

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EQUIREG: EQUIVARIANCE REGULARIZED DIFFUSION FOR INVERSE PROBLEMS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Diffusion models represent the state-of-the-art for solving inverse problems such as image restoration tasks. Diffusion-based inverse solvers incorporate a likelihood term to guide prior sampling, generating data consistent with the posterior distribution. However, due to the intractability of the likelihood, most methods rely on isotropic Gaussian approximations, which can push estimates off the data manifold and produce inconsistent, poor reconstructions. We propose *Equivariance Regularized* (EquiReg) diffusion, a general plug-and-play framework that improves posterior sampling by penalizing those that deviate from the data manifold. EquiReg formalizes manifold-preferential equivariant functions that exhibit low equivariance error for on-manifold samples and high error for off-manifold ones, thereby guiding sampling toward symmetry-preserving regions of the solution space. We highlight that such functions naturally emerge when training non-equivariant models with augmentation or on data with symmetries. EquiReg is particularly effective under reduced sampling and measurement consistency steps, where many methods suffer severe quality degradation. By regularizing trajectories toward the manifold, EquiReg implicitly accelerates convergence and enables high-quality reconstructions. EquiReg consistently improves performance in linear and nonlinear image restoration tasks and solving partial differential equations.

## 1 INTRODUCTION

Inverse problems aim to recover an unknown signal  $\mathbf{x}^* \in \mathbb{R}^d$  from undersampled noisy measurements:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}^*) + \boldsymbol{\nu} \in \mathbb{R}^m, \quad (1)$$

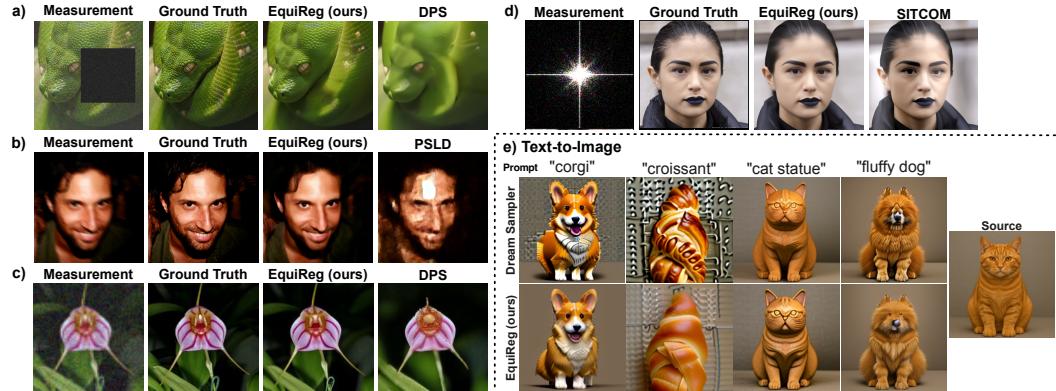


Figure 1: **EquiReg's broad applicability.** a-d) image restoration inverse problems and e) text-guided image generation, resulting in artifact reduction and more realistic generation. Here, EquiReg refers to our regularization being applied to the diffusion sampling method on the same row.

where  $\mathcal{A}$  is a known measurement operator, and  $\boldsymbol{\nu}$  is an unknown noise (Groetsch, 1993). Inverse problems are widely studied in science and engineering, including imaging and astrophotography.

Inverse problems are ill-posed, i.e., the inversion process can have many solutions; hence, they require prior information about the desired solution (Kabanikhin, 2008). In the Bayesian formulation, the solution maximizes the posterior distribution  $p(\mathbf{x}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ , where  $p(\mathbf{y}|\mathbf{x})$  is the likelihood

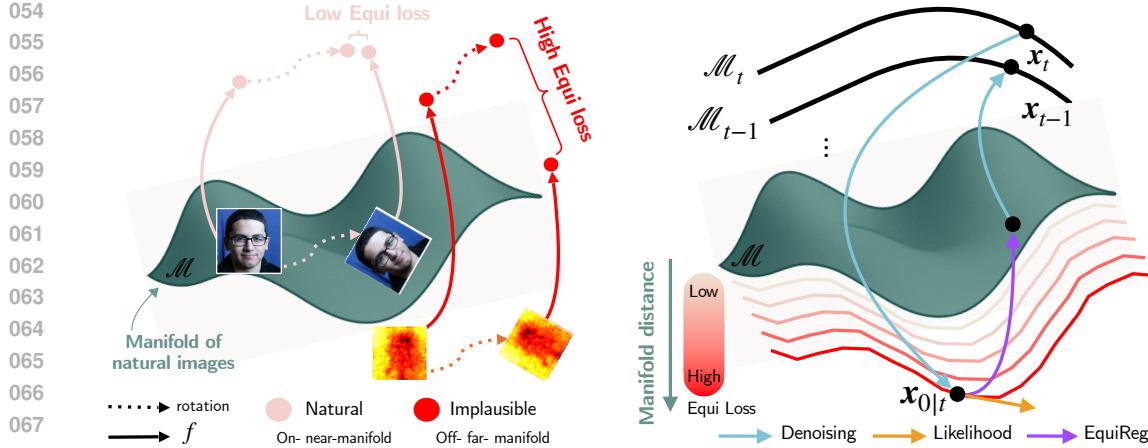


Figure 2: **Equivariance Regularized (EquiReg) diffusion for inverse problems.** (left) Manifold preferential equivariance (MPE) functions whose equivariance error is lower for on-manifold and higher for off-manifold data. (right) EquiReg regularizes the posterior sampling trajectory for improved performance. It penalizes off-manifold trajectories via MPE-based regularization.

of the measurements and  $p(\mathbf{x})$  is a prior describing the signal structure (Stuart, 2010). Examples of handcrafted priors include sparsity (Donoho, 2006) and low-rankness (Candès et al., 2011).

This paper focuses on methods that leverage unconditionally pre-trained score-based generative diffusion models as learned priors (Ho et al., 2020; Song and Ermon, 2019) with applications in image restoration (Chung et al., 2023), medical imaging (Chung et al., 2022a), and solving partial differential equations (PDEs) (Huang et al., 2024; Yao et al., 2025). These methods define a sequential noising process  $\mathbf{x}_0 \sim p_{\text{data}} \rightarrow \mathbf{x}_t \rightarrow \mathbf{x}_T \sim p_T(\mathbf{x}) \approx \mathcal{N}(\mathbf{0}, \mathbf{I})$  and a reverse denoising process parameterized by a neural network score  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$  (Vincent, 2011). During sampling, these approaches incorporate gradient signals carrying likelihood information to solve inverse problems.

Solving inverse problems with diffusion (Zhang et al., 2025a; Alkhouri et al., 2025) requires computing the conditional score  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t | \mathbf{y})$ , decomposed into  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \mathbf{x}_t)$ . This introduces challenges, as the likelihood score  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \mathbf{x}_t) = \nabla_{\mathbf{x}_t} \log \int p(\mathbf{y} | \mathbf{x}_0) p_t(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0$  is only computationally tractable when  $t = 0$ . To handle the likelihood for  $t > 0$ , many methods approximate the posterior  $p_t(\mathbf{x}_0 | \mathbf{x}_t)$  with the isotropic Gaussian distribution (Zhang et al., 2025a), where the distribution expectation is computed using the optimal denoising score (Robbins, 1956). The Gaussian approximation can be inaccurate for complex distributions (Figure 3), leading to errors in likelihood computation, especially with point estimations (Chung et al., 2023). Since the posterior expectation is a conditional expectation, a linear combination of all possible  $\mathbf{x}_0$ , it may lie off the data manifold even when individual samples remain on it. These issues are further amplified in latent diffusion models (LDMs), introducing artifacts (Rout et al., 2023).

Prior work has attempted to address this challenge via projection-based (He et al., 2024; Zirvi et al., 2025) or decoupled optimization strategies (Zhang et al., 2025a), aimed at reducing the propagation of measurement consistency errors during sampling. However, they still rely on the isotropic Gaussian assumption, which can lead to failures on difficult tasks or when the number of sampling steps is reduced. While higher-order statistics can reduce errors (Boys et al., 2024), most approaches still rely on the approximation for its efficiency, scalability, and simplicity (Alkhouri et al., 2025), often coupled with large-scale LDMs (Peebles and Xie, 2023). This raises a key question: how can we ensure the reliability and practicality of conditional diffusion models under this approximation?

Equivariance provides a natural mechanism to keep sampling trajectories close to the data manifold. We therefore address this challenge with a regularization scheme that leverages equivariance to improve posterior sampling by guiding diffusion trajectories toward symmetry-preserving solution spaces. Prior studies have enforced equivariance directly on the generation or denoising process (Chen et al., 2023a; Terris et al., 2024), with extensions to probabilistic symmetries (Bloem-Reddy et al., 2020) enabling more sample-efficient diffusion models (Wang et al., 2024).

Our approach differs as follows: rather than strictly enforcing equivariance within denoising architectures, which can hinder tasks requiring symmetry breaking (Lawrence et al., 2025), we employ equivariance as a plug-and-play regularizer to guide diffusion trajectories toward the data manifold.

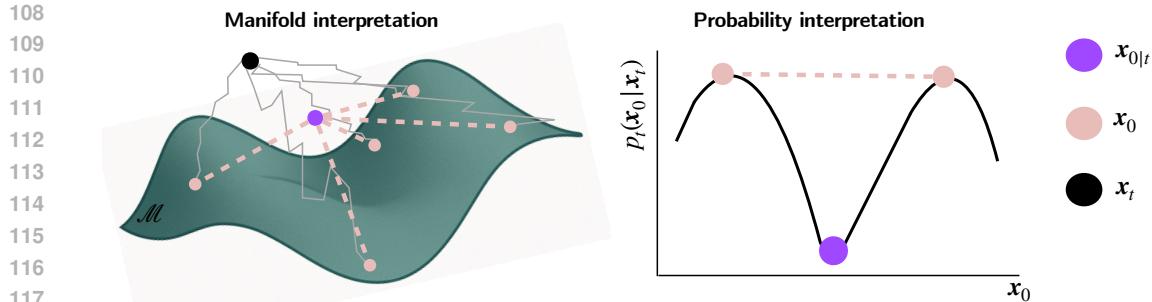


Figure 3: **Off-manifold posterior expectation.** This impacts the likelihood score  $p_t(y|x_t) = \int p(y|x_0)p_t(x_0|x_t)dx_0$  computation achieved via isotropic Gaussian modelling of  $p_t(x_0|x_t)$ .

**Our contributions.** We propose *Equivariance Regularized* (EquiReg) diffusion, an equivariance-based regularization framework for solving inverse problems with diffusion models (Figure 2). EquiReg leverages equivariance to *regularize* likelihood-induced errors during posterior sampling, guiding diffusion trajectories toward more consistent, on-manifold solutions. Crucially, it employs *Manifold-Preferential Equivariant* (MPE) functions, which discriminate on-manifold from off-manifold data by exhibiting low equivariance error in-distribution and higher error out-of-distribution. We formalize that an effective regularizer should capture such a global property, and MPE functions provide a principled way to direct sampling toward plausible solutions. This design makes EquiReg architecture-agnostic: the regularizer operates independently of the diffusion model itself. With a suitable MPE function, EquiReg improves performance across models, including those with equivariant scores, where likelihood guidance may otherwise push trajectories off the manifold.

We observe that many practical functions behave as MPEs: their equivariance error is small on the training or data manifold but grows off-manifold. This behavior arises in learned models trained with data augmentation, as well as in data with inherent symmetries such as those from physical systems. Rather than treating the degradation off-manifold as a limitation, we exploit it as a signal: equivariance error serves as a natural discriminator of equivariance for identifying undesirable states during diffusion sampling. Building on this idea, we construct pre-trained MPEs as the foundation of our EquiReg loss. The choice of this function is independent of the denoiser in diffusion models and can be derived separately. For instance, if the diffusion model architecture is itself equivariant, it cannot be leveraged for regularization as it cannot discriminate between on- and off-manifold samples. Instead, a separate non-equivariant architecture can be used to train to derive EquiReg.

We validate the effectiveness of EquiReg through extensive experiments across diverse diffusion models, inverse problems, and datasets. We demonstrate that EquiReg improves perceptual image quality and remains effective in cases where baselines fail. We show that EquiReg improves the performance of SITCOM (Alkhouri et al., 2025) and DPS (Chung et al., 2023) when the number of measurement consistency and sampling steps are reduced, thus moving toward more efficient diffusion-based solvers. Our method is particularly useful when applied to LDMs. EquiReg reduces failure cases, and consistently improves PSLD (Rout et al., 2023), ReSample (Song et al., 2023a), and DPS (Chung et al., 2023) on linear and nonlinear image restoration tasks. For example, EquiReg significantly improves the LPIPS (Song et al., 2023a) of ReSample by 51% for motion deblur and the FID of DPS (Chung et al., 2023) by 59% on super-resolution. **We also include diversity analyses, demonstrating that EquiReg maintains diversity without collapse of single mode reconstruction.**

We extend EquiReg’s applicability to function-space diffusion models and demonstrate its added benefit for solving PDEs. EquiReg achieves a 7.3% relative reduction in the  $\ell_2$  error of FunDPS (Mammadov et al., 2024a; Yao et al., 2025) on the Helmholtz equation and a 7.5% relative reduction on the Navier-Stokes equation. Lastly, we include preliminary experiments on EquiReg improving the realism and plausibility of text-guided image generation, emphasizing that the benefits of EquiReg extend beyond image restorations. Overall, the flexibility of EquiReg as a plug-and-play regularization framework suggests that its utility will extend well beyond the specific methods studied in this paper.

