolter, and Kaiming He, “Mean Flows for One-step Generative Modeling,” in *Proc. NeurIPS*, 2025.
- [26] Xiao He, Yang Li, Peizhen Zhang, Songtao Liu, Zhao Zhong, and Nannan Wang, “Stabilizing, Scaling & Enhancing MeanFlow for Large-scale Diffusion Distillation,” *arXiv preprint arXiv:2605.17834*, 2026.
- [27] Yexiong Lin, Jia Shi, Shanshan Ye, Wanyu Wang, Yu Yao, and Tongliang Liu, “SubFlow: Sub-mode Conditioned Flow Matching for Diverse One-Step Generation,” *arXiv 2604.12273*, 2026.
- [28] Xiao Shi, Edward P Chandler, Chicago Y Park, Shirin Shoushtari, and Ulugbek S Kamilov, “NullFlow: One-Step Generative Reconstruction,” *arXiv:2606.22696*, 2026.
- [29] Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky T. Q. Chen, “Multisample Flow Matching: Straightening Flows with Minibatch Couplings,” in *Proc. ICML*, 2023.
- [30] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio, “Improving and generalizing flow-based generative models with minibatch optimal transport,” *TMLR*, pp. 1–34, 2024.
- [31] Michael Samuel Albergo, Mark Goldstein, Nicholas Matthew Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden, “Stochastic Interpolants with Data-Dependent Couplings,” in *Proc. ICML*, 2024.
- [32] Sang gil Lee, Heeseung Kim, Chaehun Shin, Xu Tan, Chang Liu, Qi Meng, Tao Qin, Wei Chen, Sungroh Yoon, and Tie-Yan Liu, “PriorGrad: Improving Conditional Denoising Diffusion Models with Data-Dependent Adaptive Prior,” in *Proc. ICLR*, 2022.
- [33] Abbas Mammadov, So Takao, Bohan Chen, Ricardo Baptista, Morteza Mardani, Yee Whye Teh, and Julius Berner, “Variational Flow Maps: Make Some Noise for One-Step Conditional Generation,” in *Proc. ICML*, 2026.
- [34] George Papandreou and Alan L Yuille, “Gaussian Sampling by Local Perturbations,” in *Proc. NeurIPS*, 2010.
- [35] Johnathan M Bardsley, Antti Solonen, Heikki Haario, and Marko Laine, “Randomize-then-optimize: A method for sampling from posterior distributions in nonlinear inverse problems,” *SIAM J. Sci. Comput.*, vol. 36, no. 4, pp. A1895–A1910, 2014.
- [36] Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis, “Gradient Step Denoiser for Convergent Plug-and-Play,” in *Proc. ICLR*, 2022.
- [37] Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir, “Robust Compressed Sensing MRI with Deep Generative Priors,” in *Proc. NeurIPS*, 2021.
- [38] Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine Bouman, “Principled Probabilistic Imaging using Diffusion Models as Plug-and-Play Priors,” in *Proc. NeurIPS*, 2024.
- [39] Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy Feng, Caifeng Zou, Yu Sun, Nikola Borislavov Kovachki, Zachary E. Ross, Katherine Bouman, and Yisong Yue, “InverseBench: Benchmarking Plug-and-Play Diffusion Priors for Inverse Problems in Physical Sciences,” in *Proc. ICLR*, 2025.
- [40] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli, “Image Quality Assessment: From Error Visibility to Structural Similarity,” *IEEE Trans. Image Proc.*, vol. 13, no. 4, pp. 600–612, 2004.
- [41] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang, “The Unreasonable Effectiveness of Deep Features as a Perceptual Metric,” in *Proc. CVPR*, 2018.
- [42] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter, “Gans Trained by a Two Time-scale Update Rule Converge to a Local Nash Equilibrium,” in *Proc. NeurIPS*, 2017.
- [43] Jeongsol Kim, Bryan Sangwoo Kim, and Jong Chul Ye, “Flowdps: Flow-driven Posterior Sampling for Inverse Problems,” in *Proc. ICCV*, 2025.

- [44] Ruiqing Wang, Kai Zhang, Yuanzhi Zhu, Hanshu Yan, Shilin Lu, and Jian Yang, “MFSR: MeanFlow Distillation for One Step Real-World Image Super Resolution,” *arXiv:2603.20690*, 2026.
- [45] Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar, “I2SB: Image-to-image Schrödinger Bridge,” in *Proc. ICML*, 2023.
- [46] Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon, “Denoising Diffusion Bridge Models,” in *Proc. ICLR*, 2024.
- [47] Mauricio Delbracio and Peyman Milanfar, “Inversion by Direct Iteration: An Alternative to Denoising Diffusion for Image Restoration,” *TMLR*, 2023.
- [48] Zhiqi Li, Yuchen Sun, Greg Turk, and Bo Zhu, “Functional Mean Flow in Hilbert Space,” in *Proc. CVPR*, 2026.

## A Proofs

### A.1 Proof of Proposition 1

*Proof.* Let  $x_0 = Ky + \xi$  as in equation 6.

Since  $x_1 - x_0 = (x_1 - Ky) - \xi$  and  $\xi$  has mean 0 and is independent of  $x_1 - Ky$ , for every unit vector  $e$ ,

$$\mathbb{E}[(e^\top (x_1 - x_0))^2] = e^\top R(K) e + e^\top \Sigma e, \quad \mathbb{E}[\|x_1 - x_0\|_2^2] = \text{tr } R(K) + \text{tr } \Sigma.$$

For fixed  $\Sigma$ , only  $R(K)$  depends on  $K$ .

Let  $C_{xy} := \mathbb{E}[x_1 y^\top] = CA^\top$  and  $C_y := \mathbb{E}[yy^\top] = ACA^\top + \sigma_n^2 I$ , which is positive definite since  $\sigma_n > 0$ , so that  $K_C = C_{xy} C_y^{-1}$ . Expanding  $R(K) = C - KC_{xy}^\top - C_{xy} K^\top + KC_y K^\top$  and using  $K_C C_y = C_{xy}$ ,

$$R(K) = R(K_C) + (K - K_C) C_y (K - K_C)^\top, \quad R(K_C) = C - C_{xy} C_y^{-1} C_{xy}^\top = C - K_C A C. \quad (16)$$

By equation 16, for every unit vector  $e$ ,

$$e^\top R(K) e = e^\top R(K_C) e + \|C_y^{1/2} (K - K_C)^\top e\|_2^2 \geq e^\top R(K_C) e,$$

so  $K_C$  minimizes the displacement along every direction. Taking the trace,

$$\text{tr } R(K) = \text{tr } R(K_C) + \|C_y^{1/2} (K - K_C)^\top\|_F^2,$$

and since  $C_y^{1/2}$  is invertible, the second term vanishes if and only if  $K = K_C$ . The total displacement is therefore minimized if and only if  $K = K_C$ , for every  $\Sigma$ .

*Isotropic case.* For  $C = \tau^2 I$  and  $\lambda = \sigma_n^2 / \tau^2$ , dividing by  $\tau^2$  gives  $K_C = A^\top (AA^\top + \lambda I)^{-1}$ , so  $K_C y = m_y$ , and  $R(K_C) = \tau^2 (I - A^\top (AA^\top + \lambda I)^{-1} A) = \tau^2 W$ .  $\square$

**Remarks.** (i) The proof uses only the first two moments of  $x_1$ ,  $n$  and  $\xi$ ; neither the images, the noise, nor the perturbation need to be Gaussian. (ii) If  $x_1 \sim \mathcal{N}(0, C)$ , Gaussian conditioning gives  $x_1 | y \sim \mathcal{N}(K_C y, \Sigma_C)$ , so the source with center  $K_C y$  and Gaussian perturbation of covariance  $\Sigma_C$  is the exact posterior; for  $C = \tau^2 I$  it is the working posterior equation 9. (iii) For  $C = \tau^2 I$ , as  $\sigma_n \rightarrow 0$  with  $\tau$  fixed,  $\lambda \rightarrow 0$ , and the standard limit of Tikhonov regularization gives  $K_C \rightarrow A^\dagger$  and  $W \rightarrow I - A^\dagger A = P$ , so the source tends to  $\mathcal{N}(A^\dagger y, \tau^2 P)$ .

### A.2 Proof of Proposition 2

*Proof.* Let  $p_t(\cdot | y)$  denote the distribution of  $z_t = (1 - t)x_0 + tx_1$ , so that  $p_0(\cdot | y) = q_\tau(\cdot | y)$  and  $p_1(\cdot | y) = p(\cdot | y)$ . The marginal velocity

$$v_t(z, y) = \mathbb{E}[x_1 - x_0 | z_t = z, y] \quad (17)$$

generates  $\{p_t(\cdot | y)\}_{t \in [0,1]}$  through the continuity equation [7, 9].

Since  $\sigma_n > 0$ , we have  $\lambda > 0$  and  $w_i \in (0, 1]$  for all  $i$ , so  $W$  is positive definite. For  $t < 1$ , the interpolant  $z_t$  is therefore the sum of  $tx_1$  and an independent Gaussian with nondegenerate covariance  $(1 - t)^2 \tau^2 W$ . Its density  $p_t(\cdot | y)$  is thus smooth and strictly positive. Since  $x_1$  has finite second moments, standard arguments show that  $v_t(\cdot, y)$  is locally Lipschitz, uniformly for  $t \in [0, T]$  and every  $T < 1$  [9]. The ODE  $\frac{d}{dt} z_t = v_t(z_t, y)$  then has a unique solution on  $[0, 1]$ , and its flow map satisfies  $(\psi_{0 \rightarrow t})_\# q_\tau(\cdot | y) = p_t(\cdot | y)$  for every  $t < 1$ .

