

# 000 001 002 003 004 005 **FLOWER: A FLOW-MATCHING SOLVER FOR INVERSE** 006 **PROBLEMS**

007  
008  
009  
010  
011 **Anonymous authors**  
012 Paper under double-blind review

## ABSTRACT

013 We introduce *Flower*, a solver for linear inverse problems. It leverages a pre-  
014 trained flow model to produce reconstructions that are consistent with the observed  
015 measurements. *Flower* operates through an iterative procedure over three steps:  
016 (i) a flow-consistent destination estimation, where the velocity network predicts a  
017 denoised target; (ii) a refinement step that projects the estimated destination onto  
018 a feasible set defined by the forward operator; and (iii) a time-progression step  
019 that re-projects the refined destination along the flow trajectory. We provide a  
020 theoretical analysis that demonstrates how *Flower* approximates Bayesian pos-  
021 terior sampling, thereby unifying perspectives from plug-and-play methods and  
022 generative inverse solvers. On the practical side, *Flower* achieves state-of-the-art  
023 reconstruction quality while using nearly identical hyperparameters across various  
024 linear inverse problems.

## 1 INTRODUCTION

025 Inverse problems are central to computational imaging and computer vision (McCann & Unser,  
026 2019; Zeng, 2001). Their goal is to reconstruct an underlying signal  $\mathbf{x} \in \mathbb{R}^d$  from its observed  
027 measurements  $\mathbf{y} \in \mathbb{R}^M$ . Here, we focus on linear inverse problems, such that the acquisition of the  
028 measurements follows the model

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n} \quad (1)$$

029 for some linear forward operator  $\mathbf{H}: \mathbb{R}^d \rightarrow \mathbb{R}^M$  and additive white Gaussian noise  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$ .  
030 From a Bayesian perspective, the simplest reconstruction approach is to obtain the maximum-  
031 likelihood estimation

$$\hat{\mathbf{x}}_{\text{MLE}} = \arg \max_{\mathbf{x} \in \mathbb{R}^d} p_{\mathbf{Y}|\mathbf{X}=\mathbf{x}}(\mathbf{y}) = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2. \quad (2)$$

032 However, this problem is ill-posed and yields poor-quality solutions. Another approach is to obtain  
033 the maximum a posteriori estimation (MAP)

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x} \in \mathbb{R}^d} p_{\mathbf{X}|\mathbf{Y}=\mathbf{y}}(\mathbf{x}) = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \left( \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 - \log p_{\mathbf{X}}(\mathbf{x}) \right), \quad (3)$$

034 which requires the knowledge of the prior distribution  $p_{\mathbf{X}}$  of images, a quantity that is generally  
035 unknown. The minimization problem of equation 3 is consistent with the variational perspective  
036 of inverse problems, where the term  $(-\log p_{\mathbf{X}}(\mathbf{x}))$  is replaced by a regularizer  $\mathcal{R}(\mathbf{x})$  that encodes  
037 some properties of the images. From classic signal processing to the advent of deep learning, the  
038 design of a good regularizer  $\mathcal{R}$  has been of interest. Classic signal processing relies on the smoothness  
039 or sparsity of images to introduce wavelet- or total-variation-based regularizers (Rudin et al., 1992;  
040 Figueiredo & Nowak, 2003; Beck & Teboulle, 2009). Some methods build upon classical models  
041 and try to learn such criteria in a data-driven manner (Roth & Black, 2009; Goujon et al., 2024;  
042 Ducotterd et al., 2025; Pourya et al., 2025). Plug-and-play (PnP) approaches focus on the implicit  
043 replacement of  $\mathcal{R}$  by its proximal operator, with a learned neural network that serves as a denoiser  
044 (Venkatakrishnan et al., 2013; Zhang et al., 2022; Hurault et al., 2022b;a). Although MAP estima-  
045 tions tend to have a good reconstruction quality, they do not necessarily provide the minimum-mean-  
046 square estimator  $\hat{\mathbf{x}}_{\text{MMSE}}$  that is best in terms of the peak signal-to-noise ratio (PSNR). To estimate  
047  $\hat{\mathbf{x}}_{\text{MMSE}}$ , one would have to compute the posterior mean  $\hat{\mathbf{x}}_{\text{MMSE}} = \mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}]$ . Moreover, for  
048

054 perceptual metrics, it is better to generate a sample from  $p_{\mathbf{X}|\mathbf{Y}=\mathbf{y}}$  instead of an estimator of the  
055 distribution.

056 The objective of generative modeling is to sample from a target distribution  $p_{\mathbf{X}}$ . In practice, this  
057 distribution is unknown, and one typically only has access to a finite collection of its samples.  
058 Numerous approaches have been proposed to address this issue. Among them, diffusion models  
059 (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020) and, more recently, flow-matching  
060 methods (Lipman et al., 2023) represent the state of the art in scalable generative modeling for  
061 images.

062 Flow matching, introduced by Lipman et al. (2023), constructs a continuous-time generative process  
063 by parameterizing the velocity field of an ordinary differential equation (ODE) as a neural network.  
064 It takes inspiration from optimal transport and continuous normalizing flows (Ambrosio et al., 2008;  
065 Hagemann et al., 2022) and transports an initial source distribution  $p_{\mathbf{X}_0}$  to a target distribution  
066  $p_{\mathbf{X}_1} \approx p_{\mathbf{X}}$ . The choice of probability paths from  $p_{\mathbf{X}_0}$  to  $p_{\mathbf{X}_1}$  are numerous, with Gaussian paths  
067 recovering diffusion as a special case (Albergo & Vanden-Eijnden, 2023). However, flow matching  
068 mostly focuses on straight-line paths, which yields competitive performance and improved sampling  
069 efficiency (Liu et al., 2023; Liu, 2022).

070 The remarkable success of generative models in image generation motivates their extension to in-  
071 verse problems, where the goal shifts from the sampling of the prior distribution  $p_{\mathbf{X}1}$  to the sampling  
072 of the posterior  $p_{\mathbf{X}_1|\mathbf{Y}=\mathbf{y}}$ . Several inverse solvers based on diffusion models have been introduced  
073 (Chung et al., 2023; 2024; Kawar et al., 2022; Song et al., 2023; Zhu et al., 2023; Zhang et al., 2025;  
074 Mardani et al., 2024). Recent efforts also focus on flow-based solvers (Pokle et al., 2024; Martin  
075 et al., 2025). Existing approaches can be broadly grouped into two categories: (i) methods that ap-  
076 proximate the posterior score (velocity field) with gradient corrections along the generative path; and  
077 (ii) PnP strategies that alternate between generative (diffusion or flow) updates and data-consistency  
078 steps.

079 In this work, we introduce a novel solver based on flow matching that achieves state-of-the-art results  
080 among flow-based methods for linear inverse problems. Our approach departs from existing methods  
081 by framing the problem through a Bayesian ancestral-sampling perspective, which gives rise to a  
082 simple three-step procedure with a natural plug-and-play interpretation. Our main contributions are  
083 as follows.

- 085 **1. Flow-matching solver for inverse problems.** We introduce *Flower*, an inverse problem  
086 solver that consists of three steps: (1) a *flow-consistent destination estimation*, where  
087 the velocity network is used to predict a destination, interpretable as denoising; (2) a  
088 *measurement-aware refinement*, in which the estimated destination is projected onto the  
089 feasible set defined by the forward operator; and (3) a *time progression*, where the refined  
090 destination is re-projected along the flow path.
- 091 **2. Bayesian analysis and relation to PnP.** We provide a Bayesian analysis in which we  
092 demonstrate how and under what considerations *Flower* generates approximate poste-  
093 rior samples from  $p_{\mathbf{X}_1|\mathbf{Y}=\mathbf{y}}$ . Specifically, we show that Step 1 computes the con-  
094 ditional expectation  $\mathbb{E}[\mathbf{X}_1|\mathbf{X}_t = \mathbf{x}_t]$ , which we then use for the approximation  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t}$   
095 of  $p_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t}$ . Through this approximation, we show that Step 2 generates a sample  
096  $\tilde{\mathbf{x}}(\mathbf{x}_t, \mathbf{y}) \sim \tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ . Step 3 then updates the trajectory given the refined desti-  
097 nation  $\tilde{\mathbf{x}}(\mathbf{x}_t, \mathbf{y})$  and draws a sample  $\mathbf{x}_{t+\Delta t} \sim \tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ , which by induction and  
098 ancestral sampling, produces the final sample  $\mathbf{x}_1 \sim \tilde{p}_{\mathbf{X}_1|\mathbf{Y}=\mathbf{y}}$ . These steps rely on three as-  
099 sumptions: the velocity network is optimally trained for unconditional flow matching; the  
100 acquisition of measurements follows the forward model in equation 1; and the source and  
101 target distributions are independent. To the best of our knowledge, this Bayesian construc-  
102 tion is novel within flow-based solvers for inverse problems. Although this construction is  
103 key to the derivation of our solver, the resulting procedure closely mirrors the PnP meth-  
104 ods. Thus, our Bayesian justification establishes a link between the PnP approach and  
105 approximate posterior sampling with generative models for linear inverse problems. We  
106 also discuss a possible extension of *Flower* to nonlinear inverse problems.
- 107 **3. Numerical validation.** We first examine a controlled setup with Gaussian mixture mod-  
108 els and show that *Flower* successfully recovers posterior samples. We then evaluate our

108 method on standard inverse problem benchmarks for flow matching. We achieve competitive  
 109 performance, with nearly identical hyperparameters across all tasks.  
 110

111 The remainder of this paper is organized as follows. In Section 2, we review the fundamentals of  
 112 flow matching along with the mathematical tools required for the development of our method. We  
 113 then introduce *Flower* in Section 3 and present the associated theoretical analysis. In Section 4, we  
 114 discuss related work and highlight their similarities and differences with our approach. We report  
 115 our numerical results in Section 5. Finally, we provide a potential nonlinear extension of *Flower* in  
 116 Section 6.

## 117 2 BACKGROUND

### 118 2.1 FLOW MATCHING

121 Let  $p_{\mathbf{X}_0}$  be a source distribution that is easy to sample and let  $p_{\mathbf{X}_1}$  be a target distribution that we  
 122 want to sample from. A time-dependent flow  $\psi_t$  transports  $p_{\mathbf{X}_0}$  to  $p_{\mathbf{X}_1}$  via the ODE  
 123

$$124 \frac{d\psi_t(\mathbf{x})}{dt} = \mathbf{v}_t(\psi_t(\mathbf{x})), \quad t \in [0, 1], \quad (4)$$

125 for some velocity field  $\mathbf{v}_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ . The intermediate variables  $\mathbf{X}_t = \psi_t(\mathbf{X}_0)$  follow a distribution  
 126  $p_{\mathbf{X}_t}$ . The objective of flow matching is to approximate  $\mathbf{v}_t$  with a neural network  $\mathbf{v}_t^\theta$ , which will  
 127 allow us to sample from  $p_{\mathbf{X}_1}$ . However, the determination of the flow-matching loss  
 128

$$129 \mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0,1]} \mathbb{E}_{\mathbf{x}_t \sim p_{\mathbf{X}_t}} \left[ \left\| \mathbf{v}_t^\theta(\mathbf{x}_t) - \mathbf{v}_t(\mathbf{x}_t) \right\|_2^2 \right] \quad (5)$$

130 is challenging, as it requires access to the marginal velocity field  $\mathbf{v}_t(\mathbf{x}_t)$ . To address this, we focus  
 131 on the conditional velocity  $\mathbf{v}_t(\mathbf{x}_t | \mathbf{x}_1)$  and define the conditional straight-line flow and velocity  
 132

$$133 \mathbf{x}_t = \psi_t(\mathbf{x}_0 | \mathbf{x}_1) = (1-t)\mathbf{x}_0 + t\mathbf{x}_1, \quad \mathbf{v}_t(\mathbf{x}_t | \mathbf{x}_1) = \mathbf{x}_1 - \mathbf{x}_0. \quad (6)$$

134 This leads to the practical conditional flow-matching loss

$$135 \mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0,1]} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi} \left[ \left\| \mathbf{v}_t^\theta((1-t)\mathbf{x}_0 + t\mathbf{x}_1, t) - (\mathbf{x}_1 - \mathbf{x}_0) \right\|_2^2 \right], \quad (7)$$

136 where  $\pi \in \Pi(p_{\mathbf{X}_0}, p_{\mathbf{X}_1})$  is a coupling over  $(\mathbf{X}_0, \mathbf{X}_1)$ , given by joint distributions on  $\mathbb{R}^d \times \mathbb{R}^d$   
 137 with marginals  $p_{\mathbf{X}_0}$  and  $p_{\mathbf{X}_1}$ . Lipman et al. (2023) have shown that the minimization of  $\mathcal{L}_{\text{CFM}}$  is  
 138 equivalent to the minimization of  $\mathcal{L}_{\text{FM}}$ , since their gradients with respect to  $\theta$  are equal.  
 139

140 The coupling  $\pi$  determines how  $(\mathbf{x}_0, \mathbf{x}_1)$  are paired. With the *independent* (IND) coupling  $\pi =$   
 141  $p_{\mathbf{X}_0} \otimes p_{\mathbf{X}_1}$ , the training is simple and scalable. However, the resulting interpolated paths can over-  
 142 lap, which may slow down convergence. At the other extreme, the *optimal transport* (OT) cou-  
 143 pling  $\pi^* \in \arg \min_{\pi \in \Pi(p_{\mathbf{X}_0}, p_{\mathbf{X}_1})} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi} [\|\mathbf{x}_1 - \mathbf{x}_0\|_2^2]$  produces globally aligned pairs such  
 144 that straight-line flows approximate displacement interpolation along the Wasserstein-2 geodesic. If  
 145 the Monge map  $T$  satisfying  $T \# p_{\mathbf{X}_0} = p_{\mathbf{X}_1}$  exists and is known, then no training is needed: the  
 146 sampling  $\mathbf{x}_0 \sim p_{\mathbf{X}_0}$  and the computation of  $\mathbf{x}_1 = T(\mathbf{x}_0)$  already generate a sample from  $p_{\mathbf{X}_1}$ . In  
 147 practice,  $T$  is unknown and its approximation is infeasible. A practical compromise is *mini-batch*  
 148 *OT*, which solves an entropically regularized OT problem within each batch to compute an approxi-  
 149 mate coupling  $\hat{\pi}$ . This improves alignment over independence with moderate computational over-  
 150 head. For more details on the mathematical background of OT, such as the definition and uniqueness  
 151 of the Monge map, we refer to Peyré (2025).

152 The choice of the source distribution  $p_{\mathbf{X}_0}$  is crucial for effective training and sampling. In practice,  
 153  $p_{\mathbf{X}_0}$  is often chosen as the standard normal distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . With the independent coupling,  
 154 this simply yields  $p_{\mathbf{X}_t | \mathbf{X}_1 = \mathbf{x}_1} = \mathcal{N}(t\mathbf{x}_1, (1-t)^2\mathbf{I})$ . However, the computation of  $p_{\mathbf{X}_t | \mathbf{X}_1 = \mathbf{x}_1}$  in the  
 155 OT case is challenging due to the mini-batch approach, in which it is difficult to determine the batch  
 156 a sample  $\mathbf{x}_1$  came from, as well as its associated OT paths.  
 157

### 158 2.2 PROXIMAL OPERATOR

159 The *proximal operator* of a proper, lower semi-continuous convex function  $f: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$  is  
 160 defined as

$$161 \text{prox}_f(\mathbf{x}) = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \left( \frac{1}{2} \|\mathbf{w} - \mathbf{x}\|_2^2 + f(\mathbf{w}) \right). \quad (8)$$

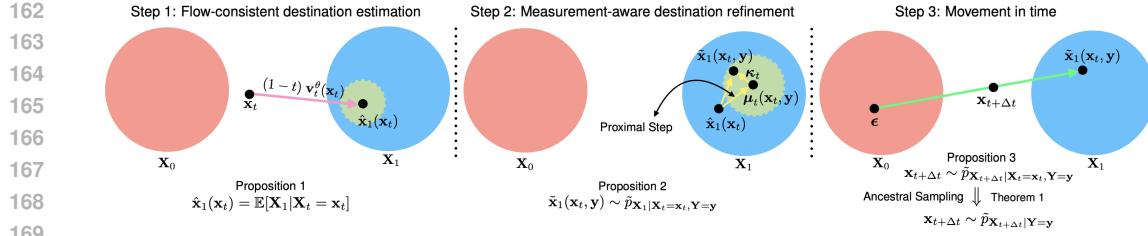


Figure 1: Overview of the three steps in *Flower*. Starting from an initial sample  $\mathbf{x}_0 \sim p_{\mathbf{X}_0}$  at time  $t$ , the method: Step 1 predicts a flow-consistent destination  $\hat{\mathbf{x}}_1(\mathbf{x}_t)$ ; Step 2 refines this destination using the measurements via a proximal step and associated uncertainty sampling to obtain  $\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y})$ ; and Step 3 updates the trajectory along time by interpolating  $\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y})$  with new noise  $\epsilon \sim p_{\mathbf{X}_0}$ . The  $N$ -time repetition of these steps yields the final reconstruction  $\mathbf{x}_1$ .