## 2 PRELIMINARIES AND RELATED WORKS

**Diffusion models.** Diffusion generative models (Ho et al., 2020; Song and Ermon, 2019; Sohl-Dickstein et al., 2015; Kadkhodaei and Simoncelli, 2021) are state-of-the-art in computer vision for image (Esser et al., 2024) and video generation (Brooks et al., 2024; Zhang et al., 2025b), with score-based methods (Song et al., 2021) being among the most widely used. Diffusion models generate data

162 via a reverse noising process. The forward noising process transforms the data sample  $\mathbf{x}_0 \sim p_{\text{data}}$   
 163 via a series of additive noise into an approximately Gaussian distribution ( $p_{\text{data}} \rightarrow p_t \rightarrow \mathcal{N}(0, I)$   
 164 as  $t \rightarrow \infty$ ), described by the stochastic differential equation (SDE)  $d\mathbf{x} = -\frac{\beta_t}{2}\mathbf{x}dt + \sqrt{\beta_t}d\mathbf{w}$ ,  
 165 where  $\mathbf{w}$  is a standard Wiener process, and the drift and diffusion coefficients are parameterized by  
 166 a monotonically increasing noise scheduler  $\beta_t \in (0, 1)$  in time  $t$  (Ho et al., 2020). Reversing the  
 167 forward diffusion process is described by (Anderson, 1982)

$$d\mathbf{x} = [-\frac{\beta_t}{2}\mathbf{x} - \beta_t \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)] dt + \sqrt{\beta_t} d\bar{\mathbf{w}} \quad (2)$$

170 with  $dt$  moving backward in time or in discrete steps from  $T$  to 0. This reverse SDE is used to  
 171 sample data from the distribution  $p_{\text{data}}$ , where the unknown gradient  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$  is approximated  
 172 by a scoring function  $s_\theta(\mathbf{x}_t, t)$ , parameterized by a neural network and learned via denoising score  
 173 matching methods (Hyvärinen and Dayan, 2005; Vincent, 2011). Solving inverse problems is  
 174 described as a conditional generation where the data is sampled from the posterior  $p(\mathbf{x}|\mathbf{y})$ :

$$d\mathbf{x} = [-\frac{\beta_t}{2}\mathbf{x}dt - \beta_t(\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t))]dt + \sqrt{\beta_t} d\bar{\mathbf{w}} \quad (3)$$

175 For solving general inverse problems where the diffusion is *pre-trained* unconditionally, the prior score  
 176  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$  can be estimated using  $s_\theta(\mathbf{x}_t, t)$ . However, the likelihood score  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t)$  is  
 177 only known at  $t = 0$ , otherwise it is computationally intractable.

178 **Diffusion models for inverse problems.** Solving inverse problems with pre-trained diffusion  
 179 models requires approximating the intractable likelihood score  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t)$ . Training-free  
 180 solvers differ in how they approximate  $p_t(\mathbf{y}|\mathbf{x}_t)$  and combine it with the sampling prior  $p_t(\mathbf{x}_t)$  (Peng  
 181 et al., 2024). Since  $p_t(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0)p_t(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$ , the common choice is to approximate  
 182  $p_t(\mathbf{x}_0|\mathbf{x}_t)$  by an isotropic Gaussian  $\mathcal{N}(\mathbf{x}_{0|t}, r_t^2 \mathbf{I})$  (Chung et al., 2023; Song et al., 2023b; Zhu  
 183 et al., 2023; Zhang et al., 2025a). With an optimal denoising score  $s_\theta(\mathbf{x}_t, t)$ , the posterior mean  
 184  $\mathbf{x}_{0|t} := \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$  follows from Tweedie's formula (Robbins, 1956; Miyasawa et al., 1961; Efron,  
 185 2011). Although this yields an MMSE estimate, for complex or multimodal distributions,  $p_t(\mathbf{x}_0|\mathbf{x}_t)$   
 186 may not be concentrated around its mean, leading to off-manifold solutions (see Figure 3).

187 **Equivariance.** Equivariance is a property describing how functions transform predictably under  
 188 group actions. It serves as a powerful strategy for incorporating symmetries into deep learning (Bron-  
 189 stein et al., 2021). Prior work has applied equivariance to graph networks (Satorras et al., 2021), con-  
 190 volutional networks (Cohen and Welling, 2016; Romero and Lohit, 2022), Lie groups for modelling  
 191 dynamical systems (Finzi et al., 2020), and diffusion models (Wang et al., 2024) with applications in  
 192 molecular generation (Hoogeboom et al., 2022; Cornet et al., 2024), autonomous driving (Chen et al.,  
 193 2023b), robotics (Brehmer et al., 2023), crystal structure prediction (Jiao et al., 2023), and audio in-  
 194 verse problems (Moliner et al., 2023). Equivariance guidance has also been used to improve temporal  
 195 consistency in video generation (Daras et al., 2024). The benefits of equivariance as a prior to solve  
 196 inverse problems (Scanvic et al., 2025) are theoretically supported in compressed sensing (Tachella  
 197 et al., 2023). An equivariant function respects symmetries under group transformations, i.e.,

198 **Definition 2.1** (Equivariance). *Let  $G$  act on  $\mathcal{Z}$  via  $T_g : \mathcal{Z} \rightarrow \mathcal{Z}$  and on  $\mathcal{X}$  via  $S_g : \mathcal{X} \rightarrow \mathcal{X}$ . A  
 199 function  $f : \mathcal{Z} \rightarrow \mathcal{X}$  is equivariant if for all  $g \in G$  and  $z \in \mathcal{Z}$ ,  $f(T_g(z)) = S_g(f(z))$ .*

200 While prior work leverages exact equivariance as in Definition 2.1 to directly incorporate symmetries  
 201 into deep neural networks, recent studies explore approximate equivariant networks to relax strict  
 202 mathematical symmetries that may not fully hold in real-world data, aiming to improve perfor-  
 203 mance (Wang et al., 2022). They propose a definition of approximate equivariance (Definition 2.2),  
 204 along with an equivariance error of functions to quantify the deviation from perfect symmetry.

205 **Definition 2.2** (Approximate Equivariant Functions). *Let  $G$  act on  $\mathcal{Z}$  via  $T_g : \mathcal{Z} \rightarrow \mathcal{Z}$  and on  $\mathcal{X}$   
 206 via  $S_g : \mathcal{X} \rightarrow \mathcal{X}$ . A function  $f : \mathcal{Z} \rightarrow \mathcal{X}$  is  $\epsilon$ -approximate equivariant if for all  $g \in G$  and  $z \in \mathcal{Z}$ ,  
 207  $\|S_g(f(z)) - f(T_g(z))\| \leq \epsilon$ . The equivariance error of the function  $f : \mathcal{Z} \rightarrow \mathcal{X}$  is defined as  
 208  $\sup_{z,g} \|S_g(f(z)) - f(T_g(z))\|$ . Hence,  $f$  is  $\epsilon$ -approximate equivariant iff its error  $< \epsilon$ .*

209 Finally, this paper uses the term manifold which refers to the data manifold hypothesis (see Ass-  
 210 sumption H.1) (Cayton et al., 2005) that assumes data is sampled from a low-dimensional manifold  
 211 embedded in a high-dimensional space. This hypothesis is popular in machine learning (Bordt et al.,  
 212 2023) and diffusion-based solvers (He et al., 2024; Chung et al., 2022b; 2023), supported by empirical  
 213 evidence for imaging (Weinberger and Saul, 2006).

---

216 

### 3 EQUIREG: EQUIVARIANCE REGULARIZED DIFFUSION

217

218 We begin by presenting a generalized regularization framework for improving diffusion-based inverse  
219 solvers. We then focus on the property of *equivariance* and introduce a new class of functions whose  
220 equivariance errors are distribution-dependent (low for on- or near-manifold samples and high for  
221 off-manifold samples). Finally, we leverage these functions to regularize diffusion models, guiding  
222 sampling trajectories toward better inverse solutions.

223 This paper addresses the propagation error introduced by the approximation of posterior  $p_t(\mathbf{x}_0|\mathbf{x}_t)$   
224 by incorporating an explicit regularization term. The proposed framework is general and can be  
225 applied as plug-in on a wide range of pixel and latent-space diffusion models. Given  $p_t(\mathbf{y}|\mathbf{x}_t) =$   
226  $\int p(\mathbf{y}|\mathbf{x}_0)p_t(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$ , let  $\tilde{p}_t(\mathbf{x}_0|\mathbf{x}_t)$  denote an approximation of the posterior to make the likeli-  
227 hood tractable. We formulate the regularized reverse diffusion dynamics as

228 
$$d\mathbf{x} = [-\frac{\beta_t}{2}\mathbf{x}dt - \beta_t \nabla_{\mathbf{x}_t}(\log p_t(\mathbf{x}_t) + \log \int p(\mathbf{y}|\mathbf{x}_0)\tilde{p}_t(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0 - \mathcal{R}(\mathbf{x}_t))]dt + \sqrt{\beta_t}d\bar{\mathbf{w}}, \quad (4)$$
229

230 where  $\mathcal{R}(\mathbf{x}_t)$  is the regularizer. Applying this to DPS (Chung et al., 2023) takes the form in [Algorithm 1](#). This formulation raises two questions: i) how to design the regularizer, and ii) how to  
231 interpret the role of  $\mathcal{R}$  in regularizing conditional diffusion models and its impact on the sampling  
232 trajectory. We gain insight into the desirable properties of an optimal regularizer by reinterpreting the  
233 reverse conditional diffusion process as a time-inhomogeneous Wasserstein gradient flow (Ferreira  
234 and Valencia-Guevara, 2018) (see [Propositions G.1 and G.2](#) in Appendix). The analysis clarifies  
235 that an ideal regularizer should yield low values for on-manifold and high values for off-manifold  
236 samples, enabling accurate posterior sampling even when the likelihood score is approximated.

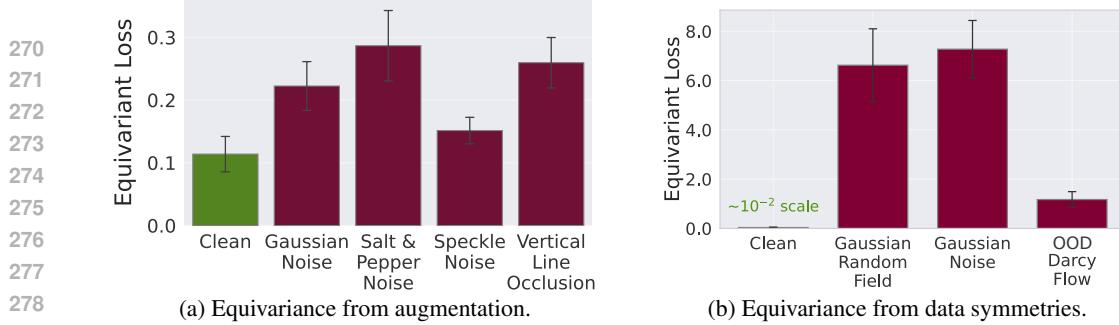
237 We further interpret this property in terms of  
238 sampling dynamics, i.e., when applied at each  
239 reverse-diffusion step, the regularizer effec-  
240 tively penalizes trajectories leaving the data  
241 manifold and reinforces those aligned with  
242 high-probability regions. This motivates de-  
243 signing a regularizer that corrects the entire  
244 functional being minimized globally, in con-  
245 trast to prior works that focus only on locally  
246 reducing likelihood error. The ideal property  
247 of a useful regularizer would be to produce  
248 high error on undesirable samples and low er-  
249 ror on desirable samples. We instantiate this  
250 ideal regularizer using equivariance, a global property that enforces geometric symmetries and guides  
251 the diffusion process toward the data manifold. To realize this idea, we seek functions that exhibit  
252 approximate equivariance and discriminate on- from off-manifold samples.

253 Thus, we propose to quantify equivariance of a function relative to a data distribution. Specifically,  
254 while the literature has primarily studied the equivariance properties of functions for general inputs,  
255 we propose a new definition for functions in which their equivariance error is distribution-dependent  
256 and defined under the support of an input data distribution ([Definition 3.1](#)).

257 **Definition 3.1** (Distribution-Dependent Equivariant Functions). *Let  $G$  act on  $\mathcal{Z}$  via  $T_g : \mathcal{Z} \rightarrow \mathcal{Z}$  and*  
258 *on  $\mathcal{X}$  via  $S_g : \mathcal{X} \rightarrow \mathcal{X}$ . The equivariance error of the function  $f : \mathcal{Z} \rightarrow \mathcal{X}$  under the distribution  $p$*   
259 *is defined as  $\sup_g \mathbb{E}_{\mathbf{z} \sim p} [\|S_g(f(\mathbf{z})) - f(T_g(\mathbf{z}))\|]$ .*

260 The above definition enables us to define functions whose equivariance error can differentiate on-  
261 manifold samples from off-manifold ones. Particularly, we aim to find functions whose equivariance  
262 error is low for on-manifold data and high elsewhere. We also introduce a constrained version of  
263 equivariance error, where the input is implicitly regularized to lie on the manifold  $\mathcal{M}$  in addition to  
264 minimizing the equivariance error ([Definition 3.2](#)). Both equivariance errors are non-local, defined  
265 at the distribution level. When used to regularize the reverse conditional diffusion process, they are  
266 computed via local evaluations over the sampled data.

267 **Definition 3.2** (Manifold-Constrained Distribution-Dependent Equivariant Functions). *Let  $G$  act on*  
268  *$\mathcal{Z}$  via  $T_g : \mathcal{Z} \rightarrow \mathcal{Z}$  and on  $\mathcal{X}$  via  $S_g : \mathcal{X} \rightarrow \mathcal{X}$ . The manifold-constrained equivariance error of the*  
269 *function  $f : \mathcal{Z} \rightarrow \mathcal{X}$  under the data distribution  $p$  is  $\sup_g \mathbb{E}_{\mathbf{z} \sim p} [\|\mathbf{z} - h(S_g^{-1}(f(T_g(\mathbf{z})))\|]$  where*  
 *$h : \mathcal{X} \rightarrow \mathcal{Z}$ , and the pair  $(f, h)$  forms a vanishing-error autoencoder (see [Appendix I](#)).*

Figure 4: **MPE function examples.**

To define our method, we term a class of *manifold-preferential equivariant (MPE)* functions, whose equivariance error is lower for samples on the data manifold than for off-manifold samples. *EquiReg* is a regularization framework, not a manifold projection method. *EquiReg* penalizes states that deviate from symmetry-preserving regions; when an MPE function is used, these regions align with the data manifold. In practice, MPE functions can emerge in different ways, which we illustrate with examples from augmented training and from data symmetries. MPE can emerge when functions are trained with symmetry-preserving mechanisms such as data augmentation. Prior work has studied equivariant properties of learned representations in deep networks (Lenc and Vedaldi, 2015), showing that data augmentations (Krizhevsky et al., 2012) and representation compression via reduced model capacity (Bruintjes et al., 2023) promote equivariant features even when equivariance is not explicitly built into the architecture. Importantly, the trained network is only approximately equivariant, and prior studies have noted that symmetry-preserving properties degrade for inputs deviating from in-distribution data (Azulay and Weiss, 2019). A few studies have leveraged this emergent MPE in trained networks for out-of-distribution detection (Zhou, 2022; Kaur et al., 2022; 2023).