Since  $\mathbb{E}[\|z_t - x_1\|_2^2] = (1 - t)^2 \mathbb{E}[\|x_1 - x_0\|_2^2] \rightarrow 0$ , the distribution  $p_t(\cdot | y)$  converges to  $p(\cdot | y)$  as  $t \rightarrow 1$ . By the terminal-limit assumption of Proposition 2,  $\psi_{0 \rightarrow 1}(x) := \lim_{t \rightarrow 1} \psi_{0 \rightarrow t}(x)$  exists for  $q_\tau$ -almost every  $x$ . Almost-sure convergence implies convergence in distribution, so

$$(\psi_{0 \rightarrow 1})_\# q_\tau(\cdot | y) = p(\cdot | y). \quad (18)$$

Integrating the ODE from 0 to  $t$  and letting  $t \rightarrow 1$  gives

$$\psi_{0 \rightarrow 1}(x_0) - x_0 = \int_0^1 v_s(\psi_{0 \rightarrow s}(x_0), y) ds = u(x_0, 0, 1 | y), \quad (19)$$



where the last equality is the average velocity in equation 3. Hence  $\psi_{0 \rightarrow 1} = T_{\mathbf{y}}$ , and equation 18 proves equation 12. The target  $p(\cdot | \mathbf{y})$  does not depend on  $\tau$ , so the conclusion holds for every  $\tau > 0$ .  $\square$

### A.3 Perturb-and-solve sampling

Direct sampling would require a square root of  $\mathbf{W}$ . The perturb-and-solve sampler avoids this square root while preserving the same Gaussian law.

**Lemma 1.** *Let  $\epsilon \sim \mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$  and  $\mathbf{n}' \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$  be independent of each other and of  $(\mathbf{x}_1, \mathbf{n})$ , and let  $\mathbf{x}_0$  be given by equation 13. Then  $\mathbf{x}_0 | \mathbf{y} \sim \mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$  for every  $\mathbf{A} \in \mathbb{R}^{m \times n}$  and every  $\sigma_n > 0$ .*

*Proof.* Write  $\mathbf{K} := \mathbf{A}^\top (\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I})^{-1}$ , so that  $\mathbf{m}_y = \mathbf{K} \mathbf{y}$  and  $\mathbf{W} = \mathbf{I} - \mathbf{K} \mathbf{A}$  by Proposition 1. Substituting these into equation 13 gives

$$\mathbf{x}_0 = \mathbf{m}_y + \mathbf{W} \epsilon - \mathbf{K} \mathbf{n}'. \quad (20)$$

Given  $\mathbf{y}$ , this is an affine function of two independent Gaussian vectors, so  $\mathbf{x}_0$  is Gaussian with mean  $\mathbf{m}_y$  and covariance  $\tau^2 \mathbf{W}^2 + \sigma_n^2 \mathbf{K} \mathbf{K}^\top$ . Both  $\mathbf{W}$  and  $\mathbf{K} \mathbf{K}^\top$  are diagonal in the right singular basis of  $\mathbf{A}$ , with entries  $w_i$  and  $s_i^2 / (s_i^2 + \lambda)^2$ . Using  $\sigma_n^2 = \tau^2 \lambda$ , the  $i$ -th diagonal entry of the covariance is

$$\tau^2 w_i^2 + \tau^2 \lambda \frac{s_i^2}{(s_i^2 + \lambda)^2} = \tau^2 \frac{\lambda^2 + \lambda s_i^2}{(s_i^2 + \lambda)^2} = \tau^2 \frac{\lambda}{s_i^2 + \lambda} = \tau^2 w_i,$$

so the covariance equals  $\tau^2 \mathbf{W}$ .  $\square$

Each source draw requires applications of  $\mathbf{A}$  and  $\mathbf{A}^\top$  and one solve with  $\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I}$ . No square root of  $\mathbf{W}$  is required. Operator-specific implementations are given in App. C.2.

## B Additional Results

### B.1 Full results on AFHQ-Cat

Table 6: AFHQ-Cat, 100 test images. We report single draw ( $M=1$ ) and the average of  $M$  draws for SNaP. Note that a single SNaP draw is competitive in LPIPS at one network evaluation, and averaging a few draws recovers distortion at a cost still far below the iterative samplers. On AFHQ-Cat the number of time steps used by PnP-Flow and Flower depends on the operator, so their NFE does too; the per-operator budgets are listed in App. C.3. The MSE regressor is omitted on box inpainting, where its single output is the blurred conditional mean of the masked region and is not comparable with the other methods. **Best** and **second best** are shown in blue and red.

| Method           | NFE      | Deblurring      |                 |                    | Super-resolution |                 |                    | Random inpainting |                 |                    | Box inpainting  |                 |                    |
|------------------|----------|-----------------|-----------------|--------------------|------------------|-----------------|--------------------|-------------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                  |          | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$  | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$   | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| Degraded         | –        | 24.39           | 0.532           | 0.543              | 11.98            | 0.212           | 0.900              | 13.25             | 0.214           | 1.091              | 21.57           | 0.736           | 0.214              |
| MSE regressor    | 1        | 29.82           | 0.806           | 0.289              | 29.16            | 0.817           | 0.201              | 33.23             | 0.915           | 0.066              | –               | –               | –                  |
| PnP-GS           | 23       | 28.39           | 0.787           | 0.387              | 24.44            | 0.639           | 0.411              | 29.89             | 0.841           | 0.123              | –               | –               | –                  |
| PnP-Flow5        | 500–2500 | 29.01           | 0.785           | 0.312              | 28.01            | 0.791           | 0.167              | 34.47             | 0.933           | 0.044              | 27.14           | 0.900           | 0.127              |
| DDRM             | 20       | 29.01           | 0.784           | 0.194              | 27.09            | 0.781           | 0.186              | 32.20             | 0.907           | 0.064              | –               | –               | –                  |
| DiffPIR          | 100      | 28.89           | 0.773           | 0.187              | 23.16            | 0.635           | 0.291              | 31.76             | 0.881           | 0.053              | 27.71           | 0.880           | 0.059              |
| OT-ODE           | 180      | 27.82           | 0.735           | 0.126              | 26.71            | 0.737           | 0.108              | 29.98             | 0.849           | 0.086              | 24.47           | 0.873           | 0.095              |
| Flower5-OT       | 500–2500 | 29.73           | 0.801           | 0.264              | 27.09            | 0.767           | 0.272              | 34.41             | 0.932           | 0.047              | 27.32           | 0.922           | 0.067              |
| DPS              | 1000     | 25.90           | 0.661           | 0.148              | 25.98            | 0.694           | 0.144              | 32.06             | 0.896           | 0.036              | 26.78           | 0.900           | 0.054              |
| SNaP ( $M=1$ )   | 1        | 26.25           | 0.674           | 0.159              | 26.07            | 0.703           | 0.130              | 30.48             | 0.853           | 0.066              | 26.46           | 0.891           | 0.054              |
| SNaP ( $M=4$ )   | 4        | 28.33           | 0.750           | 0.157              | 28.07            | 0.777           | 0.120              | 32.38             | 0.896           | 0.052              | 28.49           | 0.914           | 0.047              |
| SNaP ( $M=16$ )  | 16       | 29.09           | 0.779           | 0.231              | 28.77            | 0.802           | 0.157              | 33.04             | 0.909           | 0.056              | 29.32           | 0.922           | 0.056              |
| SNaP ( $M=100$ ) | 100      | 29.32           | 0.788           | 0.320              | 28.99            | 0.810           | 0.187              | 33.25             | 0.913           | 0.058              | 29.53           | 0.925           | 0.069              |

Table 6 reports the same protocol as Tab. 1 on AFHQ-Cat.




Figure 5: Box inpainting on CelebA.

## B.2 Additional visual results

Figure 5 shows box inpainting on CelebA.

Figure 6 shows box inpainting on AFHQ-Cat.

Figure 7 shows multi-coil fastMRI at  $\times 8$  acceleration and 30 dB input SNR.

Figure 8 shows Gaussian deblurring on CelebA.

Figure 9 shows several SNaP draws for box inpainting.

Figure 10 shows further SNaP draws for box inpainting.

Figure 11 shows further SNaP draws for box inpainting.

## B.3 Ablations

**Noise misspecification.** Since  $\sigma_n$  enters the source in closed form, a wrong value at test time changes how the source allocates uncertainty rather than breaking the method: by equation 10, over-estimating  $\sigma_n$  loosens the source toward the working prior, while under-estimating it shrinks the source toward the least-squares anchor. Table 7 tests the first case: the model and its source keep  $\sigma_n = 0.05$ , while the test measurements carry noise between 0.01 and 0.05, an over-estimate of up to a factor of 5. Reconstruction degrades gracefully. PSNR still rises as the measurements get cleaner, from 32.90 to 34.29 dB, and LPIPS stays between 0.017 and 0.020, although the gain over the input shrinks from 5.09 to 4.33 dB as the mismatch grows. Over this range the method therefore does not require  $\sigma_n$  to be known exactly; under-estimation is not tested here.




Figure 6: Box inpainting on AFHQ-Cat.

Table 7: CelebA Gaussian deblurring trained at  $\sigma_n = 0.05$  and evaluated on  $y$  with different measurement noise level.  $\dagger$  marks the matched case, where the test noise equals the training noise.

| Measurement noise $\sigma_n$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | input PSNR | Gain  |
|------------------------------|-----------------|-----------------|--------------------|------------|-------|
| 0.01                         | 34.29           | 0.940           | 0.020              | 29.96      | +4.33 |
| 0.02                         | 34.15           | 0.939           | 0.019              | 29.62      | +4.53 |
| 0.03                         | 33.89           | 0.936           | 0.018              | 29.12      | +4.77 |
| 0.04                         | 33.50           | 0.930           | 0.017              | 28.52      | +4.98 |
| 0.05 $^\dagger$              | 32.90           | 0.919           | 0.018              | 27.81      | +5.09 |

**Working-prior scale  $\tau$ .** Proposition 2 states that  $\tau$  shapes the source and the transport but not the target, so it is a design parameter rather than a modelling choice. Table 8 reports CelebA Gaussian deblurring for  $\tau \in \{0.15, 0.5, 1.0\}$ . Distortion is indistinguishable between  $\tau = 0.15$  and  $\tau = 0.5$  at every  $M$  and falls by about 0.4 dB at  $\tau = 1.0$ , while LPIPS is best at  $\tau = 0.5$ . The relative spread of the draws is 5.49%, 5.48% and 5.72% for the three settings, essentially unchanged, as the invariance predicts. The method therefore does not require  $\tau$  to be tuned precisely, and we use the per-operator values listed in App. C.3.