This operator can be interpreted as a generalized projection of  $\mathbf{x}$  onto a set associated with  $f$ , balancing proximity to  $\mathbf{x}$  and regularization by  $f$ . Proximal operators play a central role in optimization algorithms that solve inverse problems and are key components of proximal-gradient methods (Bubeck, 2015).

### 3 METHOD

Let  $\mathbf{v}_t^\theta$  denote a velocity network trained to generate samples from  $p_{\mathbf{X}_1}$  through flow matching. Therefore, starting from  $\mathbf{x}_0 \sim p_{\mathbf{X}_0}$ , we get a sample  $\mathbf{x}_1 \sim p_{\mathbf{X}_1}$  if we perform  $N$  iterations of the update equation

$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \Delta t \mathbf{v}_t^\theta(\mathbf{x}_t) \quad (9)$$

with  $\Delta t = \frac{1}{N}$ . We aim to use the pre-trained velocity network  $\mathbf{v}_t^\theta$  to generate solutions  $\mathbf{x}_1$  that are consistent with the flow and the linear forward model  $\mathbf{y} = \mathbf{H}\mathbf{x}_1 + \mathbf{n}$  for  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$ , as described in equation 1. To achieve this goal, we introduce *Flower* which, given the measurements  $\mathbf{y}$ , modifies the unconditional flow path of equation 9 and outputs  $\mathbf{x}_1$  by iterating  $N$  times over three steps. We first introduce these steps and then theoretically establish how and under what assumptions *Flower* approximates a sample  $\mathbf{x}_1$  of the conditional posterior  $p_{\mathbf{X}_1 | \mathbf{Y} = \mathbf{y}}$ . The three steps are as follows.

#### 1. Flow-consistent destination estimation

$$\hat{\mathbf{x}}_1(\mathbf{x}_t) = \mathbf{x}_t + (1-t)\mathbf{v}_t^\theta(\mathbf{x}_t). \quad (10)$$

#### 2. Measurement-aware destination refinement

$$\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y}) = \mu_t(\mathbf{x}_t, \mathbf{y}) + \gamma \kappa_t \quad (11)$$

for

$$\mu_t(\mathbf{x}_t, \mathbf{y}) = \text{prox}_{\nu_t^2 F_{\mathbf{y}}}(\hat{\mathbf{x}}_1(\mathbf{x}_t)), \quad \kappa_t \sim \mathcal{N}(\mathbf{0}, \Sigma_t), \quad (12)$$

where  $F_{\mathbf{y}}(\mathbf{x}) = \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2$  and  $\text{prox}_{\nu_t^2 F_{\mathbf{y}}}$  denotes the proximal operator of  $\nu_t^2 F_{\mathbf{y}}$  as defined in equation 8. We have that  $\nu_t = \frac{(1-t)}{\sqrt{t^2 + (1-t)^2}}$  and that  $\Sigma_t = (\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}$ . The hyperparameter  $\gamma \in \{0, 1\}$  controls the consideration of the uncertainty of the destination refinement step.

#### 3. Movement in time

$$\mathbf{x}_{t+\Delta t} = (1-t-\Delta t)\epsilon + (t+\Delta t)\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y}), \quad (13)$$

where  $\epsilon$  is newly sampled from  $p_{\mathbf{X}_0}$  at each iteration.

Here,  $\Delta t = \frac{1}{N}$  and the scheme is initialized with a sample  $\mathbf{x}_0 \sim p_{\mathbf{X}_0}$ . In Figure 2.2, we present a visual illustration of these three steps. We also summarize these steps in Algorithm 1 of the Appendix.

We now interpret *Flower* through a Bayesian lens. We assume that, at each iteration, the three steps collectively draw  $\mathbf{x}_{t+\Delta t}$  from the transition distribution  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ . Under this assumption and with proper initialization, the procedure performs ancestral sampling along the conditional trajectory. By induction, we obtain  $\mathbf{x}_{t+\Delta t} \sim p_{\mathbf{X}_{t+\Delta t}|\mathbf{Y}=\mathbf{y}}$ . We formalize this in Theorem 1, with proof in Appendix 8.2.1, which in turn implies that the final sample  $\mathbf{x}_1$  produced by *Flower* follows the desired posterior  $p_{\mathbf{X}_1|\mathbf{Y}=\mathbf{y}}$ . We then detail how, in practice, the three steps realize a draw from  $\tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ , which serves as an approximation of  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ .

**Theorem 1.** *Let  $\mathbf{x}_0$  be a sample from  $p_{\mathbf{X}_0|\mathbf{Y}=\mathbf{y}}$ . If  $\mathbf{x}_t$  is a sample from  $p_{\mathbf{X}_t|\mathbf{Y}=\mathbf{y}}$ , then the sample  $\mathbf{x}_{t+\Delta t}$  from  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  follows  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{Y}=\mathbf{y}}$ .*

**Remark 1.** *For the inductive argument to hold, *Flower* must be initialized with a sample from the conditional distribution  $p_{\mathbf{X}_0|\mathbf{Y}=\mathbf{y}}$ . When  $\mathbf{X}_0$  and  $\mathbf{X}_1$  are assumed to be independent, this reduces to a sampling from the unconditional prior  $p_{\mathbf{X}_0}$ , which is often chosen as  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ .*

Theorem 1 presupposes the existence of an ancestral-sampling scheme to generate samples from  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{Y}=\mathbf{y}}$ . This scheme requires a sampling from the transition distribution  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ . We now describe how to realize this transition in practice. We proceed sequentially and explain the details of each step of *Flower*.

First, under the assumption that  $\mathbf{v}_t^\theta$  is the optimal velocity network, we show in Proposition 1 that the predicted  $\hat{\mathbf{x}}_1(\mathbf{x}_t)$  in Step 1 equals the conditional expectation  $\mathbb{E}[\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t]$ . The proof is provided in Appendix 8.2.2.

**Proposition 1.** *If  $\mathbf{v}^\theta(\mathbf{x}_t, t) = \mathbf{v}_t^*(\mathbf{x})$  is a pre-trained velocity vector field that minimizes the conditional flow-matching loss, then*

$$\hat{\mathbf{x}}_1(\mathbf{x}_t) = \mathbb{E}[\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t] = \mathbf{x}_t + (1-t)\mathbf{v}^\theta(\mathbf{x}_t, t). \quad (14)$$

Since the distribution  $p_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t}$  is not directly available, we propose to approximate it with

$$\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t} = \mathcal{N}(\hat{\mathbf{x}}_1(\mathbf{x}_t), \nu_t^2 \mathbf{I}), \quad (15)$$

an isotropic Gaussian distribution centered at  $\hat{\mathbf{x}}_1(\mathbf{x}_t) = \mathbb{E}[\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t]$  with a time-varying covariance. As  $t \rightarrow 1$ , the distribution  $p_{\mathbf{X}_t}$  approaches the target  $p_{\mathbf{X}_1}$ . For  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t}$  to be consistent with this property,  $\nu_t$  should anneal in time. We choose  $\nu_t = (1-t)/\sqrt{t^2 + (1-t)^2}$ , which results in the valid covariance when  $p_{\mathbf{X}_1}$  is a standard Gaussian distribution. Our approximation is indeed the IIGDM approximation proposed by Song et al. (2023) within diffusion solvers and later by Pokle et al. (2024) for flow matching. However, instead of having a score-based interpretation and using this approximation to obtain  $\nabla_{\mathbf{x}_t} \log \tilde{p}_{\mathbf{Y}|\mathbf{X}_t=\mathbf{x}_t}$ , we propose to sample  $\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y})$  from the distribution  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  that approximates  $p_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ . To this end, we show in Proposition 2 that  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  is indeed a Gaussian distribution, using the IIGDM approximation and the forward model of equation 1. The proof is provided in Appendix 8.2.3.

**Proposition 2.** *Suppose that  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t} = \mathcal{N}(\hat{\mathbf{x}}_1(\mathbf{x}_t), \nu_t^2 \mathbf{I})$  (IIGDM approximation) and  $p_{\mathbf{Y}|\mathbf{X}_1=\mathbf{x}_1} = \mathcal{N}(\mathbf{H}\mathbf{x}_1, \sigma_n^2 \mathbf{I})$  (measurement operation). Then,  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}} = \mathcal{N}(\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}), \boldsymbol{\Sigma}_t)$ , where*

$$\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}) = (\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1} (\nu_t^{-2} \hat{\mathbf{x}}_1(\mathbf{x}_t) + \sigma_n^{-2} \mathbf{H}^\top \mathbf{y}), \quad (16)$$

$$\boldsymbol{\Sigma}_t = (\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}. \quad (17)$$

Proposition 2 allows us to sample from  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  as provided in Step 2 of *Flower* using the re-parameterization trick in equation 11. However, the  $\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y})$  in Step 2 (see equation 12) is described using a proximal operator which differs from equation 16. It is easy to verify the equivalence between the two, through the fact that  $\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y})$  of equation 16 can be written as the solution to the minimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left( \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 + \frac{1}{2\nu_t^2} \|\mathbf{x} - \hat{\mathbf{x}}_1(\mathbf{x}_t)\|_2^2 \right). \quad (18)$$

This directly results in  $\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}) = \text{prox}_{\nu_t^2 F}(\hat{\mathbf{x}}_1(\mathbf{x}_t))$  for  $F_{\mathbf{y}}(\mathbf{x}) = \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2$  under the definition of the proximal operator in equation 8. Moreover, the sampling from the anisotropic

270 Gaussian  $\mathcal{N}(\mathbf{0}, \Sigma_t)$  is not trivial; however, if we sample two independent  $\epsilon_1 \in \mathbb{R}^d$  and  $\epsilon_2 \in \mathbb{R}^M$   
 271 from standard Gaussian distributions, then we verify in Appendix 8.2.5 that  $\kappa_t = \Sigma_t(\nu_t^{-1}\epsilon_1 +$   
 272  $\sigma_n^{-1}\mathbf{H}^\top\epsilon_2)$  follows  $\mathcal{N}(\mathbf{0}, \Sigma_t)$ .

273 Step 3 of *Flower* aims to sample the distribution  $p_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ , which is also what we required  
 274 for our ancestral-sampling procedure to hold. We now show that if we have a sample  $\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y})$   
 275 from  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ , then we could obtain a sample from the distribution  $\tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  using  
 276 equation 13. We first compute  $\tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  under the assumption that  $p_{\mathbf{X}_0}$  is independent of  
 277  $p_{\mathbf{X}_1}$  in Proposition 3 which we prove in Appendix 8.2.4.

278 **Proposition 3.** *If from the pre-trained flow matching we have that  $p_{\mathbf{X}_0} = \mathcal{N}(\mathbf{0}, \mathbf{I})$ , if  $p_{\mathbf{X}_0}$  is inde-  
 279 pendent of  $p_{\mathbf{X}_1}$ , and if  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}(\mathbf{x}_1) = \mathcal{N}(\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}), \boldsymbol{\Sigma}_t)$ , then it holds that*

$$281 \tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}} = \mathcal{N}((t + \Delta t)\boldsymbol{\mu}_t, (t + \Delta t)^2\boldsymbol{\Sigma}_t + (1 - t - \Delta t)^2\mathbf{I}). \quad (19)$$

283 From Proposition 8.2.4 and by the means of the re-parametrization trick, it is easy to verify that

$$284 \mathbf{x}_{t+\Delta t} = (1 - t - \Delta t)\boldsymbol{\epsilon} + (t + \Delta t)\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y}) \quad (20)$$

285 follows  $\tilde{p}_{\mathbf{X}_{t+\Delta t}|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  given that  $\tilde{\mathbf{x}} \sim \mathcal{N}(\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}), \boldsymbol{\Sigma}_t)$ , which happens in Step 3 of *Flower*.

286 In our Bayesian justification, we assumed that  $p_{\mathbf{X}_0}$  and  $p_{\mathbf{X}_1}$  are independent. This assumption ex-  
 287 cludes, for example, the mini-batch optimal-transport coupling. Nevertheless, in practice, *Flower*  
 288 can still be applied in such settings and interpreted in a PnP manner. The iterative structure of  
 289 *Flower* closely resembles PnP methods: Step 1 acts as a denoising step, while Step 2 enforces data  
 290 consistency. From this viewpoint, Step 3 can be seen as a re-projection onto the flow trajectory, as  
 291 discussed in Martin et al. (2025). However, rather than relying solely on this interpretation, we pro-  
 292 vide a Bayesian justification of the procedure. This perspective highlights a conceptual link between  
 293 PnP methods and posterior sampling. We also address this empirically in our numerical results. For  
 294 exact posterior sampling to hold under our approximations,  $\gamma$  should be set to one. Interestingly,  
 295 in practice we find that the choice  $\gamma = 0$  (i.e., ignoring the uncertainty in the destination refine-  
 296 ment step) yields a better reconstruction quality. We provide further discussion on this effect with  
 297 our numerical results in Section 5. Moreover, our framework remains valid for more general noise  
 298 distributions  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_n)$  where  $\mathbf{R}_n$  is any symmetric positive-definite matrix, as detailed in  
 299 Appendix 8.3.

## 4 RELATED WORKS

301 An extensive body of work adapts pre-trained diffusion or flow priors to inverse problems by modi-  
 302 fying the dynamics to approximate the conditional posterior. We first review diffusion-based solvers,  
 303 then flow-based ones. Throughout this section, we highlight how *Flower* differs from similar meth-  
 304 ods. We use the flow-matching notation with source  $\mathbf{X}_0$  and target  $\mathbf{X}_1$ .

305 Among diffusion solvers, DPS (Chung et al., 2023) approximates  $p_{\mathbf{Y}|\mathbf{X}_t=\mathbf{x}_t}$  by  $p_{\mathbf{Y}|\mathbf{X}_t=\hat{\mathbf{x}}_1(\mathbf{x}_t)}$ , where  
 306  $\hat{\mathbf{x}}_1(\mathbf{x}_t)$  is the diffusion-based denoised version of  $\mathbf{x}_t$ , which leads to a gradient correction to the dif-  
 307 fusion dynamics. IIIGDM (Song et al., 2023) approximates  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t} = \mathcal{N}(\hat{\mathbf{x}}_1(\mathbf{x}_t), \nu_t^2 \mathbf{I})$  for some  
 308 time-annealing  $\nu_t$  and replaces the gradient correction of DPS with a pseudoinverse-based update.  
 309 *Flower* adopts the same approximation as IIIGDM, but the subsequent steps differ. DDS (Chung  
 310 et al., 2024) shares the same perspective as DPS but replaces the gradient with a proximal step mo-  
 311 tivated by a manifold-preserving gradient perspective. DiffPIR (Zhu et al., 2023) arrives at a very  
 312 similar structure through half-quadratic splitting, alternating proximal data updates with diffusion  
 313 denoising. Both DDS and DiffPIR are structurally close to *Flower* but, unlike *Flower*, they lack the  
 314 Bayesian justification that interprets the updates as posterior sampling. DAPS (Zhang et al., 2025)  
 315 also uses ancestral sampling similar to *Flower*, but instead of directly computing  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$ , it  
 316 applies Langevin updates for its evaluation, which is more computationally demanding but extends  
 317 naturally to nonlinear inverse problems.