To demonstrate the widespread MPE property of learned mappings, we have considered additional pre-trained models and quantified their equivariance loss for several set of data distributions, i.e., natural images and corrupted ones (see Section I of Appendix.) Figure 4a illustrates the MPE property, emergent via training with augmentations, of  $\mathcal{E}\mathcal{D}$  of a pre-trained autoencoder, currently used in LDMs. Specifically, it shows that the equivariance error is lower for natural images and increases when images deviate from the clean data distribution. Based on Definitions 3.1 and 3.2, we propose *Equi* and *EquiCon* losses using a pre-trained encoder-decoder for diffusion-based inverse solvers:

$$\begin{aligned} \text{Equi}_{\text{pixel}} \mathcal{R}(\mathbf{x}_t) &= \|S_g(\mathcal{E}(\mathbf{x}_{0|t})) - \mathcal{E}(T_g(\mathbf{x}_{0|t}))\|_2^2 \\ \text{Equi}_{\text{latent}} \mathcal{R}(\mathbf{z}_t) &= \|S_g(\mathcal{D}(\mathbf{z}_{0|t})) - \mathcal{D}(T_g(\mathbf{z}_{0|t}))\|_2^2 \\ \text{EquiCon}_{\text{latent}} \mathcal{R}(\mathbf{z}_t) &= \|\mathbf{z}_{0|t} - \mathcal{E}(S_g^{-1}(\mathcal{D}(T_g(\mathbf{z}_{0|t}))))\|_2^2, \end{aligned} \quad (5)$$

where  $\mathbf{x}_{0|t}$  and  $\mathbf{z}_{0|t}$  are function of  $\mathbf{x}_t$  and  $\mathbf{z}_t$ , respectively. MPE can also emerge due to symmetries present in the data itself during training. This often occurs in physics systems where coefficient functions, boundary values, and solution functions of PDEs remain valid under invertible coordinate transformations. Formally, let  $\mathcal{G}(a) \mapsto u$  be a PDE operator mapping initial condition  $a$  to solution  $u$ , and let  $T_g$  and  $S_g$  be invertible transformations that preserve PDE structure and boundary conditions. Then,  $S_g(\mathcal{G}(a)) = \mathcal{G}(T_g(a))$ . Neural operators (Kovachki et al., 2021), popular architectures for modelling physics, trained on PDEs with such inherent symmetries can learn equivariance properties. Figure 4b shows that we can construct an MPE function with Fourier Neural Operators (FNOs (Li et al., 2021)) trained on non-augmented physics data for Navier-Stokes, yielding lower error  $\|S_g(\text{FNO}(\mathbf{x}_{0|t})) - \text{FNO}(T_g(\mathbf{x}_{0|t}))\|_2^2$  on in-distribution as opposed to out-of-distribution data, with reflection as the group action.

The key message from our MPE examples is that MPE properties naturally emerge when a function (e.g., a neural network) is trained with appropriate augmentations or when the data itself exhibits inherent symmetries. Our paper leverages this property to distinguish on-manifold samples from off-manifold ones and to regularize the posterior sampling trajectory toward high-probability regions. Finally, we note that the choice of symmetry group may often be a challenge depending on application domain, a shared challenge in the broader equivariance literature. We provide guidelines on how to choose symmetry groups in Section H with literature reference on automatic symmetry discovery from data (Zhou et al.; Quessard et al., 2020; Dehmamy et al., 2021; Mohapatra et al., 2025).

Table 1: **Robustness and computational efficiency of applying EquiReg under various periods during sampling.** EquiReg maintains performance when applied every  $\{1, 2, 5, 10\}$  DDIM steps while incurring minimal computational overhead.

Method	Period	Super Resolution				Gaussian Blur			
		Runtime (s)	PSNR↑	LPIPS↓	FID↓	Runtime (s)	PSNR↑	LPIPS↓	FID↓
DPS	N/A	46.20	22.99 (1.93)	0.20 (0.05)	135.7	46.50	24.59 (2.25)	0.15 (0.03)	88.70
Equi-DPS	1	51.10	26.73 (1.99)	0.12 (0.03)	87.97	52.20	26.08 (2.25)	0.12 (0.03)	87.11
Equi-DPS	2	48.90	26.73 (1.99)	0.12 (0.03)	87.98	49.10	26.06 (2.24)	0.12 (0.03)	87.19
Equi-DPS	5	47.10	26.73 (1.99)	0.12 (0.03)	87.98	47.30	26.06 (2.24)	0.12 (0.03)	87.32
Equi-DPS	10	46.90	26.73 (1.99)	0.12 (0.03)	87.99	47.00	26.05 (2.24)	0.12 (0.03)	87.04

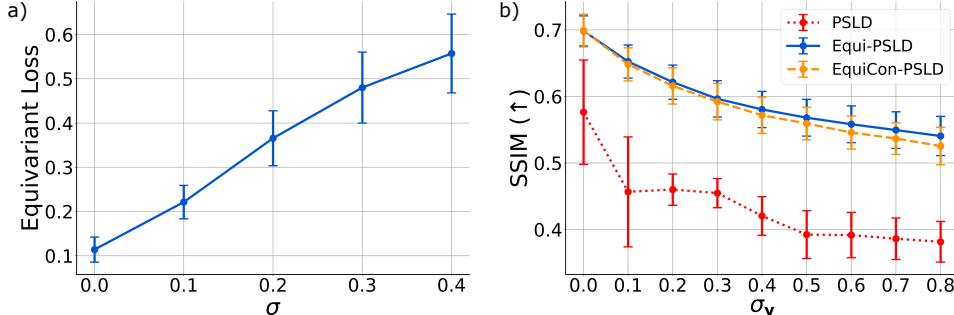


Figure 5: **EquiReg is effective across a range of measurement noise levels.** a) Equivariance error computed over a pre-trained decoder on increasingly noisy inputs. b) EquiReg performance computed over a range of measurement noise levels on the FFHQ dataset.

## 4 RESULTS

This section provides experimental results on the performance of EquiReg for inverse problems, including linear and nonlinear image restoration tasks and solving PDEs. To fairly assess EquiReg’s impact, we deliberately use a duo-setting comparison (e.g., PSLD vs. Equi-PSLD) across experiments, where all other factors (architecture, training, sampling) remain fixed. This ensures that any observed improvement can be attributed to EquiReg, not the underlying model or inference procedure. We also evaluate the impact of EquiReg under reduced measurement consistency and sampling steps, providing a path toward faster diffusion-based inverse solvers. Results emphasize the usefulness of EquiReg when the baseline performance deteriorates. Lastly, we provide preliminary analysis on EquiReg improving the realism of text-guided image generation.

**Image restoration tasks.** We evaluate the performance of EquiReg when applied to: SITCOM (Alkhouri et al., 2025), PSLD (Rout et al., 2023), ReSample (Song et al., 2023a), and DPS (Chung et al., 2023). We compare against several manifold-preserving or geometry-constraint approaches including MCG (Chung et al., 2022b), MPGD-AE (He et al., 2024), and DiffState-Grad (Zirvi et al., 2025). We measure performance via perceptual similarity (LPIPS), distribution alignment (FID), pixel-wise fidelity (PSNR), and structural consistency (SSIM). We test EquiReg on two datasets: a) the FFHQ  $256 \times 256$  validation dataset (Karras et al., 2021), and b) the ImageNet  $256 \times 256$  validation dataset (Deng et al., 2009). For pixel-based experiments, we use i) the pre-trained model from (Chung et al., 2023) on FFHQ, and ii) the pre-trained model from (Dhariwal and Nichol, 2021) on ImageNet. For latent diffusion experiments, we use i) the unconditional LDM-VQ-4 model (Rombach et al., 2022) on FFHQ, and ii) the Stable Diffusion v1.5 (Rombach et al., 2022) model on ImageNet.

Table 2: **EquiReg improves SITCOM under reduced measurement consistency steps ( $K_{\text{meas}}$ ).** We reduce  $K_{\text{meas}}$  and add an equal amount of EquiReg steps ( $K_{\text{EquiReg}}$ ). Evaluated on motion deblur for FFHQ sampled with 50 DDIM steps.

$K_{\text{meas}}$	$K_{\text{EquiReg}}$	PSNR↑	SSIM↑	Runtime (s)
10	N/A	28.06	0.81	21.57
5	5	<b>29.26</b>	<b>0.83</b>	<b>11.09</b>
20	N/A	27.04	0.79	38.85
10	10	<b>28.93</b>	<b>0.82</b>	<b>20.92</b>
30	N/A	27.79	0.80	58.84
15	15	<b>29.63</b>	<b>0.84</b>	<b>30.19</b>
40	N/A	<b>30.40</b>	<b>0.85</b>	78.08
20	20	29.50	0.83	<b>41.02</b>
60	N/A	28.35	0.81	108.57
30	30	<b>31.36</b>	<b>0.87</b>	<b>59.38</b>

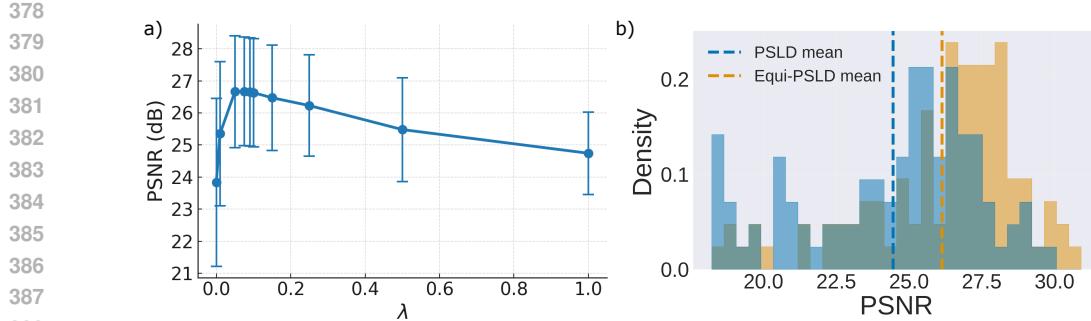


Figure 6: **Robustness of EquiReg, demonstrated on PSLD.** a) EquiReg is robust to the choice of  $\lambda_t$ . b) EquiReg reduces failure cases and enhances reconstruction fidelity for super-resolution on FFHQ.

We evaluate EquiReg on a variety of linear and nonlinear restoration tasks for natural images (see Section E for task details). We adopt the pre-trained encoder-decoder  $\mathcal{E}\text{-}\mathcal{D}$  as our MPE function. For FFHQ, we use vertical reflection as the symmetry group, which preserves upright facial orientation. For ImageNet, we define a rotation group  $G = \{0, \pi/2, \pi, 3\pi/2\}$ , and uniformly at random select the group action for each sample. Finally, the loss functions given in Equation (5) are used to regularize. We note that while our main experiment explore the reflection and rotation groups with small cardinality, EquiReg does not rely on full group coverage. Sampling even a sparse or randomly chosen subset of group actions is sufficient, as long as the function used for regularization exhibits the MPE property across the group (see additional experiment in Appendix).

First, we show that adding EquiReg optimization steps consistently enables SIT-COM to achieve superior performance with significantly faster runtime using fewer measurement consistency steps (Table 2). Next, we show that EquiReg maintains strong performance even as the number of DDIM steps is reduced, whereas DPS suffers a significant drop; Equi-DPS consistently outperforms DPS, with the performance gap widening at lower step counts (Figure 13). We also show that EquiReg is able to preserve performance when applied with lower frequency (Table 1).

**Table 3, Table 4a, and Table 5 highlights the benefits of EquiReg for latent diffusion models by consistently improving the performance of ReSample and PSLD across several tasks on both FFHQ and ImageNet.** We attribute this improvement in part to the reduction of failure cases (Figure 6b). EquiReg also significantly improves the performance of pixel-based methods (see Equi-DPS vs. DPS, Table 4b).

We observe that EquiReg achieves its largest improvements on perceptual metrics (FID and LPIPS), suggesting it generates more realistic images that lie closer to the data manifold (see Appendix E for supporting qualitative results). EquiReg improves performance under high measurement noise (Figure 5b). This result aligns with Figure 5a, which shows the equivariance error is lower on clean images than noisy ones, indicating that EquiReg enforces an effective denoising. Lastly, we note that EquiReg is robust to regularizing hyperparameter  $\lambda_t$  (Figure 6a, see Section C for details).

Table 4: **EquiReg for diffusion models on FFHQ**.  $256 \times 256$  with  $\sigma_y = 0.05$ .

Method	Gaussian deblur	Motion deblur	Super-resolution ( $\times 4$ )	Box inpainting	Random inpainting
	LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑
PSLD	0.357 106.2 22.87	<b>0.322 84.62</b>	24.25	0.313 89.72 24.51	0.158 43.02 <b>24.22</b> 0.246 49.77 29.05
Equi-PSLD	0.344 94.09 <b>24.42</b>	0.338 99.14 <b>24.83</b>	0.289 90.88 <b>26.32</b>	0.098 <b>31.54</b> 24.19	<b>0.188</b> 41.61 <b>30.43</b>
EquiCon-PSLD	<b>0.320 83.18</b> 24.38	<b>0.322 89.87</b> <b>25.14</b>	<b>0.277 79.39</b> 26.14	<b>0.092</b> 35.07 <b>24.26</b>	0.204 <b>40.75</b> 29.99

(a) Latent diffusion.