**Sampling ratio  $(r, t)$ .** The fraction  $p_{\text{ratio}} := \Pr[r=t]$  of training samples drawn with  $r = t$  controls how often the objective reduces to plain conditional flow matching. Table 9 reports CelebA Gaussian deblurring at  $\tau = 0.15$  for  $p_{\text{ratio}} \in \{0.25, 0.5, 0.75\}$ . Setting  $p_{\text{ratio}} = 0.25$  is worse at every  $M$ , by about 0.3 dB PSNR at a single draw, while 0.5 and 0.75 are indistinguishable within the precision reported. We therefore use  $p_{\text{ratio}} = 0.5$  throughout. The trend in  $M$  is the same for all three settings, with PSNR and SSIM improving as draws are averaged and LPIPS degrading, as the average moves toward the posterior mean.






Figure 7: fastMRI brain, multi-coil,  $\times 8$  acceleration at 30 dB input SNR.

**Few-step composition.** Composing  $\mathbf{u}_{r,t}$  over a uniform partition of  $[0, 1]$  gives  $k$ -step sampling with no retraining, and Table 10 reports  $k \in \{1, 2, 4, 8\}$  on CelebA Gaussian deblurring. More steps are not monotonically better. All metrics peak at  $k = 2$  and decline afterwards, and  $k = 8$  is worse than a single step for  $M \geq 16$ . The model is trained to jump from the source to the data, so discretizing the trajectory evaluates  $\mathbf{u}^\theta$  at intermediate  $(r, t)$  pairs at which its error compounds across steps rather than cancelling. The gain from  $k = 2$  is small, so we report  $k = 1$  throughout.

**Source distribution.** Table 5 isolates the source quantitatively; Figures 12 and 13 show the same comparison sample by sample and as a function of the averaging budget  $M$ . All three settings share the network, the conditioning, and the training schedule, differing only in the distribution the




Figure 8: Gaussian deblurring on CelebA.




Figure 9: Independent SNaP draws for box inpainting.

one-step map starts from: our measurement-adapted posterior source at  $\tau = 0.15$  versus isotropic Gaussian sources at  $\tau = 0.15$  and  $\tau = 1.0$ .

Figure 12 shows six draws per source for one measurement. Under the proposed posterior source the draws differ visibly in the hair, the skin texture and the region occluded by the sunglasses, and the standard deviation over draws concentrates on exactly the edges and high-frequency structures that the blur has suppressed, leaving the smooth background, which the measurement already determines, near zero. Under either isotropic Gaussian source the six draws are visually indistinguishable





Figure 10: Further independent SNaP draws for box inpainting.




Figure 11: Further independent SNaP draws for box inpainting.

and the standard-deviation map is essentially black at the same  $\times 20$  gain: the map has become a point estimator that happens to take a random input. Raising the Gaussian scale from  $\tau = 0.15$  to  $\tau = 1.0$  does not restore the spread, so the shape of the source is what dictates the spread rather than its size. Figure 13 makes the consequence quantitative.

#### B.4 Comparison with VFM

VFM [33] also places the measurement in the source, but learns it: an adapter  $q_\phi(z | y)$  is trained jointly with an unconditional flow map. Its results are reported in the latent space of a pretrained autoencoder, whereas SNaP operates in pixel space, so the published numbers are not directly comparable. We therefore compare the two constructions on the 2D benchmark of [33]. The prior is a  $4 \times 4$  checkerboard on  $[-2, 2]^2$  with 20k training samples, observed through  $y = a^\top x + \varepsilon$  with  $\varepsilon \sim \mathcal{N}(0, \sigma^2)$  and  $\sigma = 0.1$ . We use two forward operators that differ only in the direction of  $a$ , an aligned operator  $a = (1, 0)$ , which is the one used in the original paper, and a rotated operator  $a = (\cos \frac{\pi}{5}, \sin \frac{\pi}{5})$ . All methods share one backbone, a 6-layer width-512 SiLU MLP in  $x_1$ -prediction form with random-Fourier embeddings of  $r$ ,  $t$  and  $\sigma$ , and one budget, AdamW with learning rate  $2 \times 10^{-4}$ , batch 2048, 50k steps and EMA 0.999. The only thing that varies is where the measurement enters the generative process. The learned-adapter baseline follows VFM. The flow map never sees  $y$ , and all measurement information is carried by a diagonal Gaussian adapter  $q_\phi(z | y) = \mathcal{N}(\mu_\phi(y), \text{diag } \sigma_\phi^2(y))$  trained jointly with the flow map under the variational objective, with an observation term, a KL term and an EMA copy of the weights for stability. We also report a frozen- $\theta$  ablation in which the pretrained flow map is held fixed and only  $q_\phi$  is trained. SNaP



Table 8: **Working-prior scale  $\tau$ .** CelebA Gaussian deblurring, 100 test measurements, input PSNR 27.24 dB.  $M$  is the number of averaged draws

| $M$ | $\tau = 0.15$   |                 |                    | $\tau = 0.5$    |                 |                    | $\tau = 1.0$    |                 |                    |
|-----|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|     | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 1   | <b>32.941</b>   | <b>0.9198</b>   | 0.0190             | 32.922          | 0.9187          | <b>0.0186</b>      | 32.512          | 0.9060          | 0.0244             |
| 4   | <b>34.966</b>   | <b>0.9463</b>   | 0.0209             | 34.948          | 0.9457          | <b>0.0181</b>      | 34.588          | 0.9385          | 0.0197             |
| 16  | <b>35.688</b>   | <b>0.9537</b>   | 0.0275             | 35.663          | 0.9532          | <b>0.0234</b>      | 35.323          | 0.9475          | 0.0249             |
| 100 | <b>35.913</b>   | <b>0.9558</b>   | 0.0300             | 35.885          | 0.9553          | <b>0.0256</b>      | 35.552          | 0.9500          | 0.0273             |

Table 9:  **$(r, t)$  sampling ratio.** CelebA Gaussian deblurring,  $\tau = 0.15$ , 100 test measurements, input PSNR 27.24 dB.  $M$  is the number of averaged draws.

| $M$ | $p_{\text{ratio}} = 0.25$ |                 |                    | $p_{\text{ratio}} = 0.5$ |                 |                    | $p_{\text{ratio}} = 0.75$ |                 |                    |
|-----|---------------------------|-----------------|--------------------|--------------------------|-----------------|--------------------|---------------------------|-----------------|--------------------|
|     | PSNR $\uparrow$           | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$          | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$           | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 1   | 32.567                    | 0.9130          | 0.0216             | 32.941                   | 0.9198          | 0.0190             | 32.976                    | 0.9210          | 0.0189             |
| 4   | 34.679                    | 0.9429          | 0.0197             | 34.966                   | 0.9463          | 0.0209             | 34.976                    | 0.9465          | 0.0206             |
| 16  | 35.438                    | 0.9513          | 0.0266             | 35.688                   | 0.9537          | 0.0275             | 35.682                    | 0.9536          | 0.0268             |
| 100 | 35.676                    | 0.9537          | 0.0296             | 35.913                   | 0.9558          | 0.0300             | 35.902                    | 0.9557          | 0.0292             |

instead keeps the flow map conditioned on  $\mathbf{y}$  and draws its source in closed form by equation 13, with no adapter and no auxiliary losses.

The source covariance  $\mathbf{W} = \mathbf{I} - \mathbf{a}\mathbf{a}^\top / (\|\mathbf{a}\|_2^2 + \lambda)$  is diagonal in the basis  $\{\mathbf{a}, \mathbf{a}^\perp\}$ , and in the coordinate basis only when the two coincide. For  $\mathbf{a} = (1, 0)$  they coincide, so a diagonal adapter is exactly expressive and the aligned operator is a parity check. Rotating  $\mathbf{a}$  breaks the coincidence while holding the prior, the noise level, the architecture and the training budget fixed. The resulting gap can be quantified before any training: the KL divergence from  $\mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$  to its best diagonal approximation is 0 nats for the aligned operator at every  $\tau$ , and 1.576 nats for the rotated at  $\tau = 1$ .

On the aligned operator, where the diagonal adapter can represent the exact source, all three methods are close, with SNaP already the most accurate (Table 11). On the rotated operator the two constructions separate. The jointly trained adapter degrades by roughly a factor of seven and the frozen variant by more than an order of magnitude, while SNaP changes far less, since its source rotates with the operator by construction. The gaps are significant under a Welch test against the jointly trained adapter, with  $p = 0.0005$  on the aligned operator and  $p = 0.038$  on the rotated one. Figure 14 shows the mechanism, and the effect is also visible in the support of the samples, with a support accuracy of  $0.97 \pm 0.004$  for SNaP against  $0.81 \pm 0.03$  for the jointly trained adapter on the rotated operator.

## B.5 Comparison with NullFlow

NullFlow is defined only for noiseless measurements, so it cannot enter the benchmarks of Sec. 4. We instead sweep the test noise on CelebA box inpainting from  $\sigma_n = 0$ , where NullFlow is exact, to the  $\sigma_n = 0.05$  of our main experiments (Table 12). NullFlow is trained noiselessly with an LPIPS and feature loss, and SNaP at  $\sigma_n = 0.05$  with the standard MeanFlow loss, so LPIPS favors NullFlow.

Table 12 reports both methods on CelebA box inpainting, using a single draw and  $\lambda$  frozen at each model’s training value. At  $\sigma_n = 0$ , NullFlow is the stronger method, which is expected, since its source is exact in the noiseless limit and every state satisfies  $\mathbf{A}\mathbf{x} = \mathbf{y}$ . Its accuracy falls quickly as noise is added, and by  $\sigma_n = 0.05$  its SSIM has dropped from 0.977 to 0.855 and its LPIPS has more than quadrupled. This is the failure mode identified in Sec. 2, where the least-squares anchor carries  $\mathbf{A}^\dagger \mathbf{n}$  that a null-space flow cannot remove. SNaP changes little over the same range and is the better method at  $\sigma_n = 0.05$  on all three metrics. The two methods therefore occupy different regimes, with NullFlow suited to noiseless measurements and SNaP to noisy ones.