318 In the flow-matching domain, OT-ODE (Pokle et al., 2024) employs a IIIGDM-based approximation,  
 319 similar in spirit to that of *Flower*. However, instead of adopting our ancestral sampling scheme, they,  
 320 similar to the approach of IIIGDM, approximate the score of the conditional distribution  $\tilde{p}_{\mathbf{Y}|\mathbf{X}_t=\mathbf{x}_t}$   
 321 in order to construct the new velocity field. Therefore, while our method, *Flower*, relies on the same  
 322 approximation for  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t}$ , our methods act in different ways. Flow-Priors (Zhang et al., 2024)

tackle inverse problems by reformulating the MAP objective as a sequence of time-dependent MAP subproblems with closed-form evaluations by taking advantage of the velocity network. However, unlike *Flower*, their approach relies on the computation of  $\text{Tr } \nabla \mathbf{v}_t^\theta$ , which is costly. D-Flow (Ben-Hamu et al., 2024) adopts an implicit regularization strategy, replacing the data-fidelity objective  $\mathbf{x} \mapsto \|\mathbf{Hx} - \mathbf{y}\|^2$  with a latent loss  $\mathbf{z} \mapsto \|\mathbf{H}(f(1, \mathbf{z})) - \mathbf{y}\|^2$ , where  $f$  is the solution of the flow ODE. The latent loss is non-convex with an implicit regularization effect that prevents convergence to trivial solutions. The optimization is performed by back-propagation through ODE solutions, which is more computationally demanding compared to the steps of *Flower*. PnP-Flow (Martin et al., 2025) introduces a PnP framework that employs the velocity network as a denoiser. Its update steps are similar to *Flower*, but we replace their gradient update with a proximal operation, which leads to improved reconstruction quality. Moreover, *Flower* offers a Bayesian justification of the process while PnP-Flow is purely plug-and-play.

## 5 NUMERICAL RESULTS

Here, we benchmark *Flower* against state-of-the-art flow-based inverse solvers across a range of linear inverse problems. In Appendix 8.4.1, we validate the Bayesian interpretation of *Flower* through a toy experiment with Gaussian mixtures, where ground-truth posterior samples are computable. We also present further numerical results for Fourier sampling and non-isotropic Gaussian noise in Appendix 8.5.7.

We implement *Flower* as described in Algorithm 1 of the Appendix. As a practical note, the computation  $\Sigma_t \mathbf{b}$  for any  $\mathbf{b} \in \mathbb{R}^d$  requires inversion  $(\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1} \mathbf{b}$ . Wherever this inversion is required, we instead solve the corresponding linear system  $(\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}) \mathbf{z} = \mathbf{b}$  using conjugate gradients (CG) with a maximum of 50 iterations and an  $\ell_2$ -residual tolerance of  $10^{-5}$  and return  $\mathbf{z}$  as the solution of the operation. We found the CG implementation sufficiently efficient in practice due to the positive-definite structure of  $\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}$ .

### 5.1 BENCHMARK EXPERIMENTS

The goal of this section is to benchmark our method against other flow-matching-based solvers for inverse problems. For fair comparisons, we adopt the benchmark introduced by Martin et al. (2025), which also includes state-of-the-art PnP and diffusion models. We describe the datasets and experimental setup for completeness, present quantitative results in Tables 1 and 2, and provide qualitative examples in Figure 2. Finally, we discuss the key observations, highlighting the performance of *Flower* and its empirical considerations.

We use two datasets for our numerical comparisons. First, we use  $(128 \times 128)$  human-face images from Yang et al. (2015), denoted by CelebA. Second, we use resized  $(256 \times 256)$  cat images from Choi et al. (2020), denoted by AFHQ-Cat. We normalize all images to the range  $[-1, 1]$ . We train on the full training sets of both datasets. We tune the hyperparameters of different methods using a validation set. For CelebA, we use 32 images from the dataset’s validation split. AFHQ-Cat has no validation split, so following Martin et al. (2025), we construct one by selecting 32 images from the test set and removing them from that set. For reporting the metrics, we use 100 test images from each dataset, a limit imposed by the computational cost of baseline methods (D-Flow and Flow Priors). For faster methods (including *Flower*), we also report metrics on larger datasets consisting of 1000 images for CelebA and 400 for AFHQ, which are presented in Appendix 8.5.6. These extended results show the same performance trends as the 100-image evaluations presented in this section.

We compare *Flower* against flow-matching solvers OT-ODE (Pokle et al., 2024), D-Flow (Ben-Hamu et al., 2024), Flow-Priors (Zhang et al., 2024), and PnP-Flow (Martin et al., 2025), as well as two other baselines: PnP-GS (Hurault et al., 2022b), a state-of-the-art plug-and-play method, and DiffPIR Zhu et al. (2023), a diffusion-based inverse solver. All models (except DiffPIR) use the same U-Net backbone (Ronneberger et al., 2015) trained with Mini-Batch OT Flow Matching (Tong et al., 2024) and a Gaussian latent prior. Pre-trained weights for the flow models and PnP-GS are taken from Martin et al. (2025), trained with learning rate  $10^{-4}$ : on CelebA for 200 epochs (batch size 128) and on AFHQ-Cat for 400 epochs (batch size 64). For *Flower*, we additionally train a variant without latent-target coupling (*Flower-IND*) using the same hyperparameters, which corresponds to our theoretical setting. While *Flower-IND* achieves higher performance (see Appendix 8.5.2), we

378 primarily report Flower-OT for consistency with other flow-matching baselines. Training DiffPIR  
 379 with this backbone proved ineffective due to limited capacity, so following Martin et al. (2025),  
 380 we adopt a pretrained model Choi et al. (2021) from the DeepInv library (Tachella et al., 2023),  
 381 originally trained on FFHQ (Karras et al., 2019). This introduces some mismatch but provides the  
 382 fairest diffusion-based baseline. Note that we use the latest checkpoints from Martin et al. (2025),  
 383 but our averaging strategy differs from theirs. In Martin et al. (2025), results are reported by grouping  
 384 four images into one batch and then averaging across 25 such batches. In contrast, we recomputed  
 385 the results using 100 independent averages over the images themselves. Consequently, our reported  
 386 numbers differ from those in Martin et al. (2025).

387 We evaluate performance on five restoration tasks: (i) denoising with Gaussian noise with  $\sigma_n = 0.2$ ;  
 388 (ii) deblurring with a  $61 \times 61$  Gaussian kernel ( $\sigma_b = 1.0$  for CelebA,  $\sigma_b = 3.0$  for AFHQ-Cat) and  
 389 additive noise  $\sigma_n = 0.05$ ; (iii) super-resolution ( $2 \times$  downsampling for CelebA and  $4 \times$  for AFHQ-  
 390 Cat, with  $\sigma_n = 0.05$ ); (iv) random inpainting with 70% of pixels removed ( $\sigma_n = 0.01$ ); and (v) box  
 391 inpainting with a centered  $40 \times 40$  mask for CelebA and  $80 \times 80$  mask for AFHQ-Cat ( $\sigma_n = 0.05$ ).  
 392 To report quantitative results, we use peak signal-to-noise ratio (PSNR), structural similarity index  
 393 measure (SSIM), and learned perceptual image patch similarity (LPIPS). Note that for PSNR and  
 394 SSIM, higher values indicate better performance, while for LPIPS, lower values are better.

395 To ensure fair comparisons, we adopt the optimal hyperparameters reported in Martin et al. (2025)  
 396 for each method, obtained via grid search on the validation set to maximize PSNR. For PnP-Flow,  
 397 we report two variants: PnP-Flow1 and PnP-Flow5, which apply one and five evaluations of the  
 398 velocity network per denoising step, respectively. For *Flower*, we follow the same procedure, re-  
 399 porting in Tables 1 and 2 either the output of a single evaluation (Flower1-OT) or the average of five  
 400 evaluations (Flower5-OT). A key property of *Flower* is that, apart from the number  $N$  of iterations  
 401 and the knowledge of the noise level  $\sigma_n$ , it uses the same hyperparameters across different inverse  
 402 problems, unlike other flow models. In particular, aside from  $N$ , the only hyperparameter of *Flower*  
 403 is  $\gamma$ , which controls the uncertainty of the destination refinement. Across all setups,  $\gamma = 0$  yields  
 404 higher reconstruction quality. As discussed in Appendix 8.5.3, the choice  $\gamma = 1$  produces sam-  
 405 ples that appear realistic but requires the averaging of multiple runs to achieve competitive PSNR,  
 406 whereas  $\gamma = 0$  attains better quality with fewer averages. This observation is consistent with our toy  
 407 experiments, where  $\gamma = 0$  encouraged sampling from higher-probability regions. For  $N$ , we always  
 408 match the number of steps used by our main competitor, PnP-Flow. We provide further ablation  
 409 studies on the effect of the number of evaluations for the averaging and different time discretiza-  
 410 tions within *Flower* in Appendix 8.5.4 and 8.5.5, respectively. The full hyperparameter details are  
 411 reported in Appendix 8.5.8.

412 **Key Observations.** On CelebA (Table 1), *Flower* achieves the best or near-best results across all  
 413 tasks, with clear gains in deblurring and box inpainting. The five-step averaging further improves the  
 414 results. On AFHQ-Cat (Table 2), *Flower* remains highly competitive and outperforms baselines in  
 415 deblurring, box inpainting, and random inpainting, while PnP-GS is strongest in denoising. In these  
 416 tables, bold numbers indicate the best results among single-average results of methods. Underlined  
 417 numbers indicate the second best. Blue numbers highlight the overall best across all methods. We  
 418 illustrate in Figure 2 representative reconstructions across denoising, deblurring, super-resolution,  
 419 and inpainting tasks. Compared to OT-ODE, D-Flow, and Flow-Priors, *Flower* consistently pro-  
 420 duces fewer artifacts, while also avoiding the over-smoothing often observed in PnP-Flow. These  
 421 visual trends align with the quantitative results and highlight the robustness of *Flower* across diverse  
 422 degradations. In Figure 3, we illustrate the solution path of *Flower* for the box inpainting task shown  
 423 in Figure 2. As expected, Step 1 produces flow-based denoised images, while Step 2 enforces con-  
 424 sistency with the measurements. In this specific box-inpainting setup, Step 2 primarily aligns the  
 425 region outside the box with the measurements and preserves the result of Step 1 inside the box.  
 426 Step 3 then mixes the refined destination with fresh source noise; as  $t$  increases, the noise decreases  
 427 and the reconstruction emerges. The injected noise is essential to prevent the velocity network from  
 428 getting stuck at the previous iterate and to allow it to predict improved destinations. As shown in  
 429 Table 3 of Appendix 8.5.1, *Flower* has a runtime that is similar to PnP-Flow and OT-ODE, with  
 430 only a slight overhead relative to PnP-Flow due to the proximal-projection step replacing a simple  
 431 gradient update, while requiring the same minimal memory. In contrast, D-Flow and Flow-Priors  
 432 are substantially slower and more memory-intensive.

432

433

Table 1: Results on 100 test images of the dataset CelebA.

| Method            | Denoising    |              |              | Deblurring   |              |              | Super-resolution |              |              | Random inpainting |              |              | Box inpainting |              |              |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|--------------|--------------|-------------------|--------------|--------------|----------------|--------------|--------------|
|                   | PSNR         | SSIM         | LPIPS        | PSNR         | SSIM         | LPIPS        | PSNR             | SSIM         | LPIPS        | PSNR              | SSIM         | LPIPS        | PSNR           | SSIM         | LPIPS        |
| Degraded          | 20.00        | 0.348        | 0.372        | 27.83        | 0.740        | 0.126        | 10.26            | 0.183        | 0.827        | 11.95             | 0.196        | 1.041        | 22.27          | 0.742        | 0.214        |
| PnP-GS            | <b>32.64</b> | 0.910        | 0.035        | 34.03        | 0.924        | 0.041        | 31.31            | 0.892        | 0.064        | 29.22             | 0.875        | 0.070        | -              | -            | -            |
| DiffPIR           | 31.20        | 0.885        | 0.060        | 32.77        | 0.912        | 0.060        | <u>31.52</u>     | 0.895        | 0.033        | 31.74             | 0.917        | 0.025        | -              | -            | -            |
| OT-ODE            | 30.54        | 0.859        | <b>0.032</b> | 33.01        | 0.921        | <u>0.029</u> | 31.46            | 0.907        | <b>0.025</b> | 28.68             | 0.871        | 0.051        | 29.40          | 0.920        | 0.038        |
| D-Flow            | 26.04        | 0.607        | 0.092        | 31.25        | 0.854        | 0.038        | 30.47            | 0.843        | 0.026        | <b>33.67</b>      | 0.943        | <b>0.015</b> | 30.70          | 0.899        | 0.026        |
| Flow-Priors       | 29.34        | 0.768        | 0.134        | 31.54        | 0.858        | 0.056        | 28.35            | 0.713        | 0.102        | 32.88             | 0.871        | 0.019        | 30.07          | 0.858        | 0.048        |
| PnP-Flow1         | 31.80        | 0.905        | 0.044        | <u>34.48</u> | 0.936        | 0.040        | 31.09            | 0.902        | 0.045        | 33.05             | <b>0.944</b> | 0.018        | 30.47          | 0.933        | 0.037        |
| Flower1-OT (ours) | <b>32.28</b> | <b>0.914</b> | 0.034        | <b>34.98</b> | <b>0.947</b> | <b>0.026</b> | <b>32.36</b>     | <b>0.923</b> | 0.034        | <u>33.08</u>      | <b>0.944</b> | 0.018        | <b>31.19</b>   | <b>0.945</b> | <b>0.022</b> |
| PnP-Flow5         | 32.30        | 0.911        | 0.056        | 34.80        | 0.940        | 0.047        | 31.49            | 0.906        | 0.056        | <b>33.98</b>      | <b>0.953</b> | 0.022        | 31.09          | 0.940        | 0.043        |
| Flower5-OT (ours) | <b>33.14</b> | <b>0.926</b> | 0.038        | <b>35.67</b> | <b>0.954</b> | 0.032        | <b>33.09</b>     | <b>0.932</b> | 0.040        | 33.95             | <b>0.953</b> | 0.020        | <b>31.87</b>   | <b>0.952</b> | 0.023        |

442

443

444

445

Table 2: Results on 100 test images of the dataset AFHQ-Cat.