Method	Gaussian deblur	Motion deblur	Super-resolution ( $\times 4$ )	Box inpainting	Random inpainting
	LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑	LPIPS↓ FID↓ PSNR↑ LPIPS↓ FID↓ PSNR↑
DPS	0.145 104.8 25.48	0.132 99.75 26.75	0.191 125.4 24.38	0.133 56.89 23.10	0.113 51.32 29.63
Equi-DPS (ours)	<b>0.114 48.76</b> <b>26.32</b>	<b>0.094 41.71</b> <b>28.23</b>	<b>0.120 51.00</b> <b>27.15</b>	<b>0.099</b> 40.47 23.39	<b>0.068</b> 33.65 <b>32.16</b>
DiffStateGrad-DPS	<b>0.128 52.73</b> <b>26.29</b>	<b>0.118 50.14</b> <b>27.61</b>	0.186 73.02 24.65	<b>0.114</b> 47.53 <b>24.10</b>	0.107 49.42 <b>30.15</b>
MCG	0.340 101.2 6.72	0.702 310.5 6.72	0.520 87.64 20.05	0.309 <b>40.11</b> 19.97	0.286 <b>29.26</b> 21.57
MPGD-AE	0.150 114.9 24.42	0.120 104.5 25.72	0.168 137.7 24.01	0.138 248.7 21.59	0.172 339.0 25.22

(b) Pixel-based diffusion.

**Diversity analysis.** To study posterior sampling diversity of EquiReg, we generated  $K = 10$  posterior samples for 20 test images across three inverse problems of box inpainting, Gaussian deblurring,  $4\times$  super-resolution, and measured diversity using two complementary metrics: Intra-LPIPS for perceptual diversity and Pixel-Std for spatial diversity. Table 11 demonstrates that Equi-DPS achieves favorable fidelity-diversity trade-offs. We further investigated diversity scaling by varying box inpainting mask size from  $128 \times 128$  to  $192 \times 192$  pixels (Figure 18). Results show that diversity metrics increase linearly with task difficulty, demonstrating that Equi-DPS naturally expands sampling as problems become more ill-posed rather than artificially constraining solutions. This linear relationship indicates healthy, predictable posterior sampling behavior across the difficulty spectrum. Lastly, Figure 7 provides qualitative confirmation through visual examples showing four posterior samples per image. Observable variations in facial features, expressions, and eye gaze validate our quantitative measurements, confirming EquiReg can generate genuinely diverse reconstructions rather than collapsing to a single solution.

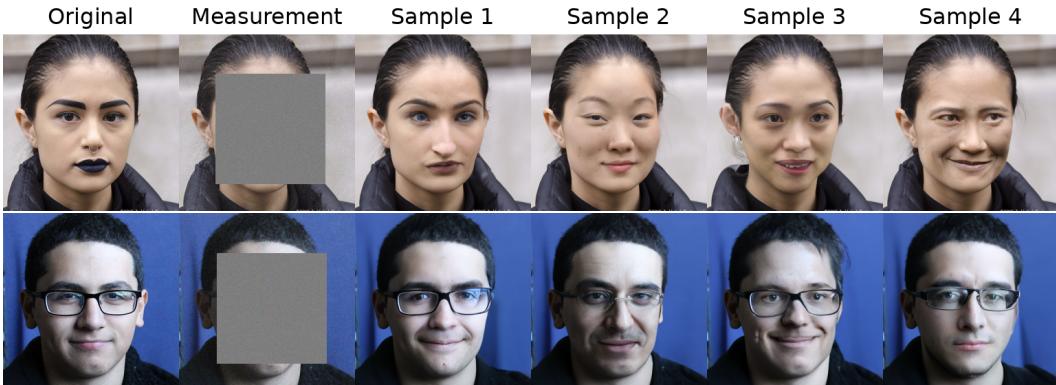


Figure 7: **Qualitative diversity examples for box inpainting.** We show  $K = 4$  posterior samples for two test images with  $160 \times 160$  masks. Each sample exhibits perceptually distinct facial features (expressions, eye gaze, facial structure) while maintaining high fidelity to the ground truth, demonstrating EquiReg generates diverse plausible reconstructions rather than collapsing to a single mode.

**Solving PDEs from sparse observations.** EquiReg is evaluated on two important PDE problems: the Helmholtz and Navier-Stokes equations (see Section F). The objective is to solve both forward and inverse problems in sparse sensor settings. The forward problem involves predicting the solution function or the final state using measurements from only 3% of the coefficient field or the initial state. The inverse problem, conversely, aims to predict the input conditions from observations of 3% of the system’s output. This task is challenging due to the nonlinearity of the equations, the complex structure of Gaussian random fields, and the sparsity of observations.

486 Table 5: **EquiReg for latent diffusion models on ImageNet**.  $256 \times 256$  with  $\sigma_y = 0.05$ .  
487

Method	Gaussian deblur		Motion deblur		Super-resolution (x4)		Box inpainting		Random inpainting	
	FID↓	PSNR↑	FID↓	PSNR↑	FID↓	PSNR↑	FID↓	PSNR↑	FID↓	PSNR↑
PSLD	263.9	20.70	252.1	21.26	224.3	22.29	151.4	16.28	83.22	26.56
EquiCon-PSLD	<b>214.5</b>	<b>22.01</b>	<b>196.3</b>	<b>22.69</b>	<b>198.5</b>	<b>22.34</b>	<b>137.6</b>	<b>19.25</b>	<b>65.14</b>	<b>27.03</b>

493 Recent studies (Huang et al., 2024; Mammadov et al., 2024a; Yao et al., 2025) have demonstrated  
494 the superiority of diffusion models over deterministic single-forward methods for solving PDEs.  
495 DiffusionPDE (Huang et al., 2024) decomposes the conditional log-likelihood into a learned diffusion  
496 prior and a measurement score. FunDPS (Yao et al., 2025) extends the sampling process to a more  
497 natural infinite-dimensional spaces, achieving better accuracy and speed via function space models.  
498

499 We integrate EquiReg into the state-of-  
500 the-art FunDPS framework (Mammadov  
501 et al., 2024a; Yao et al., 2025), where  
502 we compute the Equi loss with respect to  
503 equivariance learned by an FNO trained  
504 on the corresponding inverse problem.  
505 We use reflection symmetry (i.e., flipping  
506 along the  $y = x$  axis), and observe no  
507 significant performance difference when using other transformations such as rotations or alternating  
508 flips. Equi-FunDPS improves performance (Table 6), measured by relative  $\ell_2$  loss, across various  
509 tasks, especially in inverse problems where a strong data prior is critical.

510 **Text-to-image guidance.** Given the “source” image, DreamSampler (Kim et al., 2024) transforms  
511 the source image using the prompt. Applying EquiReg to DreamSampler, we observe perceptual  
512 improvement of generated images as well as artifact reduction. Figure 1 shows the “source” cat, being  
513 transformed into the prompt (e.g., “corgi”). Equi-DreamSampler generates more realistic images  
514 than DreamSampler. Notably, EquiReg resolves the three-front-legged corgi into a two-front-legged  
515 one (for an implicit acceleration of image generation when EquiReg is imposed, see Section A).

## 516 5 CONCLUSION

517 We introduce *Equivariance Regularized* (EquiReg) diffusion for inverse problems. EquiReg regularizes  
518 sampling trajectories to stay closer to the data manifold, leveraging manifold-preferential  
519 equivariance (MPE): functions with low equivariance error on-manifold and high error off-manifold.  
520 Such functions arise naturally in trained networks and can serve as plug-and-play regularizers without  
521 modifying the diffusion denoiser. EquiReg is agnostic across pixel- and latent-space diffusion models  
522 and remains robust under reduced sampling, effectively accelerating convergence. Across diverse  
523 inverse problems, it consistently improves perceptual and reconstruction metrics while reducing  
524 failure cases, highlighting its generality and efficiency.

525 **Limitations and future work.** EquiReg’s effectiveness depends on the quality of the pre-trained  
526 backbone diffusion. EquiReg is a plug-and-play regularization framework that can be applied to a  
527 variety of guidance-based diffusion models; thus, it does not directly address the approximations  
528 of the underlying diffusion models, but instead regularizes for improved performance. Also, since  
529 EquiReg is a regularization mechanism, it improves performance precisely in regimes where baseline  
530 methods degrade or fail. Hence, one cannot expect EquiReg to improve the performance of a  
531 diffusion model beyond the capability of a regularization framework. Finally, applying EquiReg  
532 requires task-specific design choices: selecting an appropriate symmetry group and identifying  
533 suitable MPE functions for the problem at hand. While we presented two systematic approaches to  
534 construct MPE functions for imaging and PDEs, the process of identifying MPE functions varies  
535 across applications and represents an important area for methodological development, which we have  
536 provided guidelines for in this paper. This task-specific design also makes EquiReg broadly adaptable  
537 across diverse domains beyond the considered applications. Finally, while our paper formalizes  
538 distribution-dependent equivariant functions and MPE functions, a full theoretical characterization  
539 of the conditions under which MPE properties emerge in trained networks or its joint training with  
diffusion remains an important and valuable direction for future work.

540 REFERENCES  
541

542 Charles W Groetsch. *Inverse problems in the mathematical sciences*, volume 52. Springer, 1993.

543 SI Kabanikhin. Definitions and examples of inverse and ill-posed problems. *Journal of Inverse and*  
544 *Ill-Posed Problems*, 16(4):317–357, 2008.

545

546 Andrew M Stuart. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.

547

548 David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306,  
549 2006.

550 Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis?  
551 *Journal of the ACM (JACM)*, 58(3):1–37, 2011.

552

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

555 Yang Song and Stefano Ermon. Generative Modeling by Estimating Gradients of the Data Distribution.  
556 In *Advances in Neural Information Processing Systems*, volume 32, 2019.

557

558 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye.  
559 Diffusion Posterior Sampling for General Noisy Inverse Problems. In *The Eleventh International*  
560 *Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=OnD9zGAGT0k>.

561

562 Hyungjin Chung, Eun Sun Lee, and Jong Chul Ye. MR Image Denoising and Super-Resolution Using  
563 Regularized Reverse Diffusion. *IEEE Transactions on Medical Imaging*, 42(4):922–934, 2022a.

564

565 Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. DiffusionPDE: Generative PDE-  
566 solving under partial observation. In *The Thirty-eighth Annual Conference on Neural Information*  
567 *Processing Systems*, 2024.

568

569 Jiachen Yao, Abbas Mammadov, Julius Berner, Gavin Kerrigan, Jong Chul Ye, Kamyar Azizzade-  
570 nesheli, and Anima Anandkumar. Guided diffusion sampling on function spaces with applications  
571 to pdes, 2025.

572

573 Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural*  
574 *Computation*, 23(7):1661–1674, 2011. doi: 10.1162/NECO\_a\_00142.

575

576 Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song.  
577 Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of*  
578 *the Computer Vision and Pattern Recognition Conference*, pages 20895–20905, 2025a.

579

580 Ismail Alkhouri, Shijun Liang, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar, and  
581 Rongrong Wang. Sitcom: Step-wise triple-consistent diffusion sampling for inverse problems. In  
582 *Forty-second International Conference on Machine Learning*, 2025.

583

584 Herbert Ellis Robbins. An empirical bayes approach to statistics. In *Proc. Third Berkley Symposium*  
585 *on Mathematical Statistics*, pages 157–163, 1956.

586

587 Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai.  
588 Solving Linear Inverse Problems Provably via Posterior Sampling with Latent Diffusion Models.  
589 In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=XKBFdYwfRo>.

590

591 Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-  
592 Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, and Stefano Ermon. Manifold  
593 preserving guided diffusion. In *International Conference on Learning Representations*, 2024.

594

595 Rayhan Zirvi, Bahareh Tolooshams, and Anima Anandkumar. Diffusion state-guided projected gradi-  
596 ent for inverse problems. In *The Thirteenth International Conference on Learning Representations*,  
597 2025. URL <https://openreview.net/forum?id=kRBQwlkFSP>.

594 Benjamin Boys, Mark Girolami, Jakub Pidstrigach, Sebastian Reich, Alan Mosca, and Omer Deniz  
 595 Akyildiz. Tweedie moment projected diffusions for inverse problems. *Transactions on Machine*  
 596 *Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=4unJi0qrTE>. Featured Certification.

597

598 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
 599 *the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.

600

601 Dongdong Chen, Mike Davies, Matthias J Ehrhardt, Carola-Bibiane Schönlieb, Ferdia Sherry, and  
 602 Julián Tachella. Imaging with equivariant deep learning: From unrolled network design to fully  
 603 unsupervised learning. *IEEE Signal Processing Magazine*, 40(1):134–147, 2023a.

604

605 Matthieu Terris, Thomas Moreau, Nelly Pustelnik, and Julian Tachella. Equivariant plug-and-play  
 606 image reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
 607 *Pattern Recognition*, pages 25255–25264, 2024.

608

609 Benjamin Bloem-Reddy, Yee Whye, et al. Probabilistic symmetries and invariant neural networks.  
*Journal of Machine Learning Research*, 21(90):1–61, 2020.

610

611 Dian Wang, Stephen Hart, David Surovik, Tarik Kelestemur, Haojie Huang, Haibo Zhao, Mark  
 612 Yeatman, Jiuguang Wang, Robin Walters, and Robert Platt. Equivariant diffusion policy. In *8th*  
 613 *Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=wD2kUVLT1g>.

614

615 Hannah Lawrence, Vasco Portilheiro, Yan Zhang, and Sékou-Oumar Kaba. Improving equivariant  
 616 networks with probabilistic symmetry breaking. *International Conference on Learning Representations*, 2025.

617

618 Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving Inverse  
 619 Problems with Latent Diffusion Models via Hard Data Consistency. In *Conference on Parsimony*  
 620 *and Learning (Recent Spotlight Track)*, 2023a. URL <https://openreview.net/forum?id=iHcarDCZLn>.

621

622 Abbas Mammadov, Julius Berner, Kamyar Azizzadenesheli, Jong Chul Ye, and Anima  
 623 Anandkumar. Diffusion-based inverse solver on function spaces with applications to  
 624 pdes. *Machine Learning and the Physical Sciences Workshop at NeurIPS*, 2024a. URL  
 625 [https://ml4physicalsciences.github.io/2024/files/NeurIPS\\_ML4PS\\_2024\\_253.pdf](https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_253.pdf).

626

627 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised  
 628 learning using nonequilibrium thermodynamics. In *International conference on machine learning*,  
 629 pages 2256–2265. pmlr, 2015.