Table 10: **Trajectory steps  $k$** . CelebA Gaussian deblurring,  $\tau = 0.15$ , 100 test measurements, input PSNR 27.24 dB.  $M$  is the number of averaged draws. **Best** in each row and metric is in bold.

|     | $k = 1$         |                 |                    | $k = 2$         |                 |                    | $k = 4$         |                 |                    | $k = 8$         |                 |                    |
|-----|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
| $M$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 1   | 32.941          | 0.9198          | 0.0190             | <b>33.084</b>   | <b>0.9211</b>   | <b>0.0172</b>      | 33.037          | 0.9205          | <b>0.0172</b>      | 32.992          | 0.9200          | 0.0173             |
| 4   | 34.966          | 0.9463          | 0.0209             | <b>35.044</b>   | <b>0.9469</b>   | <b>0.0206</b>      | 35.001          | 0.9465          | <b>0.0206</b>      | 34.942          | 0.9458          | 0.0210             |
| 16  | 35.688          | 0.9537          | <b>0.0275</b>      | <b>35.729</b>   | <b>0.9540</b>   | <b>0.0275</b>      | 35.688          | 0.9536          | 0.0276             | 35.621          | 0.9529          | 0.0281             |
| 100 | 35.913          | 0.9558          | <b>0.0300</b>      | <b>35.943</b>   | <b>0.9561</b>   | 0.0302             | 35.902          | 0.9557          | 0.0303             | 35.833          | 0.9550          | 0.0308             |



Figure 12: **Source distribution, qualitative**. Six independent draws from a fixed CelebA Gaussian-deblurring measurement under the three sources, with the per-pixel standard deviation over draws on the same  $\times 20$  scale in every row. The posterior source varies where the measurement is uninformative; both Gaussian sources collapse onto a single reconstruction.

## B.6 Irreducible regression error: SNaP vs. isotropic sources

Proposition 2 fixes the target of the exact one-step map, but not how hard that map is to learn. We therefore compare sources through the regression problem that training solves. At  $r = t$  the Mean-Flow target equation 14 reduces to the conditional velocity  $x_1 - x_0$ , which the network regresses on  $(z_t, y)$  with  $z_t = (1 - t)x_0 + tx_1$ . The error left by the best predictor,

$$\mathcal{E}_t(y) := \mathbb{E}[\text{tr Cov}(x_1 - x_0 | z_t, y) | y], \quad (21)$$

cannot be removed by any network, so a source with a smaller  $\mathcal{E}_t$  poses an easier regression problem at time  $t$ .

**Lemma 2.** *Let  $\sigma_n > 0$ , fix  $y$  and  $t \in (0, 1)$ , and let  $x_1 \sim p(\cdot | y)$  have finite second moments. Let  $x_0 \sim q_\tau(\cdot | y)$ , and let  $\tilde{x}_0 = \tilde{K}y + \tilde{\xi}$  be a source of the form equation 6 with  $\tilde{\xi} \sim \mathcal{N}(0, \tilde{\Sigma})$  and  $\tilde{\Sigma} \succeq \tau^2 W$ , each conditionally independent of  $x_1$  given  $y$ . Let  $\mathcal{E}_t(y)$  and  $\tilde{\mathcal{E}}_t(y)$  be the errors equation 21 for the interpolants  $z_t = (1 - t)x_0 + tx_1$  and  $\tilde{z}_t = (1 - t)\tilde{x}_0 + tx_1$ . Then  $\mathcal{E}_t(y) \leq \tilde{\mathcal{E}}_t(y)$ . In particular, this holds for the isotropic source  $\mathcal{N}(0, \sigma^2 I)$  for every  $\sigma \geq \tau$ , since  $\tau^2 W \preceq \tau^2 I \preceq \sigma^2 I$ .*

The lemma requires nothing of the posterior beyond finite second moments, and the center of the source does not enter it: given  $y$ , the center only shifts  $z_t$  by a constant, which conditioning removes. The center is what Proposition 1 chooses to shorten the displacement; the lemma shows that the covariance  $\tau^2 W$  is what lowers the regression error. Both isotropic sources of Table 5 satisfy its assumption, since  $\tau = 0.15 \leq 1$  there.

*Proof.* Given  $y$ , write  $x_0 = m_y + \eta$  and  $\tilde{x}_0 = c + \tilde{\xi}$  with  $c := \tilde{K}y$ , where  $\eta \sim \mathcal{N}(0, \tau^2 W)$  and  $\tilde{\xi}$  are independent of  $x_1$ . Since  $m_y$  and  $c$  are constants given  $y$ , the maps  $z \mapsto (z - (1 - t)m_y)/t$  and  $z \mapsto (z - (1 - t)c)/t$  are invertible, so conditioning on  $z_t$  and on  $\tilde{z}_t$  is the same as conditioning on

$$\bar{z}_t := x_1 + \frac{1-t}{t} \eta \quad \text{and} \quad \tilde{\bar{z}}_t := x_1 + \frac{1-t}{t} \tilde{\xi},$$




Figure 13: **Source distribution, averaging budget.** PSNR (left, higher is better) and LPIPS (right, lower is better) against the number of averaged draws  $M$  on CelebA Gaussian deblurring. Only the posterior source responds to  $M$ ; the Gaussian sources are flat because their draws carry almost no variance to average away.




Figure 14: One-step ( $K=1$ ) posterior samples on the  $4 \times 4$  checkerboard. Rows are the two forward operators and columns are the learned adapter of VFM, the SNaP source, and the exact rejection-sampled posterior. The green line is the measurement constraint  $\mathbf{a}^\top \mathbf{x} = \mathbf{y}$ , and the dashed ellipse on each method panel is that method's  $2\sigma$  source covariance, the learned diagonal  $q_\phi(\mathbf{z} | \mathbf{y})$  for VFM and the closed-form  $\tau^2 \mathbf{W}$  for SNaP. In the top row the measurement direction is axis-aligned, both sources align with the constraint, and both recover the posterior. In the bottom row the constraint tilts. The SNaP source tilts with it, whereas the diagonal adapter cannot and degenerates to a near-isotropic blob, so its samples concentrate on one mode and spill off the checkerboard support.

respectively. Each is  $\mathbf{x}_1$  observed through additive Gaussian noise, and the two differ only in that noise.

Since  $\tilde{\Sigma} - \tau^2 \mathbf{W} \succeq \mathbf{0}$ , let  $\zeta \sim \mathcal{N}(\mathbf{0}, \tilde{\Sigma} - \tau^2 \mathbf{W})$  be independent of  $(\mathbf{x}_1, \eta)$  given  $\mathbf{y}$ . Then  $\eta + \zeta \sim \mathcal{N}(\mathbf{0}, \tilde{\Sigma})$ , and because  $\tilde{\mathcal{E}}_t(\mathbf{y})$  depends only on the joint law of  $\mathbf{x}_1$  and  $\tilde{\mathbf{z}}_t$  given  $\mathbf{y}$ , we may take  $\tilde{\xi} = \eta + \zeta$ . This gives  $\tilde{\mathbf{z}}_t = \tilde{\mathbf{z}}_t + \frac{1-t}{t} \zeta$ , so  $\mathbf{x}_1 \rightarrow \tilde{\mathbf{z}}_t \rightarrow \tilde{\mathbf{z}}_t$  is a Markov chain given  $\mathbf{y}$ .

Let  $\mu_t := \mathbb{E}[\mathbf{x}_1 | \tilde{\mathbf{z}}_t, \mathbf{y}]$  and  $\tilde{\mu}_t := \mathbb{E}[\mathbf{x}_1 | \tilde{\mathbf{z}}_t, \mathbf{y}]$ . The Markov property gives  $\mathbb{E}[\mathbf{x}_1 | \tilde{\mathbf{z}}_t, \tilde{\mathbf{z}}_t, \mathbf{y}] = \mu_t$ , and the tower rule gives  $\tilde{\mu}_t = \mathbb{E}[\mu_t | \tilde{\mathbf{z}}_t, \mathbf{y}]$ . Hence  $\mathbf{x}_1 - \mu_t$  is orthogonal to  $\mu_t - \tilde{\mu}_t$ , and

$$\mathbb{E}[\|\mathbf{x}_1 - \tilde{\mu}_t\|_2^2 | \mathbf{y}] = \mathbb{E}[\|\mathbf{x}_1 - \mu_t\|_2^2 | \mathbf{y}] + \mathbb{E}[\|\mu_t - \tilde{\mu}_t\|_2^2 | \mathbf{y}]. \quad (22)$$

Table 11: Posterior MMD<sup>2</sup> ( $\downarrow$ ) on the two-dimensional checkerboard benchmark at  $K=1$ , mean  $\pm$  standard deviation over three seeds, each at its validation-selected setting. The aligned operator is a parity check, where a diagonal adapter is exactly expressive, and the rotated operator differs only in the direction of  $\mathbf{a}$ .

| Method                 | Aligned (parity)    | Rotated (discriminating) |
|------------------------|---------------------|--------------------------|
| SNaP                   | $0.0027 \pm 0.0003$ | $0.0171 \pm 0.0012$      |
| VFM (joint)            | $0.0109 \pm 0.0007$ | $0.0737 \pm 0.0199$      |
| VFM (frozen $\theta$ ) | $0.0102 \pm 0.0003$ | $0.1900 \pm 0.0015$      |

Table 12: NullFlow vs. SNaP on CelebA box inpainting ( $M=1$ ) as test noise grows.