| Method            | Denoising    |              |              | Deblurring   |              |              | Super-resolution |              |              | Random inpainting |              |              | Box inpainting |              |              |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|--------------|--------------|-------------------|--------------|--------------|----------------|--------------|--------------|
|                   | PSNR         | SSIM         | LPIPS        | PSNR         | SSIM         | LPIPS        | PSNR             | SSIM         | LPIPS        | PSNR              | SSIM         | LPIPS        | PSNR           | SSIM         | LPIPS        |
| Degraded          | 20.00        | 0.314        | 0.509        | 23.94        | 0.517        | 0.444        | 11.70            | 0.208        | 0.873        | 13.36             | 0.223        | 1.081        | 21.80          | 0.740        | 0.198        |
| PnP-GS            | <b>32.58</b> | <b>0.894</b> | <b>0.072</b> | <u>27.91</u> | 0.753        | 0.349        | 24.15            | 0.632        | 0.362        | 29.42             | 0.836        | 0.126        | -              | -            | -            |
| DiffPIR           | 30.58        | 0.835        | 0.189        | 27.56        | 0.728        | 0.342        | 23.65            | 0.624        | 0.402        | 31.70             | 0.881        | 0.062        | -              | -            | -            |
| OT-ODE            | 30.03        | 0.815        | <u>0.076</u> | 27.06        | 0.713        | <b>0.123</b> | 25.91            | 0.716        | <b>0.108</b> | 29.40             | 0.839        | 0.090        | 24.62          | 0.875        | 0.085        |
| D-Flow            | 26.13        | 0.574        | 0.175        | 27.82        | 0.721        | <u>0.164</u> | 24.64            | 0.601        | 0.190        | 32.20             | 0.894        | <u>0.040</u> | <b>26.26</b>   | 0.842        | <u>0.077</u> |
| Flow-Priors       | 29.41        | 0.763        | 0.153        | 26.47        | 0.700        | 0.181        | 23.51            | 0.570        | 0.272        | 32.37             | 0.906        | 0.047        | 26.20          | 0.818        | 0.118        |
| PnP-Flow1         | 31.18        | 0.863        | 0.135        | 27.87        | <u>0.760</u> | 0.304        | <b>26.94</b>     | <b>0.763</b> | <u>0.171</u> | <b>33.00</b>      | <b>0.918</b> | <b>0.037</b> | 26.00          | <u>0.897</u> | 0.103        |
| Flower1-OT (ours) | <u>31.69</u> | <b>0.879</b> | 0.102        | <b>28.64</b> | <b>0.775</b> | 0.255        | <u>26.23</u>     | <u>0.741</u> | 0.272        | <u>32.97</u>      | <b>0.918</b> | 0.040        | 26.19          | <b>0.915</b> | <b>0.063</b> |
| PnP-Flow5         | 31.43        | 0.864        | 0.168        | 28.19        | 0.766        | 0.332        | <b>27.37</b>     | <b>0.774</b> | 0.183        | <b>33.75</b>      | <b>0.929</b> | 0.048        | 26.68          | 0.901        | 0.120        |
| Flower5-OT (ours) | 32.35        | 0.891        | 0.116        | <b>28.97</b> | <b>0.784</b> | 0.283        | 26.57            | 0.075        | 0.282        | 33.70             | 0.927        | 0.045        | <b>26.88</b>   | <b>0.922</b> | 0.066        |

455

456

457

458

## 6 POTENTIAL EXTENSION TO NONLINEAR INVERSE PROBLEMS

459

460

*Flower* as it can handle different linear forward operators. To extend it to nonlinear inverse problems, Proposition 2 needs to be revisited. When the measurement operator  $\mathbf{H}$  is linear, the likelihood  $p_{\mathbf{Y}|\mathbf{X}_1=\mathbf{x}_1} = \mathcal{N}(\mathbf{H}\mathbf{x}_1, \sigma_n^2 \mathbf{I})$  is Gaussian, and, combined with the PIGDM Gaussian prior  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t} = \mathcal{N}(\hat{\mathbf{x}}_1(\mathbf{x}_t), \nu_t^2 \mathbf{I})$ , the approximate posterior remains Gaussian:  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}} = \mathcal{N}(\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}), \boldsymbol{\Sigma}_t)$ . This yields the closed-form mean  $\boldsymbol{\mu}_t$  and covariance  $\boldsymbol{\Sigma}_t$  in Proposition 2.

If the measurement model is nonlinear, i.e.,

$$p_{\mathbf{Y}|\mathbf{X}_1=\mathbf{x}_1} = \mathcal{N}(\mathbf{h}(\mathbf{x}_1), \sigma_n^2 \mathbf{I}) \quad (21)$$

for some nonlinear function  $\mathbf{h}: \mathbb{R}^d \rightarrow \mathbb{R}^M$ , then it still holds that

$$\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}(\mathbf{x}_1) \propto \exp\left(-\frac{1}{2\sigma_n^2} \|\mathbf{y} - \mathbf{h}(\mathbf{x}_1)\|^2 - \frac{1}{2\nu_t^2} \|\mathbf{x}_1 - \hat{\mathbf{x}}_1(\mathbf{x}_t)\|^2\right), \quad (22)$$

which is no longer Gaussian due to the nonlinearity of  $\mathbf{h}$ . Nevertheless, sampling from this distribution can be done using iterative sampling schemes, since the score (gradient of the log density) is available in closed form:

$$\nabla_{\mathbf{x}_1} \log \tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}(\mathbf{x}_1) = \sigma_n^{-2} \mathbf{J}_{\mathbf{h}}(\mathbf{x}_1)^\top (\mathbf{y} - \mathbf{h}(\mathbf{x}_1)) - \nu_t^{-2} (\mathbf{x}_1 - \hat{\mathbf{x}}_1(\mathbf{x}_t)), \quad (23)$$

where  $\mathbf{J}_{\mathbf{h}}$  denotes the Jacobian of  $\mathbf{h}$ . This score function could be used to generate samples of  $\tilde{p}_{\mathbf{X}_1|\mathbf{X}_t=\mathbf{x}_t, \mathbf{Y}=\mathbf{y}}$  via schemes such as Langevin dynamics which provides a valid substitute for Step 2 of *Flower*, which enables the handling of nonlinear cases without requiring modifications to the remaining steps.

## 7 CONCLUSION

481

483

484

485

We introduced *Flower*, a method that leverages pre-trained flow-matching models to solve linear inverse problems through a simple three-step iterative procedure. By combining flow-consistent predictions, measurement-aware refinement, and time evolution, *Flower* provides a principled Bayesian

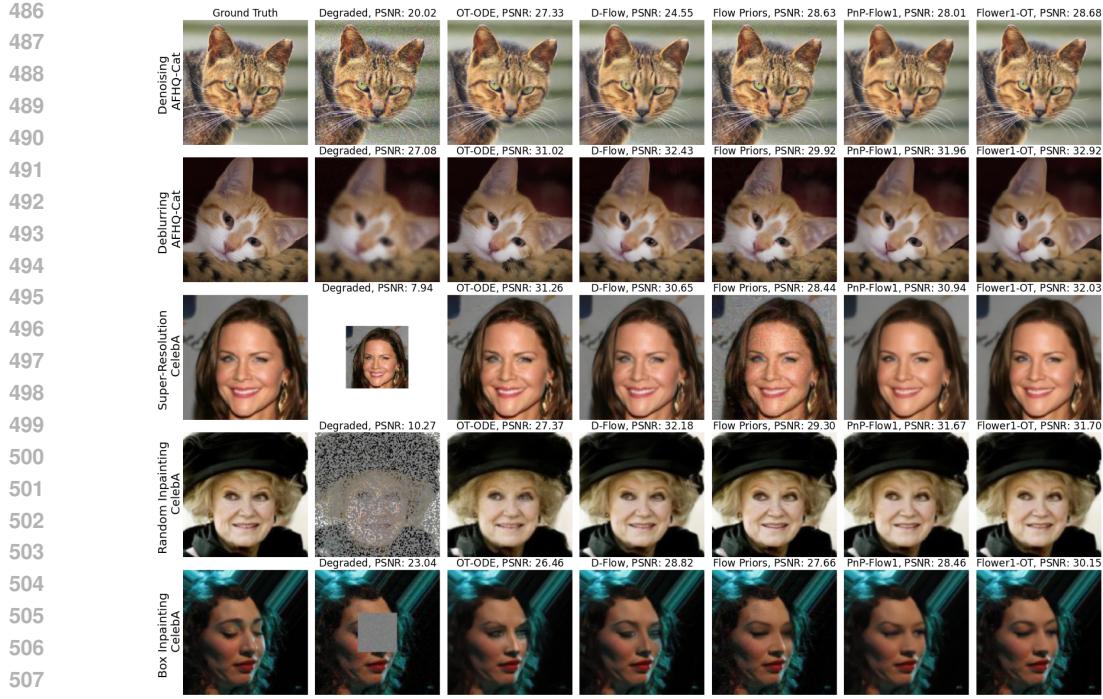
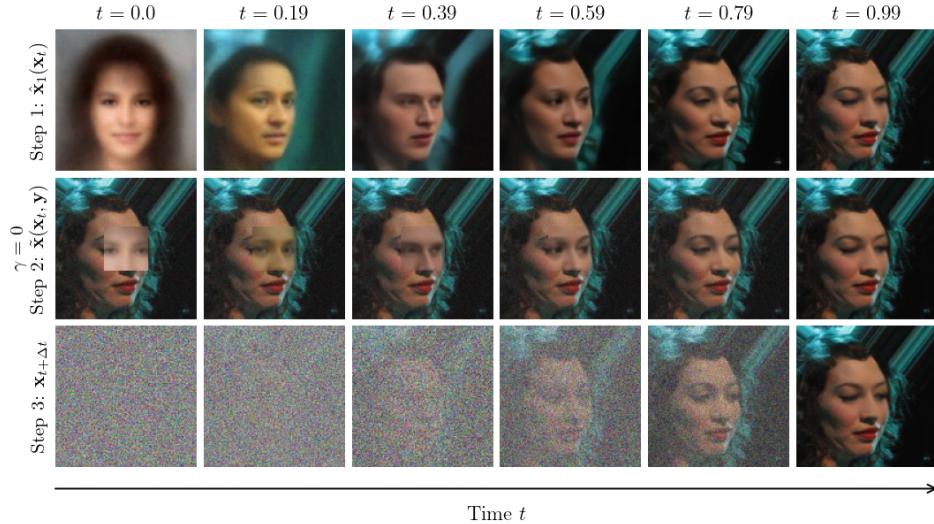


Figure 2: Visual comparison for flow-matching inverse solvers.

Figure 3: Solution path of *Flower* for box inpainting.

interpretation while retaining the plug-and-play flexibility of existing approaches. Our analysis established the conditions under which the method recovers approximate samples from the conditional posterior. Our experiments demonstrated both validity on toy data and state-of-the-art performance across diverse inverse problems.

540 **Ethics Statement.** This work proposes a methodology for solving inverse problems in imaging using  
 541 pre-trained generative models. Our method is designed as a general-purpose solver and does not  
 542 target specific sensitive domains. Our experiments are conducted exclusively on publicly available  
 543 datasets (CelebA and AFHQ-Cat) that are commonly used in the literature. No private or otherwise  
 544 sensitive data were collected or used. Our method has potential positive applications in areas such  
 545 as medical imaging, but, as with other generative techniques, should be applied responsibly to avoid  
 546 misuse in creating misleading content.

547 **Reproducibility Statement.** We aim to ensure that our work is fully reproducible. We provide  
 548 detailed descriptions of the algorithm (Section 3) and its pseudo-code (Algorithm 1), the theoreti-  
 549 cal analysis (Appendix 8.2), and the experimental setup (Section 5 and Appendix 8.5), including  
 550 datasets, hyperparameters, training procedures, and evaluation metrics. All datasets are publicly  
 551 available and cited in the text, and our method relies on standard architectures and benchmarks,  
 552 allowing independent verification of our results. Our implementation will be released in a public  
 553 repository upon the publication of our paper.

554 **Use of LLMs.** The authors of this manuscript acknowledge the use of large language models  
 555 (LLM) for grammatical polishing and typographic corrections.

## 556 REFERENCES

557 Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic  
 558 interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.

559 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the*  
 560 *Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhauser, 2nd edition,  
 561 2008.

562 Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse  
 563 problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

564 Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-Flow:  
 565 Differentiating through flows for controlled generation. In *International Conference on Machine*  
 566 *Learning (ICML)*, 2024.

567 Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*,  
 568 8(3–4):231–357, November 2015. ISSN 1935-8237. doi: 10.1561/2200000050.

569 Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Con-  
 570 ditioning method for denoising diffusion probabilistic models. In *CVF International Conference*  
 571 *on Computer Vision (ICCV)*, volume 1, pp. 2, 2021.

572 Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis  
 573 for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 574 *recognition*, pp. 8188–8197, 2020.

575 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul  
 576 Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference*  
 577 *on Learning Representations*, 2023.

578 Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating  
 579 large-scale inverse problems. In *The Twelfth International Conference on Learning Representa-  
 580 tions*, 2024.

581 Stanislas Ducotterd, Sebastian Neumayer, and Michael Unser. Learning of patch-based smooth-  
 582 plus-sparse models for image reconstruction. In Beidi Chen, Shijia Liu, Mert Pilanci, Weijie Su,  
 583 Jeremias Sulam, Yuxiang Wang, and Zhihui Zhu (eds.), *Conference on Parsimony and Learning*,  
 584 volume 280 of *Proceedings of Machine Learning Research*, pp. 89–104. PMLR, 24–27 Mar 2025.  
 585 URL <https://proceedings.mlr.press/v280/ducotterd25a.html>.

586 M.A.T. Figueiredo and R.D. Nowak. An em algorithm for wavelet-based image restoration. *IEEE*  
 587 *Transactions on Image Processing*, 12(8):906–916, 2003. doi: 10.1109/TIP.2003.814255.

594 Alexis Goujon, Sebastian Neumayer, and Michael Unser. Learning weakly convex regularizers for  
 595 convergent image-reconstruction algorithms. *SIAM Journal on Imaging Sciences*, 17(1):91–115,  
 596 2024.

597

598 Paul Hagemann, Johannes Hertrich, and Gabriele Steidl. Stochastic normalizing flows for inverse  
 599 problems: a Markov Chains viewpoint. *SIAM Journal on Uncertainty Quantification*, 10(3):  
 600 1162–1190, 2022.

601 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint*  
 602 *arxiv:2006.11239*, 2020.

603

604 Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Proximal denoiser for convergent plug-  
 605 and-play optimization with nonconvex regularization. In *International Conference on Machine  
 606 Learning*, volume 162, pp. 9483–9505, 2022a.

607 Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Gradient step denoiser for convergent  
 608 plug-and-play. In *International Conference on Learning Representations*, 2022b.

609

610 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative  
 611 adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition  
 612 (CVPR)*, pp. 4401–4410, 2019.

613 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration  
 614 models. In *Advances in Neural Information Processing Systems*, 2022.

615

616 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow  
 617 matching for generative modeling. In *The Eleventh International Conference on Learning Repre-  
 618 sentations*, 2023.

619 Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint*  
 620 *arXiv:2209.14577*, 2022.

621

622 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and  
 623 transfer data with rectified flow. In *International Conference on Learning Representations (ICLR)*,  
 624 2023.

625 Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving  
 626 inverse problems with diffusion models. In *International Conference on Learning Representa-  
 627 tions (ICLR)*, 2024.

628

629 Ségalène Tiffany Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. Pnp-flow: Plug-  
 630 and-play image restoration with flow matching. In *The Thirteenth International Conference on  
 631 Learning Representations*, 2025.

632 Michael T. McCann and Michael Unser. Biomedical image reconstruction: From the foundations to  
 633 deep neural networks. *Foundations and Trends® in Signal Processing*, 13(3):283–359, 2019.

634

635 Gabriel Peyré. Optimal Transport for Machine Learners, 2025.

636

637 Ashwini Pokle, Matthew J. Muckley, Ricky T. Q. Chen, and Brian Karrer. Training-free linear image  
 638 inverses via flows. *Transactions on Machine Learning Research*, 2024.

639

640 Mehrsa Pourya, Erich Kobler, Michael Unser, and Sebastian Neumayer. DEALing with image re-  
 641 construction: Deep attentive least squares. In *Forty-second International Conference on Machine  
 642 Learning*, 2025.

643

644 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-  
 645 ical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention-  
 646 MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceed-  
 647 ings, Part III 18*, pp. 234–241. Springer, 2015.

648

649 Stefan Roth and Michael J Black. Fields of experts. *International Journal of Computer Vision*, 82  
 650 (2):205–229, 2009.

648 Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal  
 649 algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.  
 650

651 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised  
 652 learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. pmlr, 2015.  
 653

654 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion  
 655 models for inverse problems. In *International Conference on Learning Representations*, 2023.  
 656

657 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 658 Poole. Score-based generative modeling through stochastic differential equations. In *International  
 659 Conference on Learning Representations*, 2020.

660 Julian Tachella, Dongdong Chen, Samuel Hurault, Matthieu Terris, and Andrew Wang. DeepInverse:  
 661 A deep learning framework for inverse problems in imaging, 2023.  
 662

663 Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-  
 664 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models  
 665 with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024.

666 Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for  
 667 model based reconstruction. In *IEEE Global Conference on Signal and Information Processing*,  
 668 pp. 945–948, 2013.