630

631 Zahra Kadkhodaie and Eero Simoncelli. Stochastic solutions for linear inverse problems using the  
 632 prior implicit in a denoiser. *Advances in Neural Information Processing Systems*, 34:13242–13254,  
 633 2021.

634

635 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam  
 636 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for  
 637 high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

638

639 Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe  
 640 Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI*  
*Blog*, 1:8, 2024.

641

642 Bingliang Zhang, Zihui Wu, Berthy T Feng, Yang Song, Yisong Yue, and Katherine L Bouman.  
 643 Step: A general and scalable framework for solving video inverse problems with spatiotemporal  
 644 diffusion priors. *preprint arXiv:2504.07549*, 2025b.

645

646 Yang Song, Jascha Sohl-Dickstein, Diederik Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 647 Poole. Score-based Generative Modeling through Stochastic Differential Equations. In *The*  
*International Conference on Learning Representations*, 2021. URL <https://openreview.net/pdf/ef0eadbe07115b0853e964f17aa09d811cd490f1.pdf>.

648 Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their*  
 649 *Applications*, 12(3):313–326, 1982.  
 650

651 Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching.  
 652 *Journal of Machine Learning Research*, 6(4), 2005.  
 653

654 Xinyu Peng, Ziyang Zheng, Wenrui Dai, Nuoqian Xiao, Chenglin Li, Junni Zou, and Hongkai Xiong.  
 655 Improving diffusion models for inverse problems using optimal posterior covariance. In *Forty-first*  
 656 *International Conference on Machine Learning*, 2024.  
 657

658 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-Guided Diffusion  
 659 Models for Inverse Problems. 2023b. URL [https://openreview.net/forum?id=9\\_gSMA8MRKQ](https://openreview.net/forum?id=9_gSMA8MRKQ).  
 660

661 Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool.  
 662 Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF*  
 663 *Conference on Computer Vision and Pattern Recognition*, pages 1219–1229, 2023.  
 664

665 Koichi Miyasawa et al. An empirical bayes estimator of the mean of a normal population. *Bull. Inst.*  
 666 *Internat. Statist.*, 38(181–188):1–2, 1961.  
 667

668 Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*,  
 669 106(496):1602–1614, 2011.  
 670

671 Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning:  
 672 Grids, groups, graphs, geodesics, and gauges. *preprint arXiv:2104.13478*, 2021.  
 673

674 Víctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks.  
 675 In *International conference on machine learning*, pages 9323–9332. PMLR, 2021.  
 676

677 Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference*  
 678 *on machine learning*, pages 2990–2999. PMLR, 2016.  
 679

680 David W Romero and Suhas Lohit. Learning partial equivariances from data. *Advances in Neural*  
 681 *Information Processing Systems*, 35:36466–36478, 2022.  
 682

683 Marc Finzi, Samuel Stanton, Pavel Izmailov, and Andrew Gordon Wilson. Generalizing convolutional  
 684 neural networks for equivariance to lie groups on arbitrary continuous data. In *International*  
 685 *Conference on Machine Learning*, pages 3165–3176. PMLR, 2020.  
 686

687 Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion  
 688 for molecule generation in 3d. In *International conference on machine learning*, pages 8867–8887,  
 689 2022.  
 690

691 François Cornet, Grigory Bartosh, Mikkel Schmidt, and Christian Andersson Naesseth. Equivariant  
 692 neural diffusion for molecule generation. *Advances in Neural Information Processing Systems*, 37:  
 693 49429–49460, 2024.  
 694

695 Kehua Chen, Xianda Chen, Zihan Yu, Meixin Zhu, and Hai Yang. Equidiff: A conditional equivariant  
 696 diffusion model for trajectory prediction. In *2023 IEEE 26th International Conference on Intelligent*  
 697 *Transportation Systems (ITSC)*, pages 746–751. IEEE, 2023b.  
 698

699 Johann Brehmer, Joey Bose, Pim De Haan, and Taco S Cohen. Edgi: Equivariant diffusion for  
 700 planning with embodied agents. *Advances in Neural Information Processing Systems*, 36:63818–  
 701 63834, 2023.  
 702

703 Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal  
 704 structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing*  
 705 *Systems*, 36:17464–17497, 2023.  
 706

707 Eloi Moliner, Jaakko Lehtinen, and Vesa Välimäki. Solving audio inverse problems with a diffusion  
 708 model. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal*  
 709 *Processing (ICASSP)*, pages 1–5, 2023.  
 710

702 Giannis Daras, Weili Nie, Karsten Kreis, Alex Dimakis, Morteza Mardani, Nikola Kovachki, and  
 703 Arash Vahdat. Warped diffusion: Solving video inverse problems with image diffusion models.  
 704 *Advances in Neural Information Processing Systems*, 37:101116–101143, 2024.

705 Jérémie Scanvic, Mike Davies, Patrice Abry, and Julián Tachella. Scale-equivariant imaging: Self-  
 706 supervised learning for image super-resolution and deblurring. 2025.

708 Julián Tachella, Dongdong Chen, and Mike Davies. Sensing theorems for unsupervised learning in  
 709 linear inverse problems. *Journal of Machine Learning Research*, 24(39):1–45, 2023.

710 Rui Wang, Robin Walters, and Rose Yu. Approximately equivariant networks for imperfectly  
 711 symmetric dynamics. In *International Conference on Machine Learning*, pages 23078–23091.  
 712 PMLR, 2022.

714 Lawrence Cayton et al. *Algorithms for manifold learning*. Univ. of California at San Diego Tech.  
 715 Rep, 2005.

716 Sebastian Bordt, Uddeshya Upadhyay, Zeynep Akata, and Ulrike von Luxburg. The manifold  
 717 hypothesis for gradient-based explanations. In *Proceedings of the IEEE/CVF Conference on*  
 718 *Computer Vision and Pattern Recognition*, pages 3697–3702, 2023.

720 Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving Diffusion Models  
 721 for Inverse Problems using Manifold Constraints. *Advances in Neural Information Processing*  
 722 *Systems*, 35:25683–25696, 2022b.

723 Kilian Q Weinberger and Lawrence K Saul. Unsupervised learning of image manifolds by semidefinite  
 724 programming. *International journal of computer vision*, 70:77–90, 2006.

725 Lucas CF Ferreira and Julio C Valencia-Guevara. Gradient flows of time-dependent functionals in  
 726 metric spaces and applications to pdes. *Monatshefte für Mathematik*, 185(2):231–268, 2018.

728 Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance  
 729 and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,  
 730 pages 991–999, 2015.

731 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional  
 732 neural networks. *Advances in neural information processing systems*, 25, 2012.

734 Robert-Jan Bruintjes, Tomasz Motyka, and Jan van Gemert. What affects learned equivariance in  
 735 deep image recognition models? In *Proceedings of the IEEE/CVF Conference on Computer Vision*  
 736 *and Pattern Recognition*, pages 4839–4847, 2023.

737 Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small  
 738 image transformations? *Journal of Machine Learning Research*, 20(184):1–25, 2019.

739 Yibo Zhou. Rethinking reconstruction autoencoder-based out-of-distribution detection. In *Proceed-  
 740 ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7379–7387,  
 741 2022.

743 Ramneet Kaur, Susmit Jha, Anirban Roy, Sangdon Park, Edgar Dobriban, Oleg Sokolsky, and  
 744 Insup Lee. idecode: In-distribution equivariance for conformal out-of-distribution detection. In *Proceed-  
 745 ings of the AAAI conference on artificial intelligence*, volume 36, pages 7104–7114, 2022.

746 Ramneet Kaur, Kaustubh Sridhar, Sangdon Park, Yahan Yang, Susmit Jha, Anirban Roy, Oleg  
 747 Sokolsky, and Insup Lee. Codit: Conformal out-of-distribution detection in time-series data for  
 748 cyber-physical systems. In *Proceedings of the ACM/IEEE 14th International Conference on*  
 749 *Cyber-Physical Systems (with CPS-IoT Week 2023)*, pages 120–131, 2023.

750 Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya,  
 751 Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function  
 752 spaces. *CoRR*, abs/2108.08481, 2021.

754 Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew  
 755 Stuart, Anima Anandkumar, et al. Fourier neural operator for parametric partial differential  
 756 equations. In *International Conference on Learning Representations*, 2021.

756 Allan Zhou, Tom Knowles, and Chelsea Finn. Meta-learning symmetries by reparameterization. In  
 757 *International Conference on Learning Representations*.  
 758

759 Robin Quessard, Thomas Barrett, and William Clements. Learning disentangled representations and  
 760 group structure of dynamical environments. *Advances in Neural Information Processing Systems*,  
 761 33:19727–19737, 2020.

762 Nima Dehmamy, Robin Walters, Yanchen Liu, Dashun Wang, and Rose Yu. Automatic symmetry  
 763 discovery with lie algebra convolutional network. *Advances in Neural Information Processing*  
 764 *Systems*, 34:2503–2515, 2021.

765

766 Jeet Mohapatra, Nima Dehmamy, Csaba Both, Subhro Das, and Tommi Jaakkola. Symmetry-  
 767 driven discovery of dynamical variables in molecular simulations. In *Forty-second International*  
 768 *Conference on Machine Learning*, 2025.

769 Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative  
 770 Adversarial Networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 43(12):  
 771 4217–4228, Dec 2021.

772

773 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-  
 774 Scale Hierarchical Image Database. In *2009 IEEE Conference on Computer Vision and Pattern*  
 775 *Recognition*, pages 248–255. IEEE, 2009.

776 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*  
 777 *in neural information processing systems*, 34:8780–8794, 2021.

778

779 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 780 resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF*  
 781 *conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

782 Jeongsol Kim, Geon Yeong Park, and Jong Chul Ye. Dreamsampler: Unifying diffusion sampling  
 783 and score distillation for image manipulation. In *European Conference on Computer Vision*, pages  
 784 398–414. Springer, 2024.

785

786 Phong Tran, Anh Tran, Quynh Phung, and Minh Hoai. Explore Image Deblurring via Encoded  
 787 Blur Kernel Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 788 *Recognition (CVPR)*. IEEE, 2021.

789

790 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew  
 791 Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations.  
 792 *arXiv preprint arXiv:2010.08895*, 2020.

793

794 Hongkai Zheng, Wenda Chu, Austin Wang, Nikola Borislavov Kovachki, Ricardo Baptista, and  
 795 Yisong Yue. Ensemble kalman diffusion guidance: A derivative-free method for inverse problems.  
 796 *Transactions on Machine Learning Research*, 2025.

797

798 Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the fokker–planck  
 799 equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.

800

801 Christopher A Metzler, Arian Maleki, and Richard G Baraniuk. From denoising to compressed  
 802 sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.

803

804 Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by  
 805 denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.

806

807 Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image  
 808 restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,  
 809 pages 3929–3938, 2017.

810 Chris Metzler, Ali Mousavi, and Richard Baraniuk. Learned d-amp: Principled neural network based  
 811 compressive image recovery. *Advances in neural information processing systems*, 30, 2017.

812

813 Diederik P Kingma. Auto-encoding Variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

810 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 811 Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. *Advances in neural information*  
 812 *processing systems*, 27, 2014.

813

814 Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed Sensing using Generative  
 815 Models. In *International conference on machine learning*, pages 537–546. PMLR, 2017.

816 Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust  
 817 Compressed Sensing MRI with Deep Generative Priors. In *Advances in Neural Information*  
 818 *Processing Systems*, volume 34, pages 14938–14954. Curran Associates, Inc., 2021.

819

820 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration  
 821 models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022.

822 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.  
 823 Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the*  
 824 *IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022.

825

826 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi.  
 827 Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine*  
 828 *intelligence*, 45(4):4713–4726, 2022.

829

830 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving Inverse Problems in Medical Imaging  
 831 with Score-Based Generative Models. In *International Conference on Learning Representations*,  
 832 2022.

833

834 Hyungjin Chung and Jong Chul Ye. Score-Based Diffusion Models for Accelerated MRI. *Medical*  
 835 *Image Analysis*, page 102479, 2022.

836

837 Alex Ling Yu Hung, Kai Zhao, Haoxin Zheng, Ran Yan, Steven S Raman, Demetri Terzopoulos, and  
 838 Kyunghyun Sung. Med-CDiff: Conditional Medical Image Generation with Diffusion Models.  
 839 *Bioengineering*, 10(11):1258, 2023.

840

841 Zolnamar Dorjsembe, Hsing-Kuo Pao, Sodtavilan Odonchimed, and Furen Xiao. Conditional  
 842 Diffusion Models for Semantic 3D Brain MRI Synthesis. *IEEE Journal of Biomedical and Health*  
 843 *Informatics*, 2024.

844

845 Guangyuan Li, Chen Rao, Juncheng Mo, Zhanjie Zhang, Wei Xing, and Lei Zhao. Rethinking diffu-  
 846 sion model for multi-contrast mri super-resolution. In *Proceedings of the IEEE/CVF Conference*  
 847 *on Computer Vision and Pattern Recognition*, pages 11365–11374, 2024.

848

849 Amirkhossein Kazerouni, Ehsan Khodapanah Aghdam, Moein Heidari, Reza Azad, Mohsen Fayyaz,  
 850 Ilker Hacihaliloglu, and Dorit Merhof. Diffusion models in medical imaging: A comprehensive  
 851 survey. *Medical image analysis*, 88:102846, 2023.

852

853 Wanyu Bian, Albert Jang, Liping Zhang, Xiaonan Yang, Zachary Stewart, and Fang Liu. Diffusion  
 854 modeling with domain-conditioned prior guidance for accelerated mri and qmri reconstruction.  
 855 *IEEE Transactions on Medical Imaging*, 2024.

856

857 Victor Isakov. *Inverse problems for Partial Differential Equations*, volume 127. Springer, 2006.

858

859 Aliaksandra Shysheya, Cristiana Diaconu, Federico Bergamin, Paris Perdikaris, José Miguel  
 860 Hernández-Lobato, Richard Turner, and Emile Mathieu. On conditional diffusion models for pde  
 861 simulations. *Advances in Neural Information Processing Systems*, 37:23246–23300, 2024.