| Test $\sigma_n$ | NullFlow        |                 |                    | SNaP            |                 |                    |
|-----------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                 | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 0.00            | 33.50           | 0.977           | 0.012              | 31.37           | 0.947           | 0.025              |
| 0.01            | 33.27           | 0.970           | 0.012              | 31.35           | 0.947           | 0.024              |
| 0.02            | 32.65           | 0.952           | 0.016              | 31.30           | 0.946           | 0.023              |
| 0.05            | 30.11           | 0.855           | 0.055              | 30.74           | 0.934           | 0.021              |

Finally,  $\mathbf{x}_1 - \mathbf{x}_0 = (\mathbf{x}_1 - \mathbf{z}_t)/(1 - t)$  and  $\mathbf{z}_t$  is fixed once it is conditioned on, so  $\text{Cov}(\mathbf{x}_1 - \mathbf{x}_0 | \mathbf{z}_t, \mathbf{y}) = (1 - t)^{-2} \text{Cov}(\mathbf{x}_1 | \mathbf{z}_t, \mathbf{y})$  and  $\mathcal{E}_t(\mathbf{y}) = (1 - t)^{-2} \mathbb{E}[\|\mathbf{x}_1 - \boldsymbol{\mu}_t\|_2^2 | \mathbf{y}]$ ; likewise  $\tilde{\mathcal{E}}_t(\mathbf{y}) = (1 - t)^{-2} \mathbb{E}[\|\mathbf{x}_1 - \tilde{\boldsymbol{\mu}}_t\|_2^2 | \mathbf{y}]$ . Dividing equation 22 by  $(1 - t)^2$  gives

$$\tilde{\mathcal{E}}_t(\mathbf{y}) - \mathcal{E}_t(\mathbf{y}) = \frac{1}{(1 - t)^2} \mathbb{E}[\|\boldsymbol{\mu}_t - \tilde{\boldsymbol{\mu}}_t\|_2^2 | \mathbf{y}] \geq 0.$$

□

**Displacement.** The same source also shortens the path relative to the isotropic source. Let  $\tilde{\mathbf{x}}_0 \sim \mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$  be independent of  $(\mathbf{x}_1, \mathbf{n})$ , and let  $E_i := \mathbb{E}[(\mathbf{v}_i^\top \mathbf{x}_1)^2]$  for a right singular vector  $\mathbf{v}_i$  of  $\mathbf{A}$ . For any  $\mathbf{x}_1$  with finite second moments,

$$\mathbb{E}[(\mathbf{v}_i^\top (\mathbf{x}_1 - \tilde{\mathbf{x}}_0))^2] - \mathbb{E}[(\mathbf{v}_i^\top (\mathbf{x}_1 - \mathbf{x}_0))^2] = (1 - w_i^2)E_i + \tau^2(1 - w_i)^2 \geq 0, \quad (23)$$

with equality if and only if  $s_i = 0$ . Indeed, writing  $a_i := \mathbf{v}_i^\top \mathbf{x}_1$  and  $n_i := \mathbf{u}_i^\top \mathbf{n}$  (with  $n_i := 0$  if  $s_i = 0$ ), the Tikhonov center gives  $\mathbf{v}_i^\top (\mathbf{x}_1 - \mathbf{x}_0) = w_i a_i - \frac{s_i}{s_i^2 + \lambda} n_i - \mathbf{v}_i^\top \boldsymbol{\eta}$ , a sum of uncorrelated terms with second moments  $w_i^2 E_i$ ,  $\tau^2 w_i(1 - w_i)$  and  $\tau^2 w_i$ , whereas the isotropic source gives  $E_i + \tau^2$ .

**Evaluation.** Figure 15 evaluates  $\mathcal{E}_t$  for  $q_\tau(\cdot | \mathbf{y})$  and  $\mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$  in two settings where it is available exactly. On the two-dimensional mixture of App. B.7 at  $y = 2.5$ , the conditional covariance is computed by quadrature against the closed-form posterior. The two sources coincide in the null direction there, so the gap comes entirely from the observed direction. On CelebA Gaussian deblurring, gap is evaluated under a stationary Gaussian prior with the empirical power spectrum, for which  $\mathbf{A}$ ,  $\mathbf{W}$  and the posterior covariance are simultaneously diagonal in the Fourier basis. There the gap is present at every  $t$  and peaks at 56%, since the graded Fourier weights  $w_i$  leave the source something to exploit in every direction.

**Scope.** Three limits apply. First, the lemma concerns the  $r = t$  part of the objective. For  $r < t$  the target equation 14 also contains the network’s own derivatives, so the lemma does not transfer verbatim to the interior of the  $(r, t)$  range; a fraction  $p_{\text{ratio}}$  of training pairs have  $r = t$  (Table 9). Second, like the displacement in equation 7,  $\mathcal{E}_t$  decreases as the source covariance shrinks, so it cannot be used to choose  $\Sigma$ , and the lemma only compares  $q_\tau$  with sources that are at least as spread. A narrower source can have a smaller  $\mathcal{E}_t$ : on the mixture of App. B.7, where  $\tau = 3$ , the source  $\mathcal{N}(\mathbf{0}, \mathbf{I})$  has a smaller  $\mathcal{E}_t$  than  $q_\tau$  at every  $t$ , because it is narrower in the null direction. Third, the lemma compares regression problems, not trained samplers; Table 5 and App. B.7 compare the samplers directly, with everything but the source held fixed.






Figure 15: **Irreducible velocity-regression error**  $\mathcal{E}_t$  of equation 21 under the SNaP source  $q_\tau(\cdot | \mathbf{y})$  and the isotropic source  $\mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I})$ . **Left:** two-dimensional mixture of App. B.7 ( $\tau = 3$ ,  $\sigma_n = 0.3$ ,  $y = 2.5$ ), exact. **Right:** CelebA Gaussian deblurring ( $\tau = 0.15$ ,  $\sigma_n = 0.05$ ) under a stationary Gaussian prior. Lemma 2 guarantees the ordering; the shaded region is the gap.

App. B.6 shows that, for any posterior with finite second moments, the velocity-regression error that no network can remove is smaller under  $q_\tau(\cdot | \mathbf{y})$  than under any isotropic source  $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$  with  $\sigma \geq \tau$  (Lemma 2); Table 5 and App. B.7 measure the effect on the learned map.

### B.7 Verification on a Gaussian-mixture inverse problem

We verify the construction on a problem where the posterior is available in closed form. With a four-component Gaussian-mixture prior on  $\mathbb{R}^2$  and observation of the first coordinate ( $\mathbf{A} = \mathbf{e}_0^\top$ ,  $\sigma_n = 0.3$ ), the posterior is itself a Gaussian mixture, and for  $y \approx 2.5$  it is bimodal over  $x[1] \in \{\pm 2.5\}$  with the MMSE estimate falling between the modes. The setting therefore separates posterior sampling from posterior-mean estimation: a point estimator lands in the empty valley between the modes, and a mode-collapsed sampler reports the component variance  $\varsigma^2 = 0.12$  in the null direction rather than the posterior variance 6.36.

**Setup.** The clean data  $\mathbf{x}_1 \in \mathbb{R}^2$  is drawn from a four-component equal-weight Gaussian mixture with isotropic component covariance  $\varsigma^2 \mathbf{I}$  ( $\varsigma = 0.35$ ) and means at  $(\pm a, \pm a)$ ,  $a = 2.5$ . The forward operator observes only the first coordinate, so  $m = 1$ ,  $n = 2$  and  $\text{null}(\mathbf{A}) = \text{span}(\mathbf{e}_1)$ . The working-prior scale is  $\tau = 3$ .

**Closed-form posterior.** For each mixture component  $k$ ,

$$\tilde{\Sigma}_k = (\Sigma_k^{-1} + \mathbf{A}^\top \mathbf{A} / \sigma_n^2)^{-1}, \quad \tilde{\mu}_k = \tilde{\Sigma}_k (\Sigma_k^{-1} \mu_k + \mathbf{A}^\top \mathbf{y} / \sigma_n^2), \quad (24)$$

and

$$\tilde{\pi}_k \propto \pi_k \mathcal{N}(\mathbf{y}; \mathbf{A} \mu_k, \mathbf{A} \Sigma_k \mathbf{A}^\top + \sigma_n^2), \quad p(\mathbf{x}_1 | \mathbf{y}) = \sum_k \tilde{\pi}_k \mathcal{N}(\tilde{\mu}_k, \tilde{\Sigma}_k). \quad (25)$$

All quantities labeled “posterior” below are computed from this expression.

**Network and optimization.** An  $x$ -prediction MLP (width 256, depth 4, SiLU) with random-Fourier embeddings (dimension 64) for  $r$ ,  $t$  and  $\sigma_n$ , conditioned on  $\mathbf{y}$ , with  $\mathbf{u}_{r,t}^\theta = (\text{net}(\mathbf{z}, r, t, \mathbf{y}, \sigma_n) - \mathbf{z}) / (1 - r)$ . We use the objective equation 15 with  $p = 1$  and  $c = 10^{-3}$ ;  $t$  is logit-normal ( $\mu = 0$ ,  $\sigma = 1$ ) and  $r = t$  with probability 0.75, else  $r \sim \mathcal{U}(0, t)$ . AdamW, learning rate  $2 \times 10^{-4}$ , batch size 4096, 20,000 steps, EMA decay 0.999. One-step sampling is  $\hat{\mathbf{x}}_1 = \text{net}(\mathbf{x}_0, r=0, t=1, \mathbf{y}, \sigma_n)$ ; few-step sampling composes  $\mathbf{u}_{r,t}$  over a uniform partition of  $[0, 1]$ .

**The source is exact.** Drawing  $2 \times 10^5$  samples from the perturb-and-solve construction in equation 13 at  $y = 2.5$  and comparing against the closed form reproduces  $\mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$  to Monte-Carlo accuracy:  $\|\Delta \mathbf{m}_y\| = 3.8 \times 10^{-3}$  and  $\|\Delta \text{cov}\|_F = 4.8 \times 10^{-3}$  (Table 15). The anisotropy predicted by equation 10 is visible directly: the observed direction is pinned at variance 0.089 while the null direction carries the full working-prior variance  $\tau^2 = 9$ .