669 Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. From facial parts responses to face  
 670 detection: A deep learning approach. In *Proceedings of the IEEE international conference on  
 671 computer vision*, pp. 3676–3684, 2015.  
 672

673 Gengsheng Lawrence Zeng. Image reconstruction—A tutorial. *Computerized Medical Imaging and  
 674 Graphics*, 25(2):97–103, 2001.

675 Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song.  
 676 Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings  
 677 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20895–  
 678 20905, June 2025.  
 679

680 Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play  
 681 image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine  
 682 Intelligence*, 44(10):6360–6376, 2022.

683 Yasi Zhang, Peiyu Yu, Yaxuan Zhu, Yingshan Chang, Feng Gao, Ying Nian Wu, and Oscar Leong.  
 684 Flow priors for linear inverse problems via iterative corrupted trajectory matching. In *The Thirty-  
 685 eighth Annual Conference on Neural Information Processing Systems*, 2024.

686 Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc  
 687 Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the  
 688 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1219–1229, 2023.  
 689

690

691

692

693

694

695

696

697

698

699

700

701

702 8 APPENDIX  
703704 8.1 FLOWER ALGORITHM  
705706 We outline the steps of *Flower* in Algorithm 1.  
707708 **Algorithm 1** FLOWER: Flow Matching Solver for Inverse Problems

---

709 **Require:** Measurements  $\mathbf{y}$ , forward operator  $\mathbf{H}$ , noise level  $\sigma_n$ , pretrained velocity  $\mathbf{v}_t^\theta$ , steps  $N$ ,  
710 uncertainty flag  $\gamma \in \{0, 1\}$

711 1: Set  $\Delta t = 1/N$ ; sample  $\mathbf{x}_0 \sim p_{\mathbf{X}_0}$   $\triangleright$  e.g.,  $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$

712 2: **for**  $k = 0$  to  $N - 1$  **do**

713 3:    $t = k \Delta t$

714 4:   **(Step 1) Destination estimate:**  $\hat{\mathbf{x}}_1(\mathbf{x}_t) = \mathbf{x}_t + (1 - t) \mathbf{v}_t^\theta(\mathbf{x}_t)$

715 5:    $\nu_t = \frac{1-t}{\sqrt{t^2+(1-t)^2}}, \quad F_{\mathbf{y}}(\mathbf{x}) = \frac{1}{2\sigma_n^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2$

716 6:   **(Step 2) Refinement mean:**  $\mu_t = \text{prox}_{\nu_t^2 F_{\mathbf{y}}}(\hat{\mathbf{x}}_1(\mathbf{x}_t))$

717 7:   **(Step 2) Optional uncertainty:**  $\Sigma_t = (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}$

718 8:   sample  $\epsilon_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ ,  $\epsilon_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_M)$ ,  $\kappa_t = \Sigma_t(\nu_t^{-1} \epsilon_1 + \sigma_n^{-1} \mathbf{H}^\top \epsilon_2)$

719 9:    $\tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y}) = \mu_t + \gamma \kappa_t$

720 10:   **(Step 3) Time progression:** sample  $\epsilon \sim p_{\mathbf{X}_0}$  and set

721 11:    $\mathbf{x}_{t+\Delta t} = (1 - t - \Delta t) \epsilon + (t + \Delta t) \tilde{\mathbf{x}}_1(\mathbf{x}_t, \mathbf{y})$

722 12: **end for**

723 13: **return**  $\mathbf{x}_1$

---

725 8.2 PROOFS  
726727 8.2.1 PROOF OF THEOREM 1  
729730 To establish this result, we first recall ancestral sampling. It is a procedure that enables us to draw  
731 samples from a marginal distribution  $p_{\mathbf{Z}_K}$  when the full joint distribution over a sequence of vari-  
732 ables  $\mathbf{Z} = (\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_K)$  is defined via a chain of conditional densities and when the marginal  
733 distribution of  $\mathbf{Z}_K$  can be written as

734 
$$p_{\mathbf{Z}_K}(\mathbf{z}_K) = \int_{\mathbf{z}_{K-1}} \cdots \int_{\mathbf{z}_1} p_{\mathbf{Z}_K | \mathbf{Z}_{K-1} = \mathbf{z}_{K-1}}(\mathbf{z}_t) \cdots p_{\mathbf{Z}_2 | \mathbf{Z}_1 = \mathbf{z}_1}(\mathbf{z}_2) p_{\mathbf{Z}_1}(\mathbf{z}_1) d\mathbf{z}_1 \cdots d\mathbf{z}_{K-1}. \quad (24)$$
  
735

736 Ancestral sampling offers a practical way to generate samples from  $p_{\mathbf{Z}_K}$  without direct evaluation  
737 of this integral. The process samples sequentially from the distributions

737 
$$\mathbf{z}_1 \sim p_{\mathbf{Z}_1}, \quad \mathbf{z}_2 \sim p_{\mathbf{Z}_2 | \mathbf{Z}_1 = \mathbf{z}_1}, \quad \dots, \quad \mathbf{z}_K \sim p_{\mathbf{Z}_K | \mathbf{Z}_{K-1} = \mathbf{z}_{K-1}}. \quad (25)$$

738 By following this sequence, we obtain a valid sample from the marginal  $\mathbf{z}_K \sim p_{\mathbf{Z}_K}$ .742 *Proof of Theorem 1.* By using the marginal distributions and the general chain rule for joint proba-  
743 bility, we obtain

744 
$$\begin{aligned} p_{\mathbf{X}_{t+\Delta t} | \mathbf{Y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t}) &= \int_{\mathbf{x}_t} p_{\mathbf{X}_t, \mathbf{X}_{t+\Delta t} | \mathbf{Y} = \mathbf{y}}(\mathbf{x}_t, \mathbf{x}_{t+\Delta t}) d\mathbf{x}_t. \\ &= \int_{\mathbf{x}_t} p_{\mathbf{X}_t | \mathbf{Y} = \mathbf{y}}(\mathbf{x}_t) p_{\mathbf{X}_{t+\Delta t} | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t}) d\mathbf{x}_t. \end{aligned} \quad (26)$$
  
745

749 It then suffices to follow the ancestral sampling procedure with  $K = 2$ ,  $\mathbf{Z}_1 = \mathbf{X}_t | \mathbf{Y} = \mathbf{y}$ , and  
750  $\mathbf{Z}_2 = \mathbf{X}_{t+\Delta t} | \mathbf{Y} = \mathbf{y}$  to complete the proof.  $\square$   
751752 8.2.2 PROOF OF PROPOSITION 1  
753754 *Proof.* Lipman et al. (2023) have shown that the velocity vector field that minimizes the conditional  
755 flow-matching loss is

755 
$$\mathbf{v}_t^*(\mathbf{x}) = \mathbb{E}[\mathbf{X}_1 - \mathbf{X}_0 | \mathbf{X}_t = \mathbf{x}_t], \quad (27)$$

756 which then yields  
 757

$$\hat{\mathbf{x}}_1(\mathbf{x}_t) = \mathbb{E}[\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t] = \mathbb{E}[\mathbf{X}_t + (1-t)(\mathbf{X}_1 - \mathbf{X}_0) | \mathbf{X}_t = \mathbf{x}_t] = \mathbf{x}_t + (1-t)\mathbf{v}_t^\theta(\mathbf{x}). \quad (28)$$

759 This is the desired result under the assumption that  $\mathbf{v}_t^\theta(\mathbf{x}) = \mathbf{v}_t^*(\mathbf{x})$ .  $\square$   
 760

### 761 8.2.3 PROOF OF PROPOSITION 2

763 *Proof.* Using Bayes' rule, we can write

$$764 \quad p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) = \frac{p_{\mathbf{Y} | \mathbf{X}_1 = \mathbf{x}_1, \mathbf{X}_t = \mathbf{x}_t}(\mathbf{y}) p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{x}_1)}{p_{\mathbf{Y} | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{y})} = \frac{p_{\mathbf{Y} | \mathbf{X}_1 = \mathbf{x}_1}(\mathbf{y}) p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{x}_1)}{p_{\mathbf{Y} | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{y})}, \quad (29)$$

767 where we used the conditional independence, given  $\mathbf{X}_1$ , of  $\mathbf{X}_t$  and the measurement  $\mathbf{Y}$ . We assumed  
 768 that  $p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t}$  is approximated with  $\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t} = \mathcal{N}(\hat{\mathbf{x}}_1(\mathbf{x}_t), \nu_t^2 \mathbf{I}_d)$ , and we have  $p_{\mathbf{Y} | \mathbf{X}_1 = \mathbf{x}_1} =$   
 769  $\mathcal{N}(\mathbf{H}\mathbf{x}_1, \sigma_n^2 \mathbf{I})$  by construction. Put together, we obtain the approximate density  
 770

$$771 \quad \tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) = \frac{p_{\mathbf{Y} | \mathbf{X}_1 = \mathbf{x}_1}(\mathbf{y}) \tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{x}_1)}{p_{\mathbf{Y} | \mathbf{X}_t = \mathbf{x}_t}(\mathbf{y})}. \quad (30)$$

773 Taking the logarithm of equation 30,  $-2 \log(\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1))$  is given by  
 774

$$775 \quad (\mathbf{y} - \mathbf{H}\mathbf{x}_1)^\top \sigma_n^{-2} (\mathbf{y} - \mathbf{H}\mathbf{x}_1) + (\mathbf{x}_1 - \hat{\mathbf{x}}_1(\mathbf{x}_t))^\top \nu_t^{-2} (\mathbf{x}_1 - \hat{\mathbf{x}}_1(\mathbf{x}_t)) + C \quad (31)$$

$$776 \quad = -2\mathbf{x}_1^\top (\nu_t^{-2} \hat{\mathbf{x}}_1(\mathbf{x}_t) + \sigma_n^{-2} \mathbf{H}^\top \mathbf{y}) + \mathbf{x}_1^\top (\nu_t^{-2} \mathbf{I} + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}) \mathbf{x}_1 + C' \quad (32)$$

$$778 \quad = -2\mathbf{x}_1^\top (\nu_t^{-2} \hat{\mathbf{x}}_1(\mathbf{x}_t) + \sigma_n^{-2} \mathbf{H}^\top \mathbf{y}) + \mathbf{x}_1^\top \Sigma_t^{-1} \mathbf{x}_1 + C', \quad (33)$$

779 where  $C, C'$  are independent of  $\mathbf{x}_1$  and considered constants and  $\Sigma_t = (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}$ ,  
 780 which is well-defined because  $\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}$  is positive-definite. Completing the square with  
 781 a term independent of  $\mathbf{x}_1$ , we get  
 782

$$783 \quad -2 \log(\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1)) = (\mathbf{x}_1 - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1} (\mathbf{x}_1 - \boldsymbol{\mu}_t) + C'', \quad (34)$$

785 where  $\boldsymbol{\mu}_t = \Sigma_t (\nu_t^{-2} \hat{\mathbf{x}}_1(\mathbf{x}_t) + \sigma_n^{-2} \mathbf{H}^\top \mathbf{y})$  and  $C''$  is again a constant independent of  $\mathbf{x}_1$ . This yields  
 786

$$787 \quad \tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) = e^{-C''/2} \exp\left(-\frac{1}{2}(\mathbf{x}_1 - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1} (\mathbf{x}_1 - \boldsymbol{\mu}_t)\right). \quad (35)$$

788 As  $\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}$  is a probability density function,  $e^{-C''/2}$  corresponds to its normalization factor,  
 789 which therefore proves that  $\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) = \mathcal{N}(\mathbf{x}_1; \boldsymbol{\mu}_t, \Sigma_t)$ .  $\square$   
 790

### 792 8.2.4 PROOF OF PROPOSITION 3

793 *Proof.* By using the marginal distributions, the general chain rule for joint probability, and indepen-  
 794 dence, we obtain the expression of  $p_{\mathbf{X}_{t+\Delta t} | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t})$  as  
 795

$$796 \quad \int_{\mathbb{R}^d} p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) p_{\mathbf{X}_{t+\Delta t} | \mathbf{X}_1 = \mathbf{x}_1, \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t}) d\mathbf{x}_1 \quad (36)$$

$$798 \quad = \int_{\mathbb{R}^d} p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) p_{\mathbf{X}_{t+\Delta t} | \mathbf{X}_1 = \mathbf{x}_1, \mathbf{X}_t = \mathbf{x}_t}(\mathbf{x}_{t+\Delta t}) d\mathbf{x}_1 \quad (37)$$

$$800 \quad = \int_{\mathbb{R}^d} p_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1) \mathcal{N}(\mathbf{x}_{t+\Delta t}; (t + \Delta t)\mathbf{x}_1, (1 - t - \Delta t)^2 \mathbf{I}_d) d\mathbf{x}_1, \quad (38)$$

803 where we used the fact that, conditioned on  $\mathbf{X}_1 = \mathbf{x}_1$ ,  $\mathbf{X}_{t+\Delta t} = (1 - t - \Delta t)\mathbf{X}_0 + (t +$   
 804  $\Delta t)\mathbf{x}_1 \sim \mathcal{N}((t + \Delta t)\mathbf{x}_1, (1 - t - \Delta t)^2 \mathbf{I}_d)$ . By inserting the expression of the approximation  
 805  $\tilde{p}_{\mathbf{X}_1 | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_1)$ , the approximate density  $\tilde{p}_{\mathbf{X}_{t+\Delta t} | \mathbf{X}_t = \mathbf{x}_t, \mathbf{Y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t})$  is given by

$$806 \quad \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{x}_1; \boldsymbol{\mu}_t, \Sigma_t) \mathcal{N}(\mathbf{x}_{t+\Delta t}; (t + \Delta t)\mathbf{x}_1, (1 - t - \Delta t)^2 \mathbf{I}_d) d\mathbf{x}_1 \quad (39)$$

$$808 \quad = \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{x}_1; \boldsymbol{\mu}_t, \Sigma_t) \mathcal{N}(\mathbf{x}_{t+\Delta t} - (t + \Delta t)\mathbf{x}_1; \mathbf{0}, (1 - t - \Delta t)^2 \mathbf{I}_d) d\mathbf{x}_1. \quad (40)$$

This integral can be rewritten as a convolution of two Gaussian distributions, which also yields a Gaussian distribution. Explicitly, we have that

$$\int_{\mathbb{R}^d} \mathcal{N}(\mathbf{x}_1; \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t) \mathcal{N}(\mathbf{x}_{t+\Delta t} - (t + \Delta t)\mathbf{x}_1; \mathbf{0}, (1 - t - \Delta t)^2 \mathbf{I}_d) d\mathbf{x}_1 \quad (41)$$

$$= \int_{\mathbb{R}^d} \mathcal{N}\left(\frac{\mathbf{z}}{t + \Delta t}; \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t\right) \mathcal{N}\left(\mathbf{x}_{t+\Delta t} - \mathbf{z}; \mathbf{0}, (1 - t - \Delta t)^2 \mathbf{I}_d\right) \frac{d\mathbf{z}}{(t + \Delta t)^d} \quad (42)$$

$$= \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{z}; (t + \Delta t)\boldsymbol{\mu}_t, (t + \Delta t)^2 \boldsymbol{\Sigma}_t) \mathcal{N}(\mathbf{x}_{t+\Delta t} - \mathbf{z}; \mathbf{0}, (1 - t - \Delta t)^2 \mathbf{I}_d) d\mathbf{z}. \quad (43)$$