862

863 Ziming Liu, Di Luo, Yilun Xu, Tommi Jaakkola, and Max Tegmark. Genphys: From physical  
 864 processes to generative models. *arXiv preprint arXiv:2304.02637*, 2023.

865

866 Zijie Li, Anthony Zhou, and Amir Barati Farimani. Generative latent neural pde solver using flow  
 867 matching. *arXiv preprint arXiv:2503.22600*, 2025.

868

869 Lorenzo Baldassari, Ali Siahkoohi, Josselin Garnier, Knut Solna, and Maarten V de Hoop. Conditional  
 870 score-based diffusion models for bayesian inference in infinite dimensions. *Advances in Neural*  
 871 *Information Processing Systems*, 36:24262–24290, 2023.

864 Abbas Mammadov, Hyungjin Chung, and Jong Chul Ye. Amortized posterior sampling with diffusion  
 865 prior distillation, 2024b. URL <https://arxiv.org/abs/2407.17907>.

866

867 Gabriel Cardoso, Yazid Janati el idrissi, Sylvain Le Corff, and Eric Moulines. Monte carlo guided  
 868 denoising diffusion models for bayesian linear inverse problems. In *The Twelfth International  
 869 Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=nHESwXvxWK>.

870

871 Hang Shao, Abhishek Kumar, and P Thomas Fletcher. The riemannian geometry of deep generative  
 872 models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition  
 873 Workshops*, pages 315–323, 2018.

874

875 Christopher Anders, Plamen Pasliev, Ann-Kathrin Dombrowski, Klaus-Robert Müller, and Pan Kessel.  
 876 Fairwashing explanations with off-manifold detergent. In *International Conference on Machine  
 877 Learning*, pages 314–323. PMLR, 2020.

878

879 Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley.  
 880 Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds.  
*arXiv preprint arXiv:1802.08219*, 2018.

881

882 Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-  
 883 translation equivariant attention networks. *Advances in neural information processing systems*, 33:  
 884 1970–1981, 2020.

885

886 Minkai Xu, Jiaqi Han, Aaron Lou, Jean Kossaifi, Arvind Ramanathan, Kamyar Azizzadenesheli,  
 887 Jure Leskovec, Stefano Ermon, and Anima Anandkumar. Equivariant graph neural operator for  
 888 modeling 3d dynamics. In *International Conference on Machine Learning*, pages 55015–55032.  
 889 PMLR, 2024.

890

891 Jung Yeon Park, Sujay Bhatt, Sihan Zeng, Lawson LS Wong, Alec Koppel, Sumitra Ganesh, and  
 892 Robin Walters. Approximate equivariance in reinforcement learning. In *International Conference  
 893 on Artificial Intelligence and Statistics*, pages 4177–4185. PMLR, 2025.

894

895 Yunqian Ma and Yun Fu. *Manifold learning theory and applications*, volume 434. CRC press Boca  
 896 Raton, 2012.

897

898 Hariharan Narayanan and Sanjoy Mitter. Sample complexity of testing the manifold hypothesis.  
 899 *Advances in neural information processing systems*, 23, 2010.

900

901 Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis.  
 902 *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=MhK5aXo3gB>.

903

904 Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis.  
 905 *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.

906

907 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
 908 recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition  
 909 (CVPR)*, pages 770–778, 2016.

910

911 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 912 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever.  
 913 Learning transferable visual models from natural language supervision. In *Proceedings of the 38th  
 914 International Conference on Machine Learning (ICML)*, 2021.

915

916

917

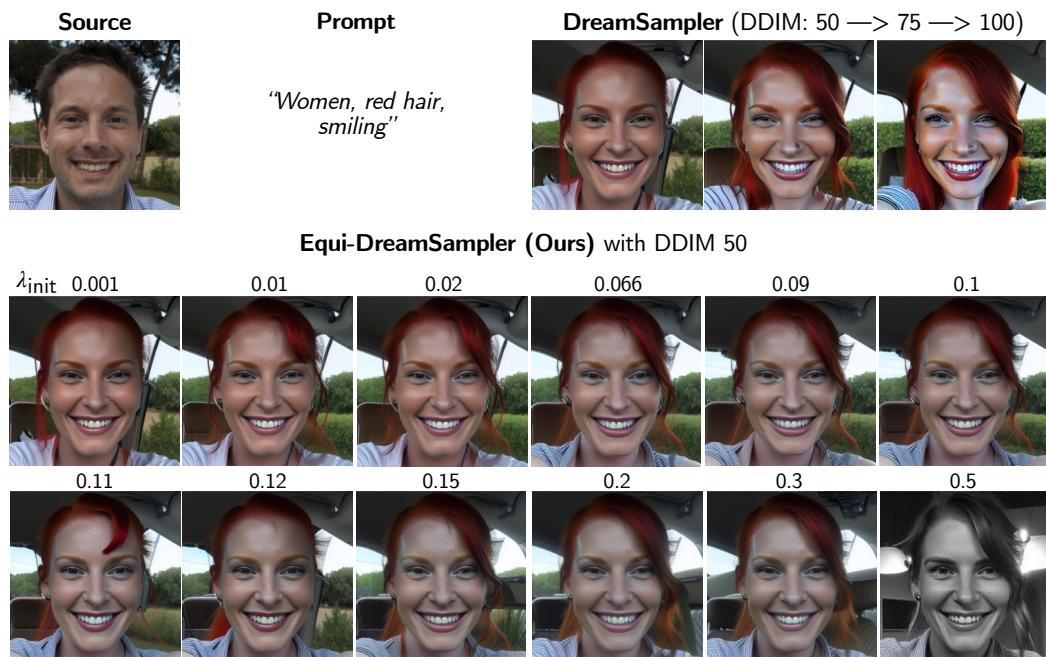
918 APPENDICES FOR “EQUIREG: EQUIVARIANCE REGULARIZED DIFFUSION FOR  
 919 INVERSE PROBLEMS”  
 920

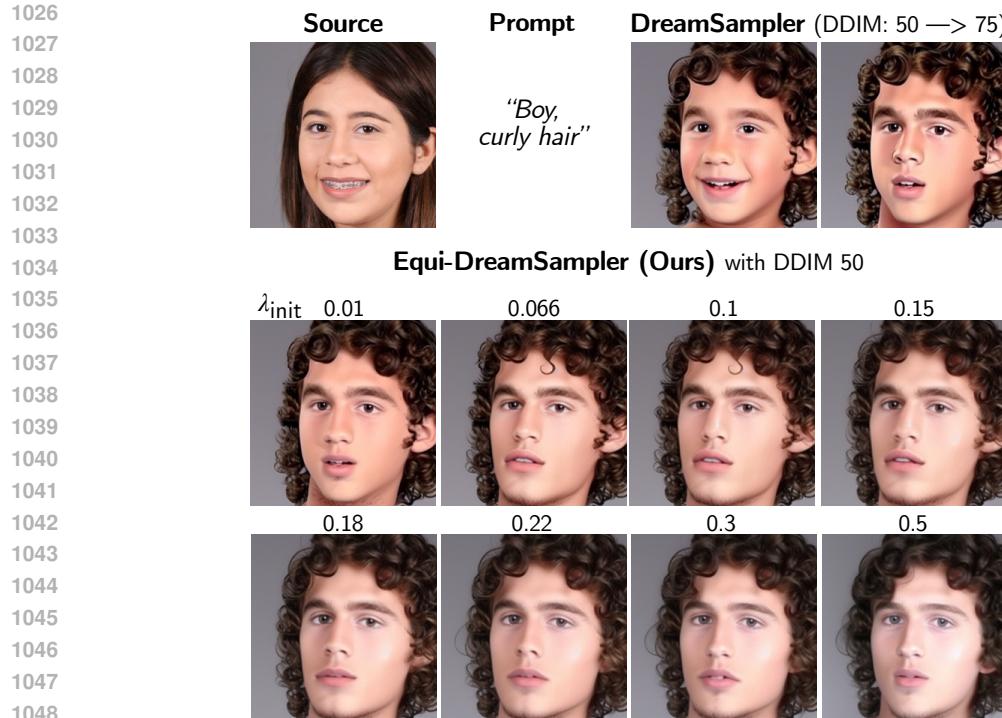
921 We provide our source code when EquiReg. We will provide a publicly available source code upon  
 922 acceptance. This supplementary materials contain the following:  
 923

- 924 • Section A includes additional experiments on text-to-image guidance. We regularize DreamSam-  
 925 pler (Kim et al., 2024) with EquiReg for an improved performance (see Figures 8 to 12).  
 926
- 927 • Section B includes additional experiments on robustness including robustness to  $\lambda_t$ , reduced  
 928 number of DDIM steps, and reduced number of measurement consistency steps.  
 929
- 930 • Section C includes qualitative analysis on the performance of methods with and without EquiReg.  
 931 Results show a reduction of artifacts and an improved perceptual quality of the solution. This  
 932 section also includes the equivariance error of a pre-trained encoder used in EquiReg (Figure 15a)  
 933 and a histogram of Equi’s improvement on DPS (Figure 14).  
 934
- 935 • Section D includes diversity experiments. Results show that EquiReg achieves favorable fidelity-  
 936 diversity tradeoffs (Table 11, Figure 18, and Figure 19).  
 937
- 938 • Section E demonstrates EquiReg experimental setup and implementation for PSLD, ReSample,  
 939 and DPS (Algorithms 2 to 6). It also contains information about the EquiReg hyperparameters for  
 940 image restoration tasks.  
 941
- 942 • Section F contains information on the PDE reconstruction experiment. It discusses the equations  
 943 along with implementation details and hyperparameters.  
 944
- 945 • Section G provides theoretical proofs of Propositions G.1 and G.2.  
 946
- 947 • Section H contains additional background information on solving inverse problems, vanishing-  
 948 error autoencoders, and equivariance.  
 949
- 950 • Section I provides additional experiment on MPE functions.  
 951
- 952 • Section J discloses computing resources used to conduct the experiments.  
 953
- 954 • Section K credits code assets used for our experiments.  
 955
- 956 • Section L discusses the broader impacts of this paper, the developed method, the conducted  
 957 experiments, and their overall implications.  
 958
- 959 • Section M concludes the appendix with a “responsible release” statement.  
 960

961 The authors acknowledge the usage of LLMs on proofreading of the manuscript. The authors have  
 962 not used LLMs for content generation.  
 963

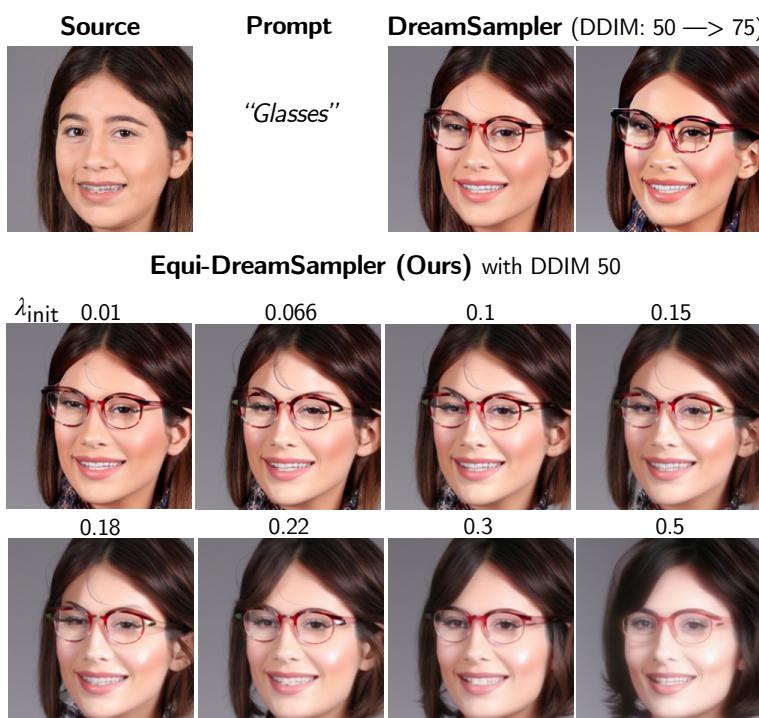
964  
 965  
 966  
 967  
 968  
 969  
 970  
 971

972 **A EQUIREG FOR TEXT-TO-IMAGE GUIDANCE**  
973  
974  
975  
976  
977  
978979 Given the “source” image, DreamSampler (Kim et al., 2024) is asked to transform the source image  
980 using the prompt. Applying EquiReg to DreamSampler, we observe perceptual improvement of  
981 generated images as well as artifact reduction.  
982  
983  
984  
985  
986  
9871009 **Figure 8: Impact of EquiReg parameter  $\lambda_t$ , implicit acceleration, and introduction of more  
1010 image details. Women, red hair, smiling.**1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019 We have observed an implicit acceleration of image generation when EquiReg is imposed (Figure 8).  
1020 Equi-DreamSampler with 50 DDIM steps can generate images that are only possible with DreamSam-  
1021 pler when the DDIM steps are increased. We attribute this to EquiReg’s ability to generate images  
1022 that are closer to the data manifold. For example, the increase of DDIM steps in DreamSampler  
1023 (from 50 to 75 to 100) has a relatively similar effect to increasing the EquiReg regularizer  $\lambda_t$  at a  
1024 fixed 50 DDIM steps. Figure 8 shows that increasing the regularization  $\lambda_t$  results in addition of a car  
1025 in the background. For DreamSampler, an early notion of the car seat in the background start to arise  
only when DDIM is increased to 100 (see also Figures 9 to 12).