**Posterior recovery.** Figure 16 shows the source, the closed-form posterior, and  $2 \times 10^4$  one-step samples for three in-distribution measurements. The flow transports the vertically-spread source onto both posterior modes at  $x[1] = \pm 2.5$  while pinning  $x[0] \approx y$ , and the MMSE estimate falls in the empty valley between them. In the null direction the one-step sampler produces variance 6.66 against the closed-form 6.36; a mode-collapsed sampler would report  $\varsigma^2 = 0.12$ . Table 13 reports mean, per-direction variance and  $\text{MMD}^2$  at three measurements.

**Per-direction calibration.** Projecting onto the right singular vectors of  $\mathbf{A}$  (here  $\mathbf{e}_0$  with  $s = 1$  and  $\mathbf{e}_1$  with  $s = 0$ ), the flow assigns almost all of the null direction’s freedom (6.66 against posterior 6.36, source 9.0) and contracts the well-measured direction from the source value 0.089 toward the posterior 0.053, reaching 0.079 (Table 16).

**Gaussian ground truth.** With a single Gaussian prior matched to the working prior ( $\Sigma = \tau^2 \mathbf{I}$ ,  $\tau = \varsigma = 0.35$ ), the source *is* the posterior and the optimal map is the identity. Retraining reproduces the source to  $\|\Delta \mathbf{m}_y\| = 5.6 \times 10^{-4}$  and matches the Gaussian posterior with mean errors  $\leq 0.05$  and  $\text{MMD}^2 \leq 3 \times 10^{-3}$  (Table 14).

Table 13: Posterior match at three in-distribution measurements.

| $y$  | mean (post) | mean (SNaP)     | $\ \Delta \text{mean}\ $ | Var (observed, null) post $\rightarrow$ SNaP | $\text{MMD}^2$       |
|------|-------------|-----------------|--------------------------|----------------------------------------------|----------------------|
| 2.5  | (2.5, 0)    | (2.474, 0.065)  | 0.069                    | (0.052, 6.37) $\rightarrow$ (0.079, 6.66)    | $4.0 \times 10^{-3}$ |
| 2.0  | (2.212, 0)  | (2.045, 0.051)  | 0.174                    | (0.052, 6.37) $\rightarrow$ (0.084, 6.64)    | $1.1 \times 10^{-2}$ |
| -2.5 | (-2.5, 0)   | (-2.444, 0.108) | 0.121                    | (0.052, 6.37) $\rightarrow$ (0.081, 6.59)    | $2.6 \times 10^{-3}$ |

Table 14: Single-Gaussian ground-truth check. The map reduces to the Wiener estimate, with the posterior mean recovered to within 0.05 and the posterior variances over-estimated by 37–39%.

| $y$ | mean (post) | mean (SNaP)     | $\ \Delta \text{mean}\ $ | Var (observed, null) post $\rightarrow$ SNaP | $\text{MMD}^2$       |
|-----|-------------|-----------------|--------------------------|----------------------------------------------|----------------------|
| 0   | (0, 0)      | (-0.040, 0.026) | 0.048                    | (0.052, 0.123) $\rightarrow$ (0.071, 0.171)  | $2.7 \times 10^{-3}$ |
| 1   | (0.576, 0)  | (0.553, 0.012)  | 0.026                    | (0.052, 0.123) $\rightarrow$ (0.069, 0.160)  | $1.7 \times 10^{-3}$ |

## C Implementation Details

### C.1 Datasets and evaluation protocol

**CelebA.** Official aligned images with the canonical partition of 162,770/19,867/19,962 for the train/validation/test images, center-cropped to  $178 \times 178$  and resized to  $128 \times 128$ . We report results on the first 100 images of the test set.

**AFHQ-Cat.** 5,153/500/100 train/validation/test images, resized to  $256 \times 256$ . AFHQ-Cat has no official validation split. We followed [20] for train/val/test split.

**fastMRI brain.** The AXT2 subset, complex multi-coil  $k$ -space at  $320 \times 320$ . We split at the volume (patient) level into 615/77/77 train/val/test volumes, giving 4,305/539/539 slices after discarding the first four and last five slices of every volume, which contain little or no anatomy. Coil handling is full multi-coil SENSE, not RSS or an emulated single-coil reduction. The ground truth is the ESPIRiT coil-combined complex image  $\mathbf{x}_1 = \sum_i \bar{\mathbf{S}}_i \text{crop}(\mathbf{F}^H \mathbf{k}_i)$  on a  $320 \times 320$  grid, retained as two real channels (real/imaginary) and normalized per slice by  $\max |\mathbf{x}_1|$ ; phase is modeled throughout and magnitude is taken only when computing PSNR/SSIM. The forward operator is  $\mathbf{A} = \mathbf{MFS}$  with the same per-slice ESPIRiT maps at train and test time, zero-padded to 20 coils. Because the coils couple, the measurement-space Gram  $\mathbf{AA}^H$  is not diagonal and the posterior-source solve admits no closed form, unlike the inpainting, super-resolution, and deblurring operators, where  $\mathbf{A}^H \mathbf{A}$  is diagonal in the pixel or Fourier basis and the solve costs  $O(n)$  or a single FFT pair. We therefore solve it with conjugate gradient in the algebraically equivalent image-space form  $\mathbf{A}^H (\mathbf{AA}^H + \lambda \mathbf{I})^{-1} = (\mathbf{A}^H \mathbf{A} + \lambda \mathbf{I})^{-1} \mathbf{A}^H$ , whose spectrum lies in  $[\lambda, \approx 1]$  and whose right-hand side already lies in  $\text{range}(\mathbf{A}^H)$ ; the measurement-space form is rank-deficient by the coil count and, at the same iteration budget, returns an anchor worse than the plain adjoint.

Table 15: The square-root-free sampler equation 13 reproduces  $\mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$  to Monte-Carlo accuracy at  $y = 2.5$ .

|                 | $\mathbf{m}_y$   | $\tau^2 \mathbf{W}$ (diag) | error                                                                                           |
|-----------------|------------------|----------------------------|-------------------------------------------------------------------------------------------------|
| closed form     | [2.4752, 0]      | [0.0891, 9.000]            | —                                                                                               |
| RTO equation 13 | [2.4752, 0.0038] | [0.0889, 8.995]            | $\ \Delta \mathbf{m}_y\  = 3.8 \times 10^{-3}$ , $\ \Delta \text{cov}\ _F = 4.8 \times 10^{-3}$ |

Table 16: Per-direction variance at  $y = 2.5$  (equation 10). The source over-disperses both directions relative to the posterior; the trained flow moves both toward it, almost exactly in the null direction and partially in the observed one.

| direction                           | Var posterior | Var one-step SNaP | Var source |
|-------------------------------------|---------------|-------------------|------------|
| $\mathbf{e}_0$ (observed, $s = 1$ ) | 0.0518        | 0.0791            | 0.0891     |
| $\mathbf{e}_1$ (null, $s = 0$ )     | 6.372         | 6.656             | 9.000      |

## C.2 Solutions of the source construction

The source equation 13 requires one application of  $\mathbf{K} = \mathbf{A}^\top (\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I})^{-1}$  per measurement, with  $\lambda = \sigma_n^2 / \tau^2$ . Writing  $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ , the covariance is  $\tau^2 \mathbf{W}$  with  $\mathbf{W} = \mathbf{V} \text{diag}(w_i) \mathbf{V}^\top$  and  $w_i = \lambda / (s_i^2 + \lambda)$ . We summarise the cases used in this work.

**Image deblurring.** For deblurring  $\mathbf{A} = \mathbf{H}$  is a spatially invariant blur. Under circular boundary conditions  $\mathbf{H} = \mathbf{F}^H \mathbf{\Lambda} \mathbf{F}$  with  $\mathbf{\Lambda}$  the Fourier coefficients  $\hat{h}$  of the kernel, so  $\mathbf{H} \mathbf{H}^H = \mathbf{F}^H |\mathbf{\Lambda}|^2 \mathbf{F}$  and

$$\mathbf{m}_y = \mathbf{F}^H \left[ \frac{\overline{\mathbf{\Lambda}} \mathbf{F} \mathbf{y}}{|\mathbf{\Lambda}|^2 + \lambda} \right], \quad w_i = \frac{\lambda}{|\hat{h}_i|^2 + \lambda}, \quad (26)$$

elementwise in the Fourier domain, at the cost of two FFTs. Because  $|\hat{h}_i| > 0$  for every  $i$ , no  $w_i$  reaches 1 and no direction is unmeasured; the spectrum instead decays smoothly.

**Image inpainting.** For inpainting  $\mathbf{A} = \mathbf{M}$  is a row-selection matrix that keeps the observed coordinates, so  $\mathbf{M} \mathbf{M}^\top = \mathbf{I}_m$  and the solve is elementwise:

$$(\mathbf{m}_y)_i = \begin{cases} \frac{y_i}{1 + \lambda}, & i \in \Omega, \\ 0, & i \notin \Omega, \end{cases} \quad w_i = \begin{cases} \frac{\lambda}{1 + \lambda}, & i \in \Omega, \\ 1, & i \notin \Omega, \end{cases} \quad (27)$$

where  $\Omega$  is the set of observed pixels. Unobserved coordinates carry the full working-prior variance  $\tau^2$ , while observed coordinates are shrunk toward the prior mean by the factor  $1/(1 + \lambda)$  and retain variance  $\tau^2 \lambda / (1 + \lambda)$ . Box and random inpainting share this structure and differ only in the geometry of  $\Omega$ .

**Image super-resolution.** We use box-decimated super-resolution by implementing strided subsampling with no anti-alias prefilter,  $\mathbf{A} \mathbf{x} = \mathbf{x}[\dots, :: s, :: s]$ , so  $\mathbf{A} \mathbf{A}^\top = \mathbf{I}$ ,  $\mathbf{A}^\top$  is zero-filled upsampling, and equation 27 applies with  $\Omega$  the sampled grid locations.