Using the Gaussian convolution identity

$$\int_{\mathbb{R}^d} \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \mathcal{N}(\mathbf{x} - \mathbf{z}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) d\mathbf{z} = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_1 + \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2), \quad (44)$$

the result simplifies to

$$\tilde{p}_{\mathbf{x}_{t+\Delta t} | \mathbf{x}_t = \mathbf{x}_t, \mathbf{y} = \mathbf{y}}(\mathbf{x}_{t+\Delta t}) = \mathcal{N}(\mathbf{x}_{t+\Delta t}; (t + \Delta t)\boldsymbol{\mu}_t, (t + \Delta t)^2 \boldsymbol{\Sigma}_t + (1 - t - \Delta t)^2 \mathbf{I}_d), \quad (45)$$

which completes the proof.  $\square$

### 8.2.5 SAMPLING FROM THE NON-ISOTROPIC GAUSSIAN

We want to show that

$$\boldsymbol{\kappa}_t = \boldsymbol{\Sigma}_t (\nu_t^{-1} \boldsymbol{\epsilon}_1 + \sigma_n^{-1} \mathbf{H}^\top \boldsymbol{\epsilon}_2) \quad (46)$$

has distribution  $\mathcal{N}(0, \boldsymbol{\Sigma}_t)$ , where  $\boldsymbol{\epsilon}_1 \sim \mathcal{N}(0, \mathbf{I}_d)$  and  $\boldsymbol{\epsilon}_2 \sim \mathcal{N}(0, \mathbf{I}_M)$  are independent, and

$$\boldsymbol{\Sigma}_t = (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}. \quad (47)$$

Observe that  $\boldsymbol{\kappa}_t$  is a Gaussian random vector because it is a linear transform of the independent Gaussians  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$ . Moreover, since  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$  are zero-mean, we have that  $\mathbb{E}[\boldsymbol{\kappa}_t] = \mathbf{0}$ . Next, we compute the covariance matrix of  $\boldsymbol{\kappa}_t$  as

$$\text{Cov}(\boldsymbol{\kappa}_t) = \boldsymbol{\Sigma}_t \text{Cov}(\nu_t^{-1} \boldsymbol{\epsilon}_1 + \sigma_n^{-1} \mathbf{H}^\top \boldsymbol{\epsilon}_2) \boldsymbol{\Sigma}_t. \quad (48)$$

By independence of  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$ , we obtain that

$$\text{Cov}(\nu_t^{-1} \boldsymbol{\epsilon}_1 + \sigma_n^{-1} \mathbf{H}^\top \boldsymbol{\epsilon}_2) = \nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}, \quad (49)$$

which implies that

$$\text{Cov}(\boldsymbol{\kappa}_t) = \boldsymbol{\Sigma}_t (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}) \boldsymbol{\Sigma}_t. \quad (50)$$

But, by definition,  $\boldsymbol{\Sigma}_t = (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H})^{-1}$ , which leads to the desired result, as

$$\text{Cov}(\boldsymbol{\kappa}_t) = \boldsymbol{\Sigma}_t (\nu_t^{-2} \mathbf{I}_d + \sigma_n^{-2} \mathbf{H}^\top \mathbf{H}) \boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_t. \quad (51)$$

### 8.3 EXTENSION TO MORE GENERAL NOISE TYPES

Throughout this manuscript, isotropic Gaussian noise  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$  was considered. Our framework remains valid in the more general setting  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_n)$ , where  $\mathbf{R}_n$  is a symmetric positive definite  $M \times M$  covariance matrix, up to the modification of some formulas detailed below.

The results of Proposition 2 become

$$\boldsymbol{\mu}_t(\mathbf{x}_t, \mathbf{y}) = (\nu_t^{-2} \mathbf{I}_d + \mathbf{H}^\top \mathbf{R}_n^{-1} \mathbf{H})^{-1} (\nu_t^{-2} \hat{\mathbf{x}}_1(\mathbf{x}_t) + \mathbf{H}^\top \mathbf{R}_n^{-1} \mathbf{y}), \quad (52)$$

$$\boldsymbol{\Sigma}_t = (\nu_t^{-2} \mathbf{I}_d + \mathbf{H}^\top \mathbf{R}_n^{-1} \mathbf{H})^{-1}, \quad (53)$$

since the measurement operation with general noise type implies that  $p_{\mathbf{y} | \mathbf{x}_1 = \mathbf{x}_1} = \mathcal{N}(\mathbf{H}\mathbf{x}_1, \mathbf{R}_n)$ .

Next, sampling from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t)$  is achieved with a method similar to the one of Appendix 8.2.5, with the updated formula

$$\boldsymbol{\kappa}_t = \boldsymbol{\Sigma}_t \left( \nu_t^{-1} \boldsymbol{\epsilon}_1 + \mathbf{H}^\top \mathbf{R}_n^{-\frac{1}{2}} \boldsymbol{\epsilon}_2 \right). \quad (54)$$

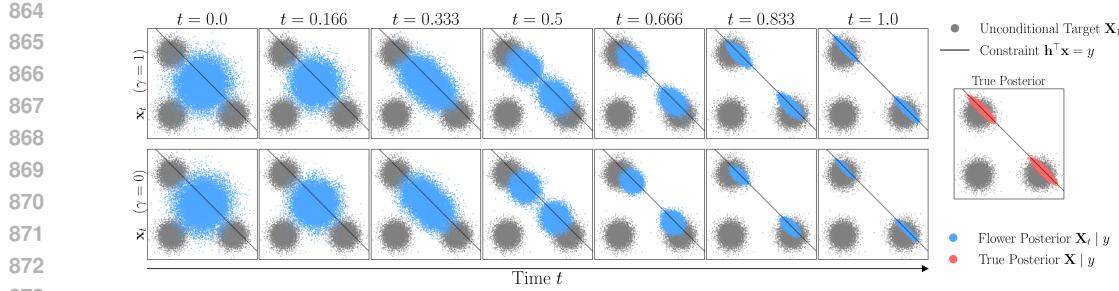


Figure 4: Temporal evolution of 2D *Flower* and comparison with true posterior for noise variance  $\sigma_n = 0.25$ .

This formula uses the standard definition of the square root of a symmetric positive-definite matrix  $\mathbf{A}$ . If  $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^\top$  is its eigenvalue decomposition, then  $\mathbf{A}^{\frac{1}{2}} = \mathbf{P}\mathbf{D}^{\frac{1}{2}}\mathbf{P}^\top$ , where  $\mathbf{D}^{\frac{1}{2}}$  is the diagonal matrix whose entries are the square roots of the nonnegative eigenvalues of  $\mathbf{A}$ .

As a result, in the measurement-aware destination refinement step of *Flower* described in Section 3, equation 12 now computes  $\text{prox}_{\nu_t^2 F_y}$  with

$$F_y(\mathbf{x}) = \frac{1}{2}(\mathbf{y} - \mathbf{H}\mathbf{x})^\top \mathbf{R}_n^{-1}(\mathbf{y} - \mathbf{H}\mathbf{x}). \quad (55)$$

In the same step, the matrix  $\Sigma_t$  is redefined to be as in equation 53.

In most imaging models, the noise is effectively isotropic. Moreover, if we assume that  $\mathbf{R}_n$  is diagonal, the resulting expressions simplify and become straightforward to compute.

## 8.4 NUMERICAL RESULTS EXTENSION

### 8.4.1 TOY EXPERIMENT

The goal of this experiment is to validate our proposed sampling perspective in a setting where samples from the true posterior are available. To this end, we consider a Gaussian mixture model (GMM) as the target distribution of the data  $\mathbf{X} \in \mathbb{R}^2$ . The forward measurement model consists of a single measurement vector  $\mathbf{h} \in \mathbb{R}^d$  (i.e., the forward operator is  $\mathbf{H} = \mathbf{h}^\top$ ) corrupted by additive white Gaussian noise  $n \sim \mathcal{N}(0, \sigma_n^2)$ , which results in  $y = \mathbf{h}^\top \mathbf{x} + n$ . In this setup, the posterior distribution  $p_{\mathbf{X}|\mathbf{Y}=y}$  is itself a GMM with known parameters, whose analytical expression is given in equation 57. Geometrically, the noiseless measurement  $y = \mathbf{h}^\top \mathbf{x}$  defines a line in the two-dimensional plane with normal vector  $\mathbf{h}$ . Consequently, when sampling from the posterior  $p_{\mathbf{X}|\mathbf{Y}=y}$ , we expect the samples to concentrate on the portions of this line that intersect regions where the prior distribution  $p_{\mathbf{X}}$  has high density.

We trained the unconditional flow-matching vector field using the source  $p_{\mathbf{X}_0} = \mathcal{N}(\mathbf{0}, \mathbf{I})$  and target  $p_{\mathbf{X}_1} = p_{\mathbf{X}}$ , with further details provided below. We ran *Flower* with  $N = 1000$  iterations, which corresponds to the step size  $\Delta t = 0.001$ . The results are shown in Figure 4, where we used  $\mathbf{h}^\top = [1.5, 1.5]$ ,  $\sigma_n = 0.25$ , and the observation  $y = 1$ , and where we report the solution paths of *Flower* with  $\gamma \in \{0, 1\}$  alongside true posterior samples. We observe, that when  $\gamma = 1$  (the setting required by our theory), *Flower* successfully recovers the samples at  $t = 1$ , which closely resemble the true posterior. When  $\gamma = 0$  (a configuration in which the uncertainty in the destination estimation step is ignored), *Flower* fails to capture samples from the tails of the distribution. In Figure 5, we present another example of *Flower* posterior sampling with  $\mathbf{h}^\top = [1.5, -1.5]$ ,  $\sigma_n = 0.75$ , and the observation  $y = 1$ . Once again, we observe that *Flower* successfully generates samples that closely match the true posterior for  $\gamma = 1$ . In contrast, for  $\gamma = 0$ , samples from the tails of the true posterior are missing. In Figures 6 and 7, the solution path is illustrated across the successive steps of *Flower* for  $\gamma = 0$  and  $\gamma = 1$ , respectively, which allows us to visualize the dynamics of each step directly.

918    **Target Prior.** The target distribution of our data  $\mathbf{X} \in \mathbb{R}^2$  is a GMM with uniform mixtures given  
 919    explicitly by  
 920

$$p_{\mathbf{X}} = \frac{1}{K} \sum_{k=1}^K \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}) \quad (56)$$

923    for some  $K \in \mathbb{N}$ ,  $\boldsymbol{\mu}_k \in \mathbb{R}^2$ , and  $\boldsymbol{\Sigma} \in \mathbb{S}_{++}^2$ . Specifically, we use  $K = 3$  with  $\boldsymbol{\mu}_1 =$   
 924     $(-0.25, -0.25)$ ,  $\boldsymbol{\mu}_2 = (-0.25, 0.25)$ ,  $\boldsymbol{\mu}_3 = (0.25, -0.25)$ , and covariance matrix  $\boldsymbol{\Sigma} = 0.25^2 \mathbf{I}_2$ .  
 925

926    **Target Posterior.** The advantage of this setup is that the posterior distribution  $p_{\mathbf{X}|\mathbf{Y}=\mathbf{y}}$  can be  
 927    computed exactly using Bayes' rule. It is given by  
 928

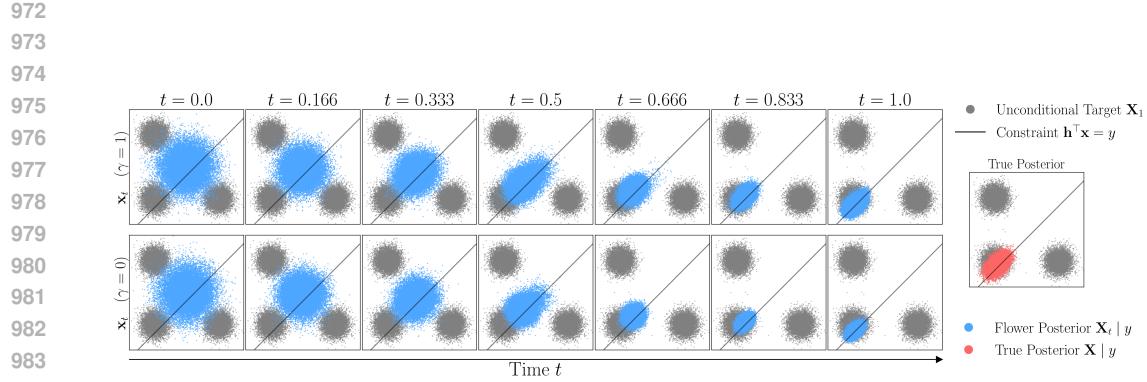
$$p_{\mathbf{X}|\mathbf{Y}=\mathbf{y}} = \sum_{k=1}^K w_k \mathcal{N}(\boldsymbol{\mu}_{k,\text{post}}, \boldsymbol{\Sigma}_{\text{post}}) \quad (57)$$

932    where, for all  $k = 1, \dots, K$ ,  $w_k \geq 0$  are some weights and  
 933

$$\boldsymbol{\mu}_{k,\text{post}} = (\boldsymbol{\Sigma}^{-1} + \sigma_n^{-2} \mathbf{h} \mathbf{h}^\top)^{-1} (\sigma_n^{-2} \mathbf{h} \mathbf{y} + \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k), \quad (58)$$

$$\boldsymbol{\Sigma}_{\text{post}} = (\boldsymbol{\Sigma}^{-1} + \sigma_n^{-2} \mathbf{h} \mathbf{h}^\top)^{-1}. \quad (59)$$

937    **Training Details.** The underlying unconditional velocity network is a fully connected network  
 938    that takes as input a 2D vector and a scalar time, concatenated into a 3D input. It consists of two  
 939    hidden layers of size 256 with SiLU activations, followed by a final linear layer that outputs a 2D  
 940    vector. For training, we use a batch size of 2048 with 20000 training steps and a learning rate of  
 941     $10^{-3}$ .  
 942  
 943  
 944  
 945  
 946  
 947  
 948  
 949  
 950  
 951  
 952  
 953  
 954  
 955  
 956  
 957  
 958  
 959  
 960  
 961  
 962  
 963  
 964  
 965  
 966  
 967  
 968  
 969  
 970  
 971



985  
986 Figure 5: Temporal evolution of 2D *Flower* and comparison with true posterior for noise variance  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

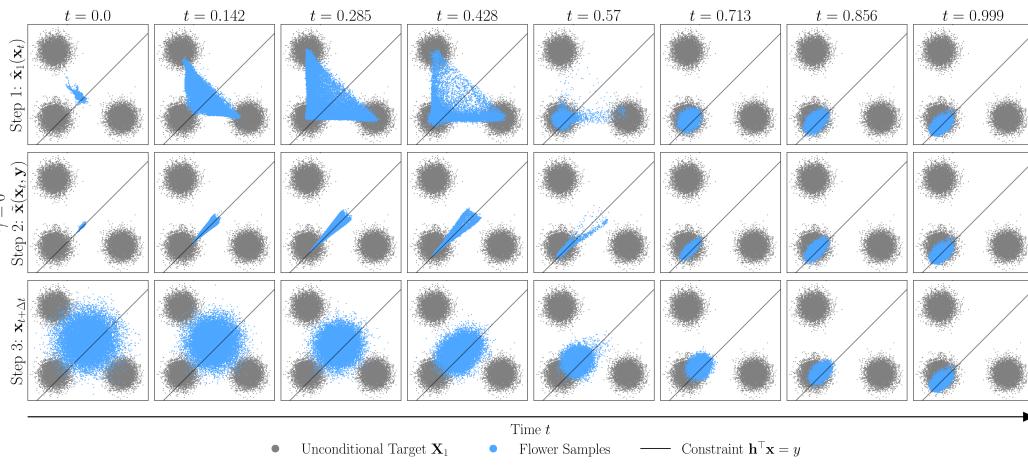


Figure 6: The three steps of 2D *Flower* with temporal evolution for  $\gamma = 0$ .