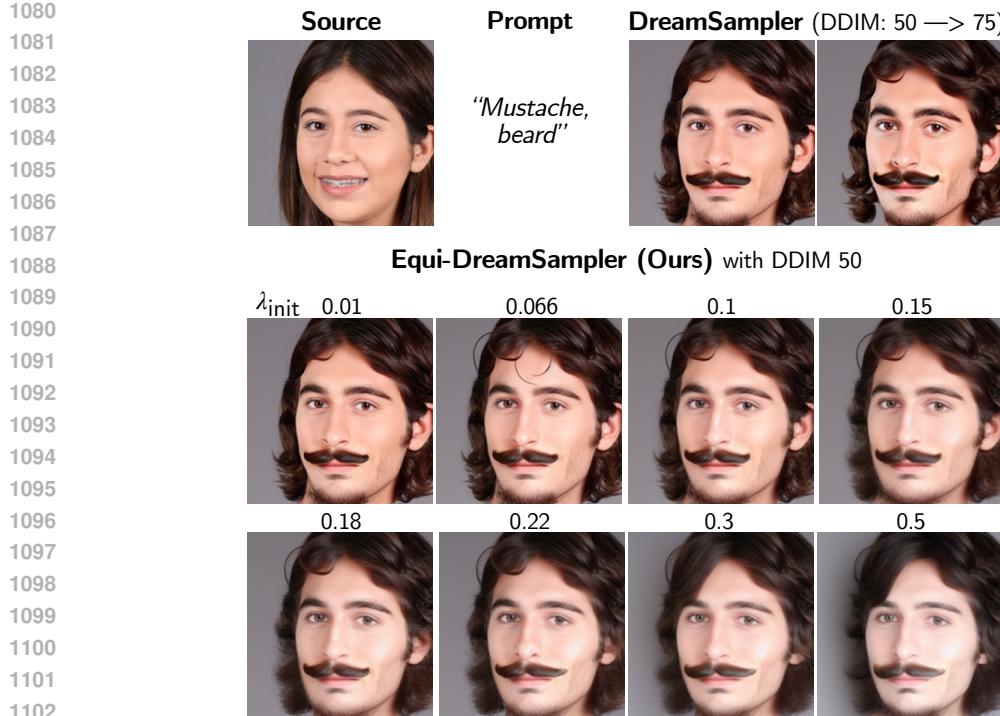


1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

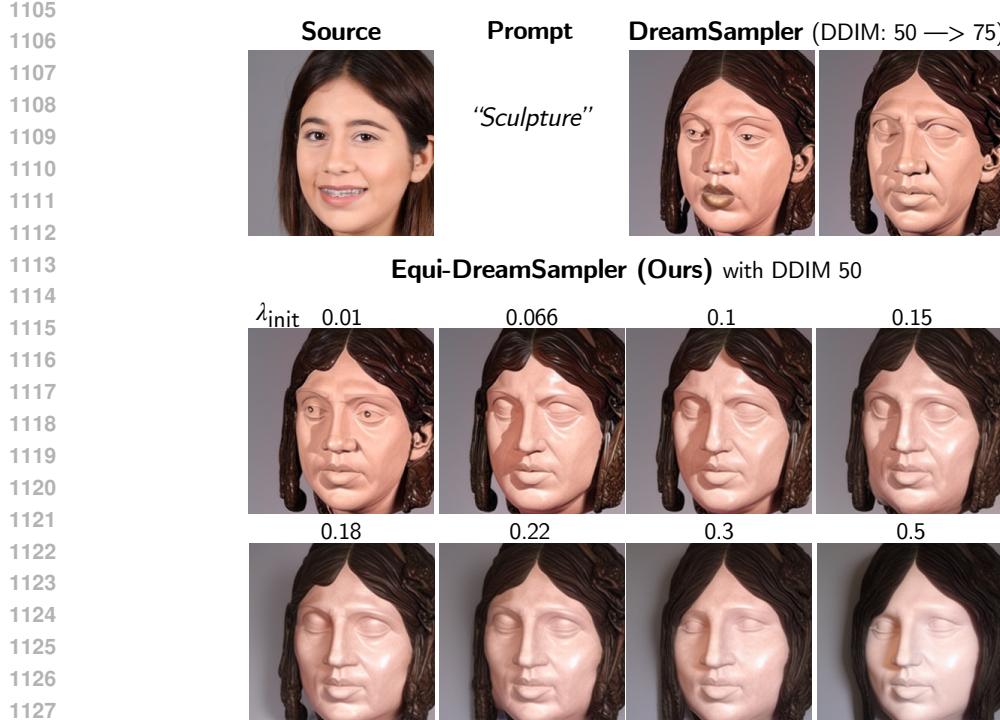
**Figure 9: Adding EquiReg into the text-to-image guidance method DreamSampler for improved performance. Boy, curly hair.**



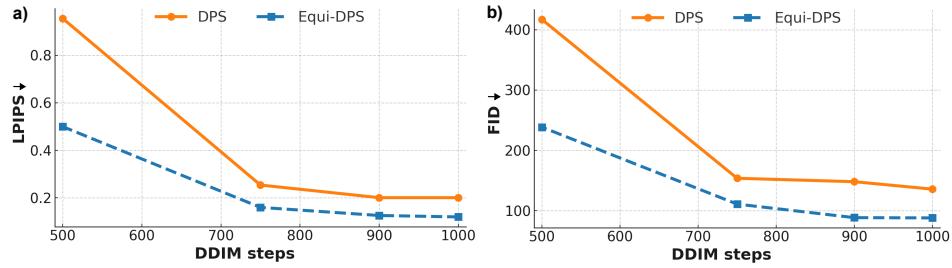
**Figure 10: Adding EquiReg into the text-to-image guidance method DreamSampler for improved performance. Glasses.**



1103 **Figure 11: Adding EquiReg into the text-to-image guidance method DreamSampler for im-  
1104 proved performance. Mustache, beard.**



1128 **Figure 12: Adding EquiReg into the text-to-image guidance method DreamSampler for im-  
1129 proved performance. Sculpture.**

1134 **B ADDITIONAL EXPERIMENTS ON ROBUSTNESS**  
11351145 Figure 13: **Advantages of EquiReg under reduced DDIM steps.** Super-resolution on FFHQ.  
11461147 Table 7: **Robustness to  $\lambda_t$ .** Sensitivity analysis for DPS and PSLD across different values of  $\lambda_t$ .  
1148

$\lambda_t$	DPS			PSLD		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
0.0	24.34 (1.03)	0.664 (0.061)	0.156 (0.051)	23.83 (2.61)	0.63 (0.12)	0.315 (0.07)
0.001	25.44 (1.22)	0.708 (0.057)	0.118 (0.038)	—	—	—
0.01	25.44 (1.22)	0.708 (0.057)	0.118 (0.038)	25.35 (2.24)	0.70 (0.09)	0.280 (0.07)
0.1	25.44 (1.22)	0.708 (0.057)	0.118 (0.038)	26.63 (1.68)	0.74 (0.08)	0.337 (0.06)
0.25	—	—	—	26.22 (1.57)	0.72 (0.08)	0.366 (0.05)
1.0	25.44 (1.22)	0.709 (0.057)	0.118 (0.038)	24.74 (1.28)	0.66 (0.07)	0.438 (0.05)

1158 Table 8: **EquiReg improves performance under reduced DDIM steps.** Pixel-based super-resolution  
1159 on FFHQ 256 × 256.  
1160

Steps	DPS				Equi-DPS (ours)			
	PSNR↑	SSIM↑	LPIPS↓	FID↓	PSNR↑	SSIM↑	LPIPS↓	FID↓
500	13.89	0.0937	0.955	417.07	20.61	0.366	0.500	238.51
750	21.77	0.540	0.254	153.74	25.60	0.704	0.160	110.89
900	22.97	0.628	0.201	148.03	26.52	0.755	0.126	88.46
1000	22.99	0.649	0.201	135.71	26.73	0.767	0.120	88.00

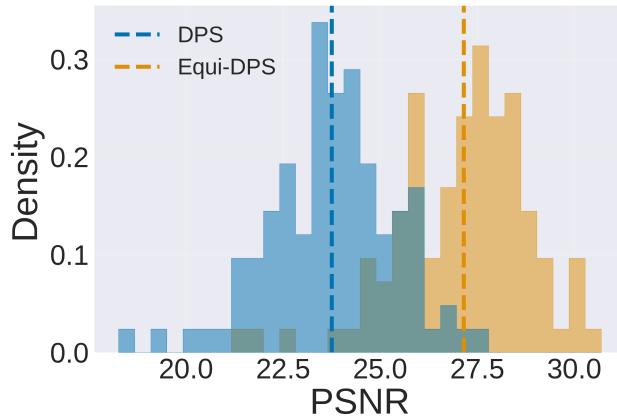
Table 9: **EquiReg improves SITCOM under reduced measurement consistency steps ( $K_{\text{meas}}$ ).**  
 Motion deblur on FFHQ sampled with 50 DDIM steps.

$K_{\text{meas.}}$	$K_{\text{EquiReg}}$	PSNR↑	SSIM↑	Runtime (s)
10	N/A	28.06	0.81	21.57
10	1	28.71	0.82	21.07
5	5	<b>29.26</b>	<b>0.83</b>	<b>11.09</b>
20	N/A	27.04	0.79	38.85
20	1	28.54	<b>0.82</b>	37.74
10	10	<b>28.93</b>	<b>0.82</b>	<b>20.92</b>
30	N/A	27.79	0.80	58.84
30	1	28.35	0.81	55.51
15	15	<b>29.63</b>	<b>0.84</b>	<b>30.19</b>
40	N/A	30.40	<b>0.85</b>	78.08
40	1	<b>30.58</b>	<b>0.85</b>	69.83
20	20	29.50	0.83	<b>41.02</b>
60	N/A	28.35	0.81	108.57
60	1	27.02	0.78	95.62
30	30	<b>31.36</b>	<b>0.87</b>	<b>59.38</b>

Table 10: **EquiReg Effectiveness with Subset of Group Actions.**

PSLD		Equi-PSLD (90, 270 deg)	
PSNR↑	SSIM↑	PSNR↑	SSIM↑
15.86 (1.19)	0.77 (0.03)	17.60 (1.60)	0.79 (0.03)

## C VISUALIZATIONS FOR IMAGE RESTORATION EXPERIMENTS

Figure 14: **Histogram of EquiReg improvement for DPS.** Super-resolution using FFHQ 256 × 256.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

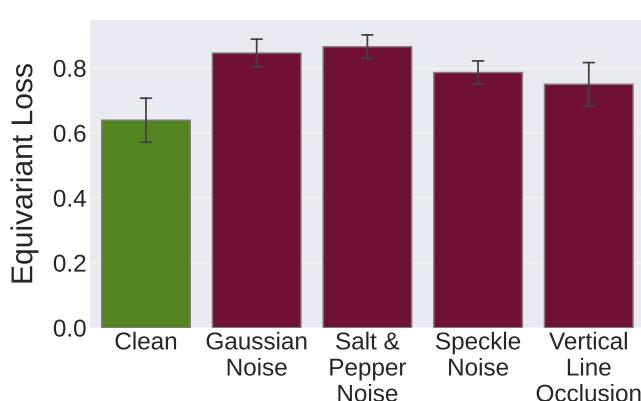
1291

1292

1293

1294

1295

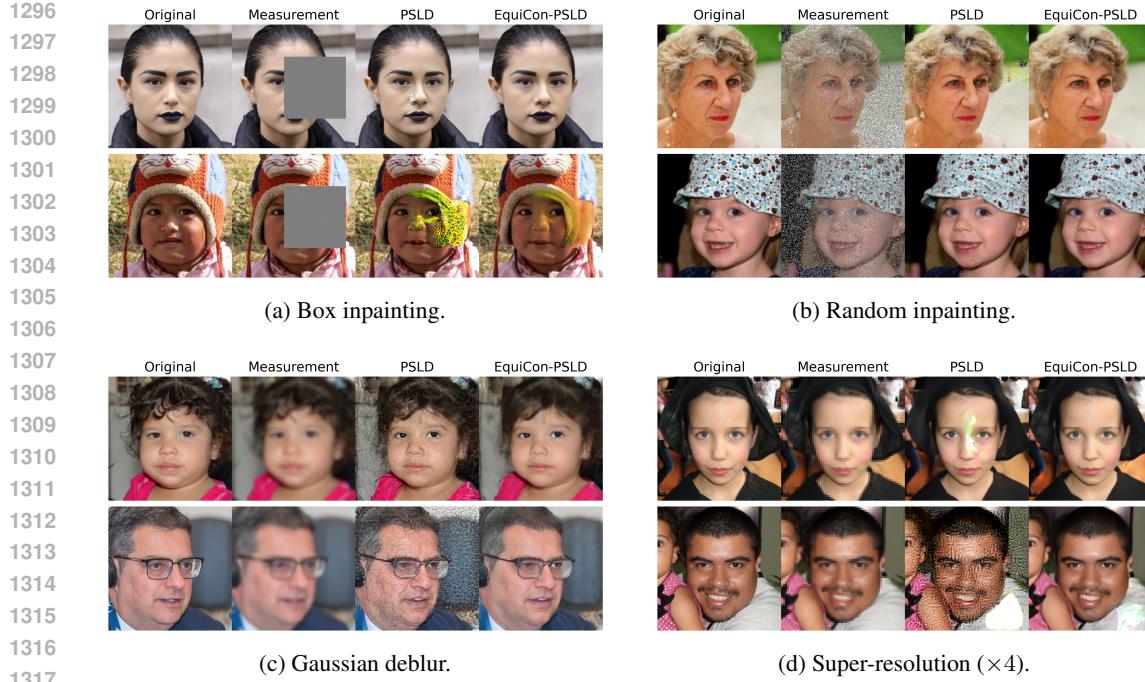
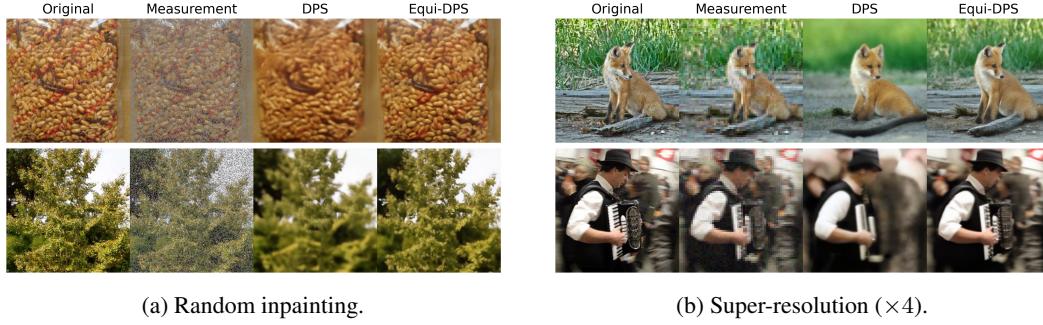


(a) The equivariance error of the encoder is lower on clean, natural images than corrupted ones.