**Multi-coil MRI.** With  $\mathbf{A} = \mathbf{M} \mathbf{F} \mathbf{S}$  the sensitivity maps  $\mathbf{S}$  are not unitary, so  $\mathbf{A} \mathbf{A}^\top$  is not diagonal and the solve with  $\mathbf{A} \mathbf{A}^\top + \lambda \mathbf{I}$  has no closed form. We solve it with conjugate gradients, warm-started at the zero-filled adjoint. Measured iteration counts and cost are reported in Table 17.

**General linear operators.** For a general  $\mathbf{A}$  with no exploitable structure, equation 13 is solved by conjugate gradients. All operators used in this paper fall under the cases above. The solve is required once per draw and does not depend on  $t$ ,  $r$ , or the network. For a fixed operator, the solve can be reduced to a diagonal inverse for diagonal (and Fourier diagonal) forward models or use a short conjugate gradient run for most others.

## C.3 Training and implementation details

**Architecture.** The backbone is a DDPM-style residual U-Net (GroupNorm<sub>32</sub>, Swish, variance-scaled initialization, zero-initialized output and residual-branch convolutions) with three total scalar




Figure 16: Posterior recovery for observations  $y = 2.5$  (left) and  $y = -2.5$  (right). The top row compares samples from the source distribution (green), the closed-form posterior (blue), and the one-step model (red). Dashed vertical lines indicate the observed values. The bottom row shows the corresponding marginals along the unobserved coordinate  $x_1$ ; the dashed line marks the posterior mean. The one-step model captures both posterior modes, although it exhibits some within-mode distortion relative to the exact posterior.

Table 17: Cost of the source construction equation 13. Larger  $\lambda$  (noisier data) improves conditioning and accelerates CG.

| Problem              | $\mathbf{A}\mathbf{A}^\top$ structure | Cost                  | Notes                                     |
|----------------------|---------------------------------------|-----------------------|-------------------------------------------|
| Inpainting (random)  | $\mathbf{I}_m$                        | free                  | $\mathbf{W}$ diagonal                     |
| Inpainting (box)     | $\mathbf{I}_m$                        | free                  | identical structure                       |
| MRI, single-coil     | $\mathbf{I}_m$                        | free                  | diagonal in $k$ -space                    |
| SR, average-pool     | $d^{-2}\mathbf{I}$                    | free                  | disjoint rows of norm $1/d$               |
| Deblurring, Gaussian | diag. in Fourier                      | 2 FFTs                | $\mathbf{A}\mathbf{A}^\top =  \hat{h} ^2$ |
| Deblurring, uniform  | diag. in Fourier                      | 2 FFTs                | $\hat{h}$ has exact zeros                 |
| SR, bicubic          | block-Fourier                         | $O(n \log n)$         | polyphase folding identity                |
| MRI, multi-coil      | not diagonal                          | CG, 10–20 it. typical | standard SENSE                            |

conditioning inputs. The time  $t$  is embedded by a sinusoidal positional encoding followed by a two-layer MLP into a 4 ch-dimensional vector; the second MeanFlow time  $r$  and the noise level  $\sigma_n$  each get their own identical embedding, and the three are *summed* before entering the body, where each residual block adds a per-channel affine projection of the sum to its first activation. The  $\sigma_n$  embedding takes  $10\sigma_n$  as input so that the small noise levels used here occupy a useful part of the sinusoidal range. Conditioning on  $\mathbf{y}$  is by *concatenation* rather than FiLM, with input  $[\mathbf{z}_t, \text{anchor}]$  of  $2C$  channels, where the anchor is the adjoint  $\mathbf{A}^\top \mathbf{y}$ .

Configurations: CelebA  $128 \times 128$  uses  $\text{ch} = 64$ ,  $\text{ch\_mult} = (1, 2, 2, 2)$ , two residual blocks per level and attention at  $16^2$  (9.1M parameters); AFHQ-Cat  $256 \times 256$  uses  $\text{ch} = 64$ ,  $\text{ch\_mult} = (1, 2, 4, 4, 8, 8)$  down to an  $8 \times 8$  bottleneck, two residual blocks, attention at  $32^2$  and  $16^2$  and dropout 0.1 (110.5M); fastMRI brain  $320 \times 320$  uses  $\text{ch} = 32$ ,  $\text{ch\_mult} = (1, 2, 4, 8, 8)$ , four residual blocks and attention at  $40^2$  (40.7M). Attention is a hand-written block of  $1 \times 1$  convolutional  $Q, K, V$  projections, a batched matrix product scaled by  $C^{-1/2}$ , and a `float32` softmax, since `torch.func.jvp` has no forward-mode rule for the fused scaled-dot-product kernels. The same

---

**Algorithm 1** SNaP training step.

---

|                                                                                                                           |                                                                              |
|---------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| 1: $\mathbf{x}_1 \sim p, \quad \mathbf{y} \leftarrow \mathbf{A}\mathbf{x}_1 + \mathbf{n}$                                 | $\triangleright$ training pair, equation 1                                   |
| 2: $t, r \leftarrow \text{SAMPLETIMES}()$                                                                                 | $\triangleright t$ logit-normal; $r = t$ with probability $p_{\text{ratio}}$ |
| 3: $\mathbf{x}_0 \leftarrow \text{SAMPLESOURCE}(\mathbf{y}, \sigma_n)$                                                    | $\triangleright$ perturb-and-solve, equation 13                              |
| 4: $\mathbf{z}_r \leftarrow (1 - r)\mathbf{x}_0 + r\mathbf{x}_1; \quad \mathbf{v} \leftarrow \mathbf{x}_1 - \mathbf{x}_0$ |                                                                              |
| 5: $\mathbf{u}, \text{d}_r \mathbf{u} \leftarrow \text{JVP}(\mathbf{u}^\theta; (\mathbf{z}_r, r, t), (\mathbf{v}, 1, 0))$ | $\triangleright$ one forward-mode pass                                       |
| 6: $\mathbf{u}_{\text{tgt}} \leftarrow \mathbf{v} + (t - r) \text{sg}[\text{d}_r \mathbf{u}]$                             | $\triangleright$ MeanFlow target, equation 14                                |
| 7: $\ell \leftarrow \ \mathbf{u} - \mathbf{u}_{\text{tgt}}\ _2^2 / n$                                                     |                                                                              |
| 8: $\mathcal{L} \leftarrow \text{sg}[\omega(\ell)] \ell$                                                                  | $\triangleright$ weighted objective, equation 15                             |
| 9: $\theta \leftarrow \text{ADAMW}(\theta, \nabla_\theta \mathcal{L})$                                                    | $\triangleright$ optimizer step                                              |

---

restriction rules out `torch.utils.checkpoint`, so activation memory is a hard ceiling and width is added only at the two deepest levels.

**Training.** Given a clean image  $\mathbf{x}_1$  we form  $\mathbf{y} = \mathbf{A}\mathbf{x}_1 + \mathbf{n}$ , draw the source  $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{m}_y, \tau^2 \mathbf{W})$  by equation 13, and set  $\mathbf{z}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$  and  $\mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0$ . Times are drawn as  $t = \sigma(\mu + s\epsilon)$  with  $\epsilon \sim \mathcal{N}(0, 1)$ ,  $\mu = 0$  and  $s = 1$  (logit-normal on  $(0, 1)$ ); with probability  $p_{\text{ratio}} = 0.5$  we set  $r = t$ , and otherwise  $r \sim \mathcal{U}(0, t)$ . A further 10% of the  $r \neq t$  samples, or 5% of the batch, is redrawn at the corner  $r \sim \mathcal{U}(0, 0.1)$ ,  $t \sim \mathcal{U}(0.9, 1)$ , because the base schedule reaches the slice at which one-step sampling is scored with probability only  $7.3 \times 10^{-4}$ .

The network is an  $\mathbf{x}_1$ -predictor, so  $\mathbf{u}(\mathbf{z}_r, r, 1) = (f_\theta(\mathbf{z}_r, r, 1) - \mathbf{z}_r)/(1 - r)$ . The total derivative is one forward-mode pass, with the differentiated arguments  $(\mathbf{z}, r, t)$  carrying tangent  $\mathbf{v}$  on  $\mathbf{z}$ , 1 on  $r$  and 0 on  $t$ , while the anchor, the coverage map and  $\sigma_n$  are closed over and carry zero tangent. The regression target is  $\mathbf{u}_{\text{tgt}} = \mathbf{v} + (t - r) \text{sg}[\text{d}\mathbf{u}/\text{d}r]$  and the loss is equation 15 with  $\Delta = \mathbf{u}^\theta - \mathbf{u}_{\text{tgt}}$ ,  $p = 1$  and  $c = 10^{-3}$ . The stop-gradient is on the weight only, and a second stop-gradient is applied to the derivative term. On CelebA and AFHQ,  $\ell$  is first divided by its all-reduced running mean (decay 0.99) so that  $c$  is dimensionless and the loss shape is stable across  $\tau$  and across training; the MRI runs use the unnormalized form. No endpoint  $\|\hat{\mathbf{x}}_1 - \mathbf{x}_1\|_2^2$  term is used, since its population minimizer is the posterior mean and would invalidate every calibration metric.

Optimization is AdamW with weight decay  $10^{-4}$ , cosine annealing to  $10^{-6}$  over the full epoch budget, no warm-up and gradient-norm clipping at 1.0. The learning rate is  $2 \times 10^{-4}$  on CelebA and AFHQ and  $1 \times 10^{-4}$  on MRI. Everything is trained in `float32`, since forward-mode AD is not autocast-safe. The effective batch is 32 throughout (CelebA  $8 \times 4$  GPUs; AFHQ  $2 \times 4$  accumulation steps  $\times 4$  GPUs; MRI  $4 \times 4$  GPUs). Budgets are 200 epochs (50,000 steps) on CelebA, 120 epochs (15,000 steps) for AFHQ super-resolution and 80 epochs (10,000 steps) for the remaining AFHQ operators, and 80 epochs (43,040 steps) for brain MRI, where an epoch is a fixed-size draw with replacement rather than a pass over the data. EMA is implemented but disabled in all reported runs, having never beaten the raw weights.