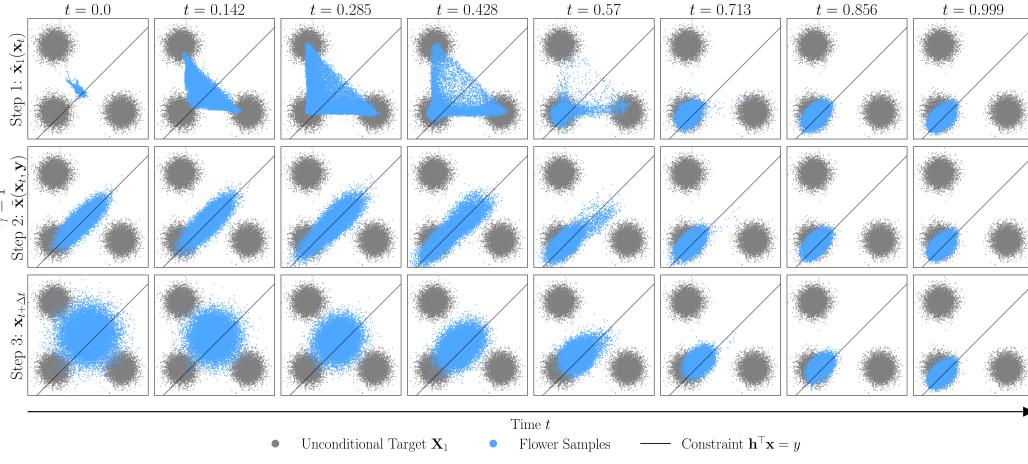


Figure 7: The three steps of 2D *Flower* with temporal evolution for  $\gamma = 1$ .

1026 8.5 BENCHMARK EXPERIMENTS  
10271028 8.5.1 COMPUTATIONAL EFFICIENCY  
10291030 We report in Table 3 the computational time and memory usage for several methods. Each entry  
1031 corresponds to the average runtime for the deblurring inverse problem, averaged over 10 CelebA  
1032 test images of size  $128 \times 128$ . All experiments were conducted on a Tesla V100-SXM2-32GB GPU.  
10331034 Table 3: Computation times and memory usage for various methods.  
1035

| Method      | OT-ODE | D-Flow | Flow Priors | PnP-Flow1 | Flower1 |
|-------------|--------|--------|-------------|-----------|---------|
| Time (s)    | 6.549  | 142.18 | 63.771      | 3.020     | 5.622   |
| Memory (GB) | 1.183  | 11.125 | 3.807       | 0.216     | 0.217   |

1040 8.5.2 EFFECT OF SOURCE-TARGET COUPLING  
10411042 Our goal in this section is to benchmark the effect of source–target coupling in the training of the  
1043 underlying (unconditional) velocity network. To this end, we consider two variants of coupling: one  
1044 based on mini-batch optimal transport (OT) coupling and the other on independent (IND) coupling.  
1045 For each variant, we report two settings: a single evaluation of *Flower*; and average over five eval-  
1046 uations. This results in four cases, summarized in Table 4, where we provide results for all inverse  
1047 problems discussed in the main text, evaluated on 100 test images from the CelebA dataset. For this  
1048 table, we fix  $\gamma = 0$  and  $N = 100$ .  
10491050 Table 4: Effect of source-target coupling of the underlying velocity network of *Flower* on different  
1051 inverse problems on 100 test images of the dataset CelebA.  
1052

| Method             | Denoising |       |       | Deblurring |       |       | Super-resolution |       |       | Random inpainting |       |       | Box inpainting |       |       |
|--------------------|-----------|-------|-------|------------|-------|-------|------------------|-------|-------|-------------------|-------|-------|----------------|-------|-------|
|                    | PSNR      | SSIM  | LPIPS | PSNR       | SSIM  | LPIPS | PSNR             | SSIM  | LPIPS | PSNR              | SSIM  | LPIPS | PSNR           | SSIM  | LPIPS |
| Degraded           | 20.00     | 0.348 | 0.372 | 27.83      | 0.740 | 0.126 | 10.26            | 0.183 | 0.827 | 11.95             | 0.196 | 1.041 | 22.27          | 0.742 | 0.214 |
| Flower1-OT (ours)  | 32.28     | 0.914 | 0.034 | 34.98      | 0.947 | 0.026 | 32.36            | 0.923 | 0.034 | 33.08             | 0.944 | 0.018 | 31.19          | 0.945 | 0.022 |
| Flower1-OT (ours)  | 33.14     | 0.926 | 0.038 | 35.67      | 0.954 | 0.032 | 33.09            | 0.932 | 0.040 | 33.95             | 0.953 | 0.020 | 31.87          | 0.952 | 0.023 |
| Flower1-IND (ours) | 32.60     | 0.918 | 0.032 | 35.22      | 0.950 | 0.026 | 32.65            | 0.927 | 0.034 | 33.23             | 0.947 | 0.017 | 31.90          | 0.950 | 0.021 |
| Flower5-IND (ours) | 33.48     | 0.930 | 0.037 | 35.90      | 0.957 | 0.031 | 33.41            | 0.935 | 0.039 | 34.24             | 0.955 | 0.020 | 32.78          | 0.958 | 0.022 |

1053 We observe that independent coupling improves the results, which is consistent with our theoretical  
1054 requirements. Nevertheless, the OT-based variant remains highly competitive, as also illustrated in  
1055 the main paper. For visual comparison, in Figure 8, we show an example from the deblurring task  
1056 on an image from the CelebA dataset.  
10571058 8.5.3 EFFECT OF  $\gamma$   
10591060 The hyperparameter  $\gamma \in \{0, 1\}$  in *Flower* controls whether the uncertainty of the refinement step  
1061 (Step 2) is taken into account. While  $\gamma = 1$  is required for a Bayesian interpretation of our method,  
1062 in practice we find that  $\gamma = 0$  yields more favorable image-reconstruction metrics. This observation  
1063 is consistent with the toy experiments, where  $\gamma = 0$  led *Flower* to generate samples concentrated in  
1064 higher-probability regions, while failing to capture the tails of the posterior.  
10651066 To illustrate this effect, we provide a visual example of the deblurring task on a CelebA image, using  
1067 the velocity network trained with mini-batch OT coupling for consistency with the numerical results  
1068 reported in this paper. A single reconstruction with  $\gamma = 1$  achieves a PSNR of only 30.84, compared  
1069 to 33.01 for a single reconstruction with  $\gamma = 0$  (Figure 8). Moreover, with  $\gamma = 1$ , it is necessary to  
1070 average over 100 reconstructions to reach a PSNR comparable to that of  $\gamma = 0$ , where we average  
1071 only over 5 reconstructions (Figure 8).  
10721073 To further illustrate this behavior, we show in Figures 6 and 7 the solution paths of the three *Flower*  
1074 steps over time for  $\gamma = 0$  and  $\gamma = 1$ , respectively. We observe that  $\gamma = 1$  leads to a noisier refine-  
1075 ment step. Consequently, we adopt  $\gamma = 0$  for all inverse problems, as this provides a more stable  
1076 refinement step and consistently yields a better quality of reconstruction despite fewer averages.  
1077



1134  
1135

## 8.5.4 EFFECT OF THE NUMBER OF EVALUATIONS FOR AVERAGING

1136  
1137  
1138  
1139  
1140  
1141  
1142

In Tables 1 and 2 of the main paper, we reported results for Flower-1 and Flower-5, which correspond to using a single evaluation and the average over five evaluations of *Flower*, respectively. Here, we ablate the effect of this averaging on both the reconstruction metrics and the computational cost. We focus on the deblurring and random inpainting tasks described in Section 5.1 for the CelebA dataset. Specifically, we vary the number of *Flower* averaging  $N_{\text{Avg}}$  from 1 to 10 and report the average PSNR, SSIM, LPIPS, and runtime (in seconds) over 100 CelebA images, measured on a Tesla V100-SXM2-32GB GPU. For the hyperparameters of *Flower*, we use  $\gamma = 0$ ,  $N = 100$ .

1143  
1144  
1145  
1146  
1147  
1148  
1149

As shown in Figure 12, PSNR and SSIM improve with a steeper gain for smaller numbers of averagings (up to around five) and tend to saturate afterward. LPIPS also increases with averaging, which is undesirable since higher LPIPS indicates worse perceptual quality. The runtime scales linearly with the number of averagings, as expected. Finally, in Figure 13, we present a visual comparison that confirms these trends: averaged reconstructions (e.g., using 5 or 10 evaluations) appear smoother than the single-evaluation result, while the difference between averaging over 5 and 10 evaluations is marginal.

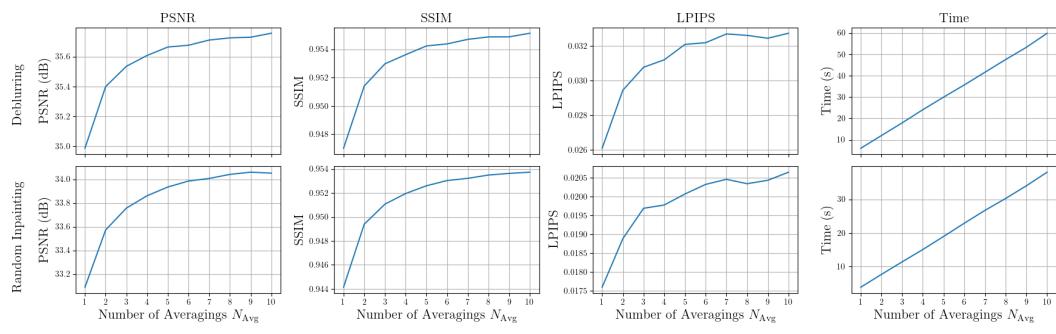
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
11601161  
1162  
1163  
1164

Figure 12: Effect of averaging over different numbers of evaluations of *Flower* on quantitative metrics and computational time.

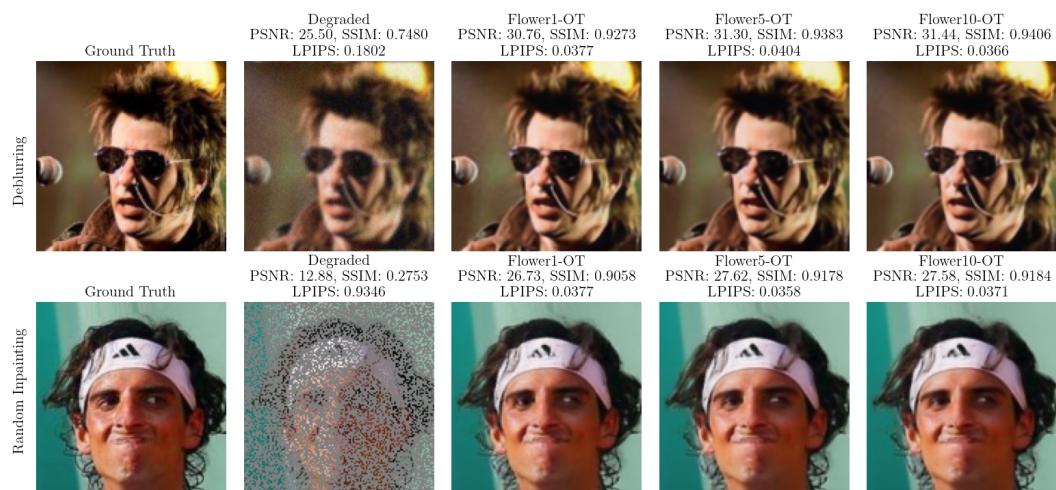
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
11801181  
1182  
1183  
1184  
1185  
1186  
1187

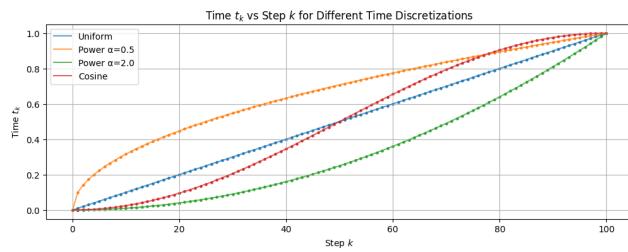
Figure 13: Visual comparison of the effect of averaging over different numbers of evaluations of *Flower*.

1188  
1189

## 8.5.5 EFFECT OF ADAPTIVE TIME STEPS

1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204

Our theoretical framework does not tie *Flower* to a particular time discretization. This raises the question of whether non-uniform time steps could improve practical performance. To explore this, we compare the uniform (Euler) time discretization with several alternatives. Specifically, we consider a power-law schedule  $t_k = \left(\frac{k}{N}\right)^\alpha$ , where  $\alpha = 1$  recovers the uniform grid,  $\alpha > 1$  concentrates steps near the start of the trajectory, and  $\alpha < 1$  allocates more steps near the end. In our experiments, we use  $\alpha = 0.5$  and  $\alpha = 2$ . We further evaluate a cosine schedule  $t_k = \frac{1 - \cos(\pi k/N)}{2}$ , which yields finer resolution at both the beginning and end of the trajectory. For clarity, Figure 14 visualizes these time grids for  $N = 100$ . We vary the total number of steps  $N \in \{10, 20, 50, 100\}$  and validate these schedules on two restoration tasks (deblurring and random inpainting) shown in Figures 15 and 16. For the hyperparameters of *Flower*, we use  $\gamma = 0$ , and  $N_{\text{Avg}} = 1$ . Our results indicate that the  $\alpha = 0.5$  schedule achieves noticeably better reconstruction quality with fewer total steps  $N$ , while  $\alpha = 2$  tends to underperform relative to the uniform case. The cosine schedule sometimes provides improvements at low step counts. For larger numbers of steps, all schedules converge to similar performance. Our findings highlight the potential of adaptive time discretizations (with more resolution toward the end of trajectory) to improve quality–compute trade-offs.

1205  
1206  
1207  
1208  
1209  
1210  
1211  
12121213  
1214  
1215Figure 14: Time  $t_k$  vs step  $k$  for different time discretizations.1216  
1217

## 8.5.6 INCREASING THE SIZE OF TEST SETS

1218  
1219  
1220  
1221  
1222  
1223  
1224

In the main paper, we evaluated all tasks for each method using 100 CelebA images and 100 AFHQ-Cat images. We chose a test size of 100 because methods such as Flow Priors and D-Flow require backpropagation during inference and are therefore slow. In this appendix, we increase the test set to 1000 images for CelebA and 400 images for AFHQ-Cat in order to further validate *Flower* compared to existing (efficient) flow-matching methods. The results, presented in Tables 5 and 6, show the same trend as in the main paper as *Flower* achieves competitive reconstruction quality.