(b) Example visualizations of used images and corresponding equivariance error computed using the decoder (see Figure 4a).

Figure 15: Training induced equivariance for a pre-trained function.

Figure 16: Qualitative comparison of EquiCon-PSLD and PSLD on FFHQ 256  $\times$  256.Figure 17: Qualitative comparison of Equi-DPS and DPS on ImageNet 256  $\times$  256.

## D DIVERSITY ANALYSIS

In the Bayesian setting, the objective of solving inverse problems with diffusion models is to sample from high-probability regions of the posterior distribution. While the goal is not to maximize “diversity”, the true diversity emerges when the posterior admits meaningful variability. In practice, diversity-related concerns in inverse problems arise when a method suffers from mode collapse, i.e., the sampler becomes biased and fails to explore multiple plausible modes of the posterior. Thus, the relevant question is whether a method properly explores the posterior rather than whether it maximizes diversity in an unconstrained sense.

Because closed-form posteriors are unavailable for real image restoration tasks, the standard practice in the diffusion inverse-problem literature is to evaluate diversity through variation among plausible reconstructions consistent with the measurement, without collapsing to a single solution. This is the notion of “diversity” our work adopts.

1350 Given the goal of posterior sampling, EquiReg is not designed to maximize diversity for its own sake.  
 1351 Its objective is to incorporate data-inherent geometric structure (equivariance) to guide sampling  
 1352 toward high-probability regions of the posterior. Hence, diversity arises naturally from the ill-  
 1353 posedness of the inverse problem; it is a consequence of posterior uncertainty, not the goal of the  
 1354 regularizer.

1355 To quantify this effect, in addition to reconstruction quality, we analyzed the diversity of posterior  
 1356 samples produced by EquiReg. We evaluate diversity metrics across multiple tasks and difficulty  
 1357 levels to characterize the sampling behavior of our method.

## 1360 D.1 EXPERIMENTAL SETUP

1361 To evaluate diversity, we generate multiple posterior samples and measure variation across these  
 1362 samples. For each of 20 test images, we generate  $K=10$  reconstructions using different random seeds.  
 1363 We evaluate diversity using two complementary metrics: Intra-LPIPS, which measures perceptual  
 1364 diversity by computing the average LPIPS distance between all pairs of samples, and Pixel-Std, which  
 1365 measures spatial diversity through pixel-wise standard deviation across samples. Higher values for  
 1366 both metrics indicate greater diversity. For Intra-LPIPS, we compute distances for all  $\binom{K}{2} = 45$   
 1367 pairs per image and average across all test images. For Pixel-Std, we compute the standard deviation  
 1368 at each pixel location across the  $K$  samples, then average across all pixels and test images. We  
 1369 evaluate diversity across three inverse problems (box inpainting, Gaussian deblurring, and  $4\times$  super-  
 1370 resolution) comparing EquiReg against DPS (Chung et al., 2023) without equivariance regularization.  
 1371 To investigate how diversity scales with task difficulty, we additionally vary the inpainting mask size  
 1372 from  $128 \times 128$  (standard) to  $160 \times 160$  to  $192 \times 192$  pixels.

1373

1374

1375

## 1376 D.2 RESULTS AND DISCUSSION

1377

1378

1379 Table 11 shows that Equi-DPS achieves favorable fidelity-diversity trade-offs across three inverse  
 1380 problems. For box inpainting and super-resolution, equivariance regularization improves both fidelity  
 1381 and diversity simultaneously. For Gaussian deblurring, Equi-DPS achieves 15-20% better fidelity  
 1382 while retaining 80-85% of baseline diversity, representing a modest but justified trade-off. These  
 1383 results demonstrate that equivariance constraints do not inherently suppress diversity; rather, they can  
 1384 guide sampling toward regions of higher data fidelity while maintaining posterior exploration.

1385

1386

1387 Table 11: **Fidelity and diversity comparison across inverse problems.** Evaluated on 20 test images  
 1388 with  $K = 10$  samples per image. Equi-DPS improves fidelity while largely preserving or enhancing  
 1389 sampling diversity.

1390

Task	Method	Fidelity Metrics		Diversity Metrics	
		LPIPS↓	FID↓	Intra-LPIPS↑	Pixel-Std↑
Box inpainting	DPS	0.140	70.89	0.112	9.286
	Equi-DPS (ours)	<b>0.112</b>	<b>59.70</b>	<b>0.118</b>	<b>10.59</b>
Gaussian deblur	DPS	0.150	76.71	<b>0.114</b>	<b>6.565</b>
	Equi-DPS (ours)	<b>0.120</b>	<b>63.02</b>	0.092	5.669
Super-resolution ( $\times 4$ )	DPS	0.683	99.11	0.134	7.956
	Equi-DPS (ours)	<b>0.703</b>	<b>87.52</b>	<b>0.187</b>	<b>23.52</b>

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

Figure 18 reveals linear diversity scaling with task difficulty. Diversity metrics grow proportionally  
 with task difficulty, indicating Equi-DPS naturally expands sampling as problems become more  
 ill-posed. This linear relationship demonstrates stable, predictable behavior across difficulty levels  
 without artificial diversity suppression.

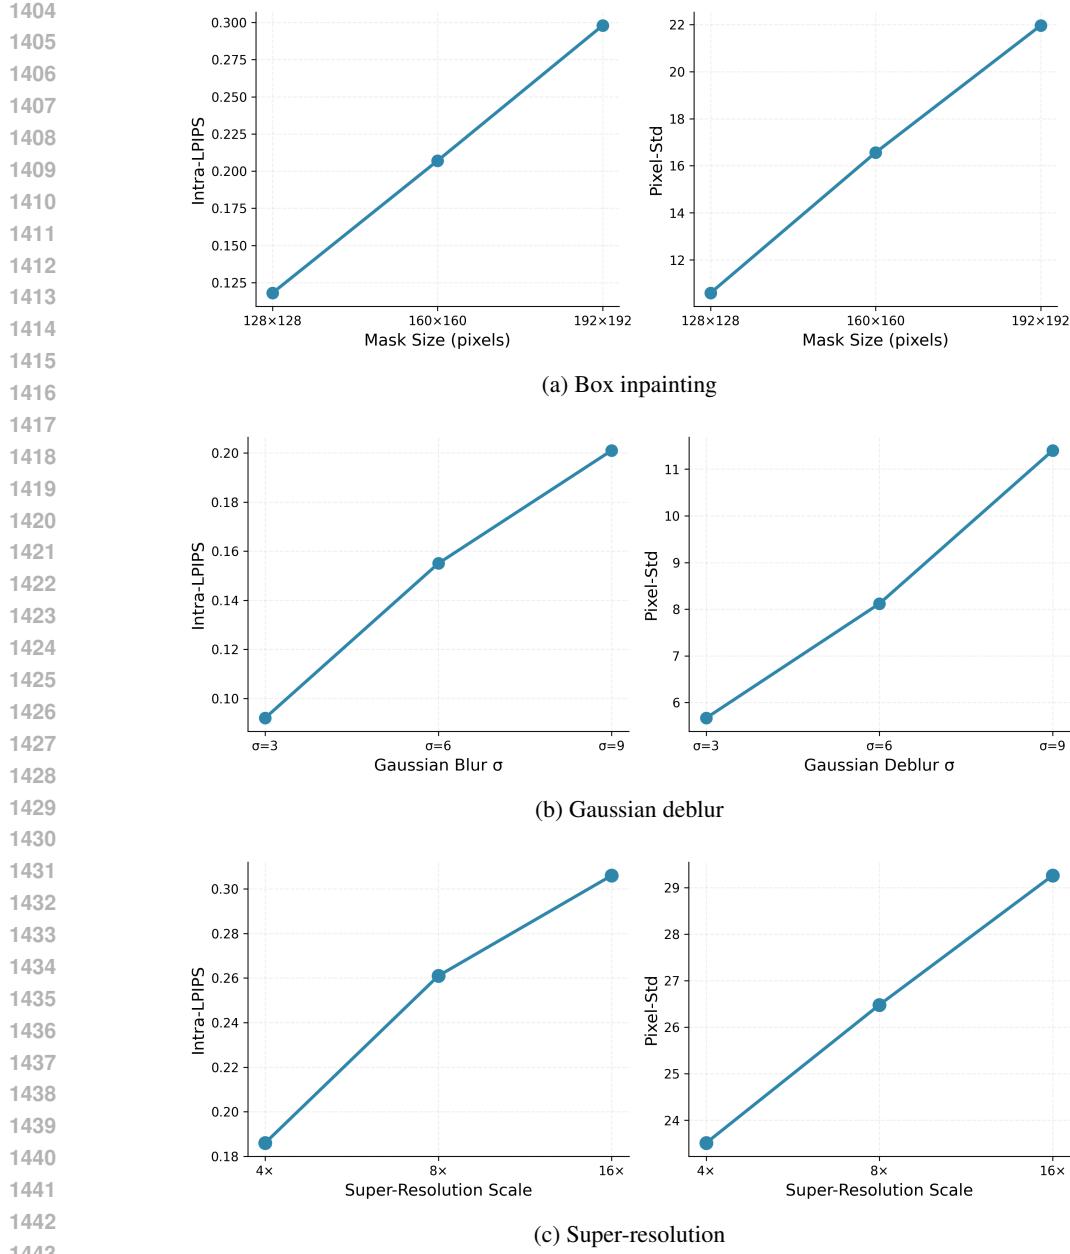


Figure 18: **Diversity vs task difficulty across three inverse problems.** As task difficulty increases (larger inpainting mask, stronger blur, higher SR scale), both diversity metrics increase proportionally, demonstrating that Equi-DPS maintains healthy posterior sampling behavior across a wide difficulty spectrum.

Figures 7 and 19 provide qualitative results.

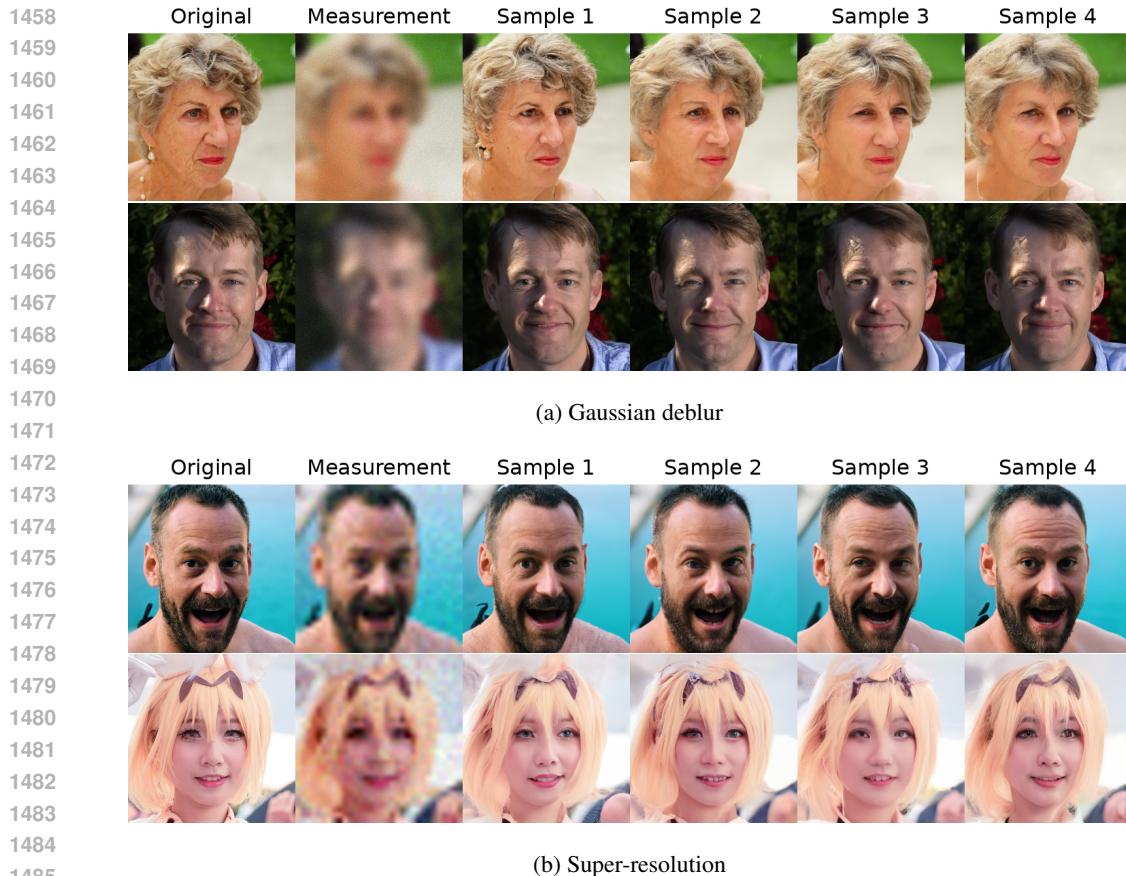


Figure 19: **Qualitative diversity examples across Gaussian deblur and super-resolution.** Each subfigure shows  $K = 4$  posterior samples for two different test images. (a) Gaussian deblur: samples differ in facial expressions and accessories (i.e., earrings in first test image). (b) Super-resolution: samples differ in facial features (i.e., teeth in the first test image, eye color and eyelashes in the second test image). Across both tasks, EquiReg produces diverse plausible reconstructions rather than collapsing to a single mode.

### D.3 CONCLUSION

Finally, we highlight that EquiReg improves both fidelity and diversity on 2 of the 3 considered tasks, an encouraging outcome that is uncommon given the general behavior of classical regularizers. Hand-crafted regularizers such as TV and  $\ell_1$  may suppress diversity by shrinking solutions toward simple structures. By contrast, EquiReg leverages data-dependent regularization that captures the richness and structural complexity of the underlying data manifold, enabling it to preserve manifold-consistent variability while suppressing implausible samples.

High diversity without fidelity is not meaningful for posterior sampling. A method that samples the entire solution space, including low-probability and artifactual regions, may score well on diversity but fail to provide useful reconstructions. Equi