**Choice of  $\tau$ .** We set  $\tau^2$  to the mean posterior variance left along the directions the measurement leaves open, estimated offline on 256 training images as  $\tau^2 = \langle \mathbf{d}, \mathbf{W}\mathbf{d} \rangle / \text{tr } \mathbf{W}$ , where  $\mathbf{d} = \mathbf{x} - \hat{\mathbf{x}}(\mathbf{y})$  is the residual of the linear-MMSE (Wiener) estimate under the empirical image power spectrum. Since  $\mathbf{W}$  depends on  $\tau$ , we solve this by fixed-point iteration. The linear estimate is suboptimal, so its residual upper-bounds the posterior variance and  $\tau$  errs toward over-dispersion. This gives  $\tau = 0.1$  for  $\times 2$  super-resolution and random inpainting on CelebA, 0.15 for deblurring and 0.3 for box inpainting;  $\tau = 0.1$  for  $\times 4$  super-resolution, random inpainting and deblurring on AFHQ and 0.5 for box inpainting; and  $\tau = 0.25$  for brain MRI. The ablation in App. B shows that reconstruction quality changes little across this range. Dataset splits and preprocessing are given in App. C.1.

**Source solves.** The draw equation 13 needs only  $\mathbf{A}$ ,  $\mathbf{A}^\top$ , and one solve with  $(\mathbf{A}\mathbf{A}^\top + \lambda \mathbf{I})^{-1}$ , where  $\lambda = \sigma_n^2 / \tau^2$  is floored at  $10^{-8}$ , and for the operators used here that solve is a diagonal inverse, except for multi-coil MRI.

Circular Gaussian deblurring is Fourier-diagonal, with  $\mathbf{A}\mathbf{A}^\top = |H(\omega)|^2$  and one division per mode. Box-decimation super-resolution has  $\mathbf{A}\mathbf{A}^\top = \mathbf{I}$  on the low-resolution grid, so the solve is the scalar division by  $1 + \lambda$ . Inpainting has  $\mathbf{A}\mathbf{A}^\top = \mathbf{M}$  for a 0/1 mask, giving a pixel-diagonal division by



$M + \lambda$ . Super-resolution and inpainting are  $\mathcal{O}(n)$  and Gaussian deblurring is  $\mathcal{O}(n \log n)$  from the FFT pair; none add any measurable time to a training step.

Multi-coil MRI is the only iterative case, and we solve it in image space using  $\mathbf{A}^\top (\mathbf{A}\mathbf{A}^\top + \lambda\mathbf{I})^{-1} = (\mathbf{A}^\top \mathbf{A} + \lambda\mathbf{I})^{-1} \mathbf{A}^\top$ , running conjugate gradient on  $\mathbf{A}^\top \mathbf{A} + \lambda\mathbf{I}$ , whose spectrum lies in  $[\lambda, \approx 1]$  and whose right-hand side already lies in  $\text{range}(\mathbf{A}^\top)$ , with a cap of 200 iterations and relative-residual tolerance  $10^{-6}$ , which is typically reached in 10–20 iterations and per-sample reductions so that batched slices do not share a step size. The measurement-space route is implemented but not used, since  $\mathbf{A}\mathbf{A}^\top$  acts on  $k$ -space while factoring through the image and is therefore rank-deficient by the coil count on top of the mask; a short CG there leaves the anchor over-scaled at 6.5 dB, worse than the plain adjoint at 24.4 dB, against 26.0 dB for the image-space form at equal cost. Requesting the anchor  $\mathbf{m}_y$  inside the sampler splits the single solve into two by linearity, which is free for the diagonal operators and doubles the CG only for MRI.

**Cost.** All models were trained on four A6000 GPUs with manual data-parallel gradient all-reduce. Measured end-to-end wall-clock: AFHQ-Cat super-resolution 6 h 07 m for 120 epochs; AFHQ random inpainting, box inpainting and deblurring 4 h 04 m, 4 h 00 m and 4 h 02 m for 80 epochs each; CelebA box inpainting 3 h 29 m for 200 epochs; and brain MRI at  $\times 8$  acceleration 19 h 05 m for 80 epochs. Throughput on the 256-pixel AFHQ configuration is 6.6 images/s/GPU at 26.1 GB peak memory with batch size 2; the 110.5M model costs 18% more memory and 21% less throughput than a 28.0M baseline, because memory is activation- rather than parameter-dominated and the added width sits at the  $16^2$  and  $8^2$  levels. At inference an  $M$ -sample posterior costs  $M$  network evaluations and  $M$  source draws. Each draw is  $\mathcal{O}(n)$  for super-resolution and inpainting,  $\mathcal{O}(n \log n)$  for Gaussian deblurring, and a CG solve for multi-coil MRI.

**Baseline budgets.** The NFE column of each table counts network evaluations per reported output. SNaP and the MSE regressor use one evaluation per draw, so an  $M$ -draw row costs  $M$  evaluations. On natural images PnP-GS uses 23, DDRM 20, DiffPIR 100, OT-ODE 180 and DPS 1000; on fastMRI, DiffPIR uses 100 and CSGM, DPS, DAPS and PnP-DM use 1000. The subscript on PnP-Flow1, PnP-Flow5, Flower1-OT and Flower5-OT is  $N_{\text{Avg}}$ , the number of network evaluations averaged at each time step, so a run of  $N$  steps costs  $N \cdot N_{\text{Avg}}$  evaluations; we use the values tabulated by [21]. On CelebA both methods use  $N = 100$  for every operator, so Tab. 1 reports the five-evaluation variants at 500 evaluations, and the single-evaluation variant Flower1-OT of Tab. 2 at 100. On AFHQ-Cat the step counts depend on the operator: PnP-Flow uses  $N = 500$  for deblurring and super-resolution, 200 for random inpainting and 100 for box inpainting, while Flower uses  $N = 100, 500, 200$  and 100 on those same four operators. The budgets behind the NFE range quoted in Tab. 6 are therefore 2500, 2500, 1000 and 500 evaluations for PnP-Flow5, and 500, 2500, 1000 and 500 for Flower5-OT. Every SNaP entry in Tab. 2 is a single draw ( $M=1$ ); averaged draws are reported only in the distortion tables and in the ablations of App. B.

#### C.4 Metric details

For natural images, PSNR is computed on  $[0, 1]$  with data range 1, SSIM on the clamped  $[0, 1]$  image, and LPIPS with an AlexNet backbone on the clamped  $[-1, 1]$  image. Images are normalized to  $[-1, 1]$  for training and mapped back to  $[0, 1]$  for scoring. For MRI, PSNR and SSIM are computed on magnitude images, per slice, with the data range set to the maximum intensity of the ground-truth slice. We do not report LPIPS on MRI: its backbone is trained on RGB natural images and its perceptual judgments do not transfer to grayscale magnitude reconstructions.

**Calibration ratio.** We check whether the spread of a sampler’s draws has the right size for a sampler of  $p(\mathbf{x} \mid \mathbf{y})$ . Let  $p_s(\mathbf{x} \mid \mathbf{y})$  be the distribution of the draws and  $\boldsymbol{\mu}_s(\mathbf{y})$  its mean, and let  $(\mathbf{x}_1, \mathbf{y})$  be a ground-truth image and its measurement. Define

$$B := \mathbb{E}_{\mathbf{x}_1, \mathbf{y}} [\|\boldsymbol{\mu}_s(\mathbf{y}) - \mathbf{x}_1\|_2^2], \quad V := \mathbb{E}_{\mathbf{y}} [\text{tr Cov}_s(\mathbf{x} \mid \mathbf{y})].$$

If  $p_s = p$ , then  $\boldsymbol{\mu}_s(\mathbf{y}) = \mathbb{E}[\mathbf{x} \mid \mathbf{y}]$  and

$$B = \mathbb{E}_{\mathbf{y}} [\mathbb{E}_{\mathbf{x}_1 \mid \mathbf{y}} [\|\mathbb{E}[\mathbf{x} \mid \mathbf{y}] - \mathbf{x}_1\|_2^2]] = \mathbb{E}_{\mathbf{y}} [\text{tr Cov}(\mathbf{x} \mid \mathbf{y})] = V.$$

Hence  $V/B = 1$  is a necessary condition for  $p_s = p$ , and we report the calibration ratio

$$\Delta_{\text{cal}} := \frac{V}{B}. \quad (28)$$

Values below 1 indicate draws with too little spread, collapsed toward a point estimate, and values above 1 indicate over-dispersed draws.

**Estimation.** For each test pair  $(\mathbf{x}_1^{(i)}, \mathbf{y}^{(i)})$ ,  $i = 1, \dots, N$ , we draw  $\mathbf{s}_{(1)}(\mathbf{y}^{(i)}), \dots, \mathbf{s}_{(M)}(\mathbf{y}^{(i)})$  with mean  $\bar{\mathbf{s}}(\mathbf{y}^{(i)})$  and use

$$\hat{V} := \frac{1}{N(M-1)} \sum_{i=1}^N \sum_{j=1}^M \left\| \mathbf{s}_{(j)}(\mathbf{y}^{(i)}) - \bar{\mathbf{s}}(\mathbf{y}^{(i)}) \right\|_2^2,$$

$$\hat{B} := \frac{1}{N} \sum_{i=1}^N \left\| \bar{\mathbf{s}}(\mathbf{y}^{(i)}) - \mathbf{x}_1^{(i)} \right\|_2^2 - \frac{\hat{V}}{M},$$

which are unbiased for  $V$  and  $B$ ; the correction  $\hat{V}/M$  removes the variance that remains in the mean of  $M$  draws. We use  $M = 100$ .  $\Delta_{\text{cal}}$  is related to the spread-skill ratio used in ensemble forecasting, which compares the standard deviation of the ensemble with the root-mean-squared error of its mean and therefore corresponds to  $\sqrt{V/B}$ .

**FID.** FID is computed with InceptionV3 pool3 features on 2048 single-draw reconstructions, quantised to uint8, against reference statistics from 2048 ground-truth images.