1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234

Table 5: Results on 1000 test images of the dataset CelebA.

| Method            | Denoising    |              |              | Deblurring   |              |              | Super-resolution |              |              | Random inpainting |              |              | Box inpainting |              |              |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|--------------|--------------|-------------------|--------------|--------------|----------------|--------------|--------------|
|                   | PSNR         | SSIM         | LPIPS        | PSNR         | SSIM         | LPIPS        | PSNR             | SSIM         | LPIPS        | PSNR              | SSIM         | LPIPS        | PSNR           | SSIM         | LPIPS        |
| Degraded          | 20.00        | 0.351        | 0.368        | 27.83        | 0.741        | 0.123        | 10.38            | 0.185        | 0.828        | 12.10             | 0.197        | 1.038        | 22.29          | 0.745        | 0.211        |
| OT-ODE            | 30.53        | 0.857        | <b>0.032</b> | 33.06        | 0.920        | <b>0.029</b> | 31.54            | <b>0.905</b> | <b>0.024</b> | 28.74             | <b>0.870</b> | 0.052        | 29.65          | 0.921        | <b>0.035</b> |
| PnP-Flow1         | <b>31.84</b> | <b>0.904</b> | 0.044        | <b>34.56</b> | <b>0.935</b> | 0.038        | 31.16            | 0.901        | 0.044        | <b>33.16</b>      | <b>0.944</b> | <b>0.019</b> | <b>30.47</b>   | <b>0.935</b> | 0.036        |
| Flower1-OT (ours) | <b>32.31</b> | <b>0.913</b> | 0.033        | <b>35.03</b> | <b>0.946</b> | <b>0.026</b> | <b>32.45</b>     | <b>0.922</b> | 0.034        | <b>33.19</b>      | <b>0.944</b> | <b>0.017</b> | <b>31.22</b>   | <b>0.946</b> | <b>0.021</b> |

1235  
1236  
1237  
1238  
1239  
1240  
1241

Table 6: Results on 400 test images of the dataset AFHQ-Cat.

| Method            | Denoising    |              |              | Deblurring   |              |              | Super-resolution |              |              | Random inpainting |              |              | Box inpainting |              |              |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|--------------|--------------|-------------------|--------------|--------------|----------------|--------------|--------------|
|                   | PSNR         | SSIM         | LPIPS        | PSNR         | SSIM         | LPIPS        | PSNR             | SSIM         | LPIPS        | PSNR              | SSIM         | LPIPS        | PSNR           | SSIM         | LPIPS        |
| Degraded          | 20.00        | 0.315        | 0.517        | 24.03        | 0.516        | 0.452        | 11.88            | 0.217        | 0.881        | 13.55             | 0.231        | 1.071        | 21.80          | 0.741        | 0.203        |
| OT-ODE            | 30.01        | 0.814        | <b>0.077</b> | 27.06        | 0.711        | <b>0.126</b> | 25.92            | 0.715        | <b>0.109</b> | 29.38             | 0.839        | 0.091        | 24.77          | 0.875        | <b>0.085</b> |
| PnP-Flow1         | <b>31.17</b> | <b>0.862</b> | 0.136        | <b>27.94</b> | <b>0.759</b> | 0.306        | <b>26.96</b>     | <b>0.762</b> | 0.170        | <b>33.01</b>      | <b>0.918</b> | <b>0.037</b> | <b>26.46</b>   | <b>0.897</b> | 0.102        |
| Flower1-OT (ours) | <b>31.66</b> | <b>0.878</b> | 0.104        | <b>28.63</b> | <b>0.773</b> | 0.255        | <b>26.24</b>     | <b>0.740</b> | 0.273        | <b>32.98</b>      | <b>0.918</b> | 0.040        | <b>26.68</b>   | <b>0.915</b> | <b>0.062</b> |

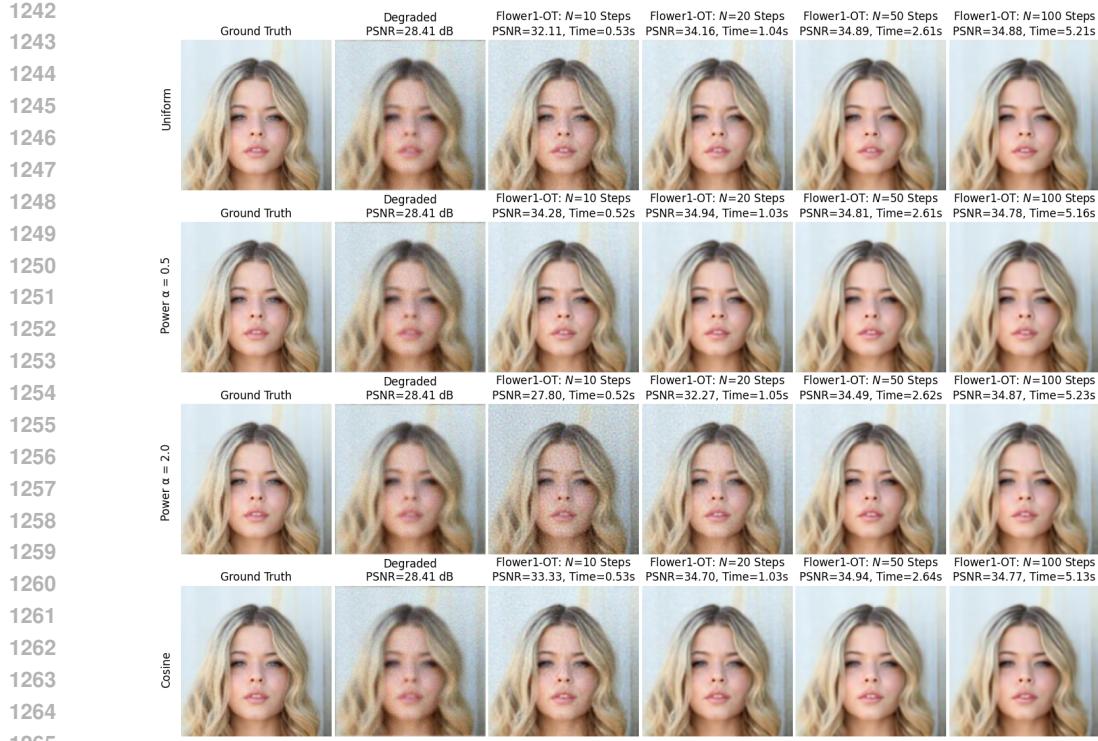


Figure 15: Deblurring example for uniform versus adaptive solvers.

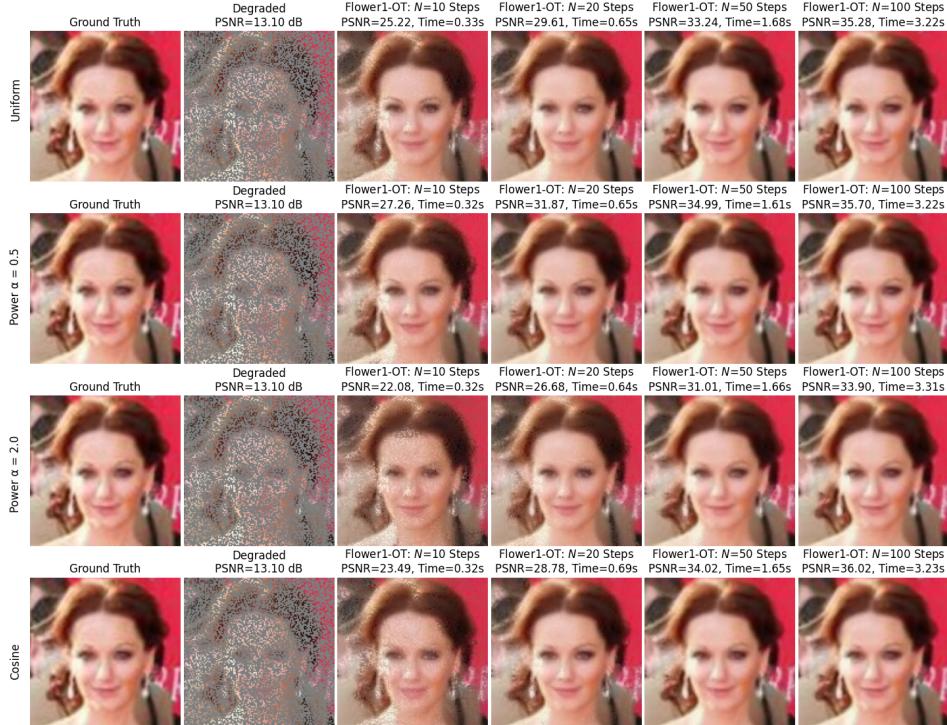
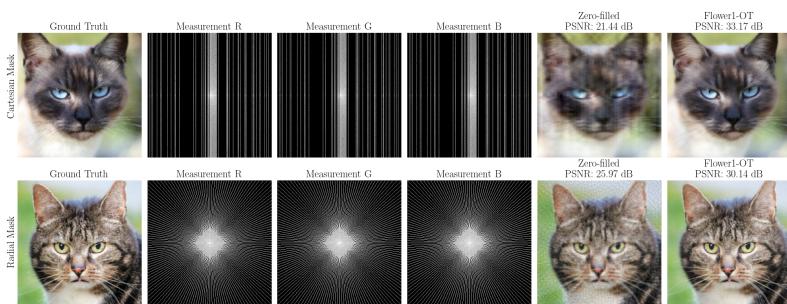
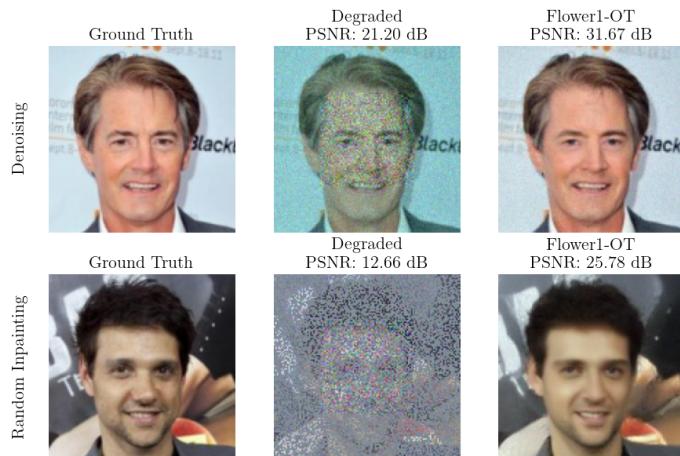


Figure 16: Random inpainting example for uniform versus adaptive solvers.

1296 8.5.7 APPLICATION TO FURTHER INVERSE PROBLEMS  
12971298 We aim to validate the generality of our framework by evaluating it on a wider range of inverse  
1299 problems.1300  
1301 **Compressed Sensing Fourier Sampling.** We choose a forward operator inspired by magnetic  
1302 resonance imaging (MRI) reconstruction: the Fourier transform followed by a binary sampling mask.  
1303 We use two masks: (i) a Cartesian mask with a sampling ratio of 0.2188, and (ii) a radial mask with  
1304 a sampling ratio of 0.2990. The measurements are corrupted with additive white Gaussian noise of  
1305 standard deviation  $\sigma_n = 0.002$ . For our examples, we use two images from the AFHQ-Cat dataset.  
1306 For the hyperparameters of *Flower*, we use  $\gamma = 0$ ,  $N = 100$ , and  $N_{\text{Avg}} = 1$ . Since the images  
1307 are RGB, we apply the forward operator channel by channel. We summarize our results in Figure  
1308 17. In both cases, we observe that *Flower* successfully handles this inverse problem and produces  
1309 reconstructions that improve on the zero-filled baseline.1320 Figure 17: Reconstruction results for compressed sensing with Cartesian and radial masks.  
1321  
13221323  
1324 **Non-Isotropic Gaussian Noise.** Here, we aim to verify the theoretical results from Appendix 8.3  
1325 for non-isotropic additive Gaussian noise using numerical experiments. For demonstration, we focus  
1326 on two tasks: image denoising and random inpainting using the same setup described in Section 5  
1327 for the CelebA images, except for the noise dynamics. Here, we add non-isotropic Gaussian noise  
1328 to the image by applying Gaussian noise with  $\sigma_n = 3$  to the central box of size  $(64 \times 64)$  of the  
1329  $(128 \times 128)$  image and  $\sigma_n = 1$  outside this box. For the *Flower* hyperparameters, we use  $\gamma = 0$ ,  
1330  $N = 100$ , and  $N_{\text{Avg}} = 1$ . We observe that, consistent with our theory, we are able to recover  
1331 good-quality reconstructions given a non-isotropic Gaussian noise.1347 Figure 18: Visual results for inverse problems with non-isotropic Gaussian noise.  
1348  
1349

1350 8.5.8 HYPERPARAMETERS FOR ALL METHODS  
13511352 In Tables 7 and 8, we report the hyperparameters that we used for all methods. Most of the hyper-  
1353 parameters are adapted from Martin et al. (2025).  
13541355 Table 7: Hyperparameters for all methods on the CelebA dataset.  
1356

| Method      | Hyperparameters  | Denoising             | Deblurring              | Super-resolution        | Random inpainting       | Box inpainting           |
|-------------|--|-----------------------|-------------------------|-------------------------|-------------------------|--------------------------|
| DiffPIR     | $\zeta$ (blending)<br>$\lambda$ (regularization)   | 1.0<br>1.0            | 1.0<br>1000.0           | 1.0<br>100.0            | 1.0<br>1.0              | N/A<br>N/A               |
| PnP-GS      | $\gamma$ (learning rate)<br>$\alpha$ (inertia param.)<br>$\sigma_f$ (factor for noise input)<br>$n_{iter}$ (number of iter.) | -<br>1.0<br>1.0<br>1  | 2.0<br>0.5<br>1.8<br>35 | 2.0<br>1.0<br>3.0<br>20 | 1.0<br>0.5<br>1.0<br>23 | N/A<br>N/A<br>N/A<br>N/A |
| OT-ODE      | $t_0$ (initial time)<br>$\gamma$   | 0.3<br>time-dependent | 0.4<br>time-dependent   | 0.1<br>constant         | 0.1<br>constant         | 0.1<br>time-dependent    |
| Flow-Priors | $\lambda$ (regularization)<br>$\eta$ (learning rate)   | 100<br>0.01           | 1,000<br>0.01           | 10,000<br>0.1           | 10,000<br>0.01          | 10,000<br>0.01           |
| D-Flow      | $\lambda$ (regularization)<br>$\alpha$ (blending)<br>$n_{iter}$ (number of iter.)  | 0.001<br>0.1<br>3     | 0.001<br>0.1<br>7       | 0.001<br>0.1<br>10      | 0.01<br>0.1<br>20       | 0.001<br>0.1<br>9        |
| PnP-Flow1   | $\alpha$ (learning-rate factor)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                            | 0.8<br>100<br>1       | 0.01<br>100<br>1        | 0.3<br>100<br>1         | 0.01<br>100<br>1        | 0.5<br>100<br>1          |
| PnP-Flow5   | $\alpha$ (learning-rate factor)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                            | 0.8<br>100<br>5       | 0.01<br>100<br>5        | 0.3<br>100<br>5         | 0.01<br>100<br>5        | 0.5<br>100<br>5          |
| Flower1-OT  | $\gamma$ (refinement uncertainty)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                          | 0<br>100<br>1         | 0<br>100<br>1           | 0<br>100<br>1           | 0<br>100<br>1           | 0<br>100<br>1            |
| Flower5-OT  | $\gamma$ (refinement uncertainty)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                          | 0<br>100<br>5         | 0<br>100<br>5           | 0<br>100<br>5           | 0<br>100<br>5           | 0<br>100<br>5            |

1378 Table 8: Hyperparameters for all methods on the AFHQ-Cat dataset.  
1379

| Method      | Denoising  | Deblurring            | Super-resolution        | Random inpainting       | Box inpainting          |
|-------------|--|-----------------------|-------------------------|-------------------------|-------------------------|
| DiffPIR     | $\zeta$ (blending)<br>$\lambda$ (regularization)   | 1.0<br>1.0            | 1.0<br>1000.0           | 1.0<br>100.0            | 1.0<br>1.0              |
| PnP-GS      | $\gamma$ (learning rate)<br>$\alpha$ (inertia param.)<br>$\sigma_f$ (factor for noise input)<br>$n_{iter}$ (number of iter.) | -<br>1.0<br>1.0<br>1  | 2.0<br>0.3<br>1.8<br>60 | 2.0<br>1.0<br>5.0<br>50 | 1.0<br>0.5<br>1.0<br>23 |
| OT-ODE      | $t_0$ (initial time)<br>$\gamma$   | 0.3<br>time-dependent | 0.3<br>time-dependent   | 0.1<br>constant         | 0.1<br>constant         |
| Flow-Priors | $\lambda$ (regularization)<br>$\eta$ (learning rate)   | 100<br>0.01           | 1,000<br>0.01           | 10,000<br>0.1           | 10,000<br>0.01          |
| D-Flow      | $\lambda$ (regularization)<br>$\alpha$ (blending)<br>$n_{iter}$ (number of iter.)  | 0.001<br>0.1<br>3     | 0.01<br>0.5<br>20       | 0.001<br>0.1<br>20      | 0.001<br>0.1<br>9       |
| PnP-Flow1   | $\alpha$ (learning-rate factor)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                            | 0.8<br>100<br>1       | 0.01<br>500<br>1        | 0.01<br>500<br>1        | 0.01<br>200<br>1        |
| PnP-Flow5   | $\alpha$ (learning-rate factor)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                            | 0.8<br>100<br>5       | 0.01<br>500<br>5        | 0.01<br>500<br>5        | 0.01<br>200<br>5        |
| Flower1-OT  | $\gamma$ (refinement uncertainty)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                          | 0<br>100<br>1         | 0<br>100<br>1           | 0<br>500<br>1           | 0<br>200<br>1           |
| Flower5-OT  | $\gamma$ (refinement uncertainty)<br>$N$ (Number of time steps)<br>$N_{Avg}$ (Number of averagings)                          | 0<br>100<br>5         | 0<br>100<br>5           | 0<br>500<br>5           | 0<br>200<br>5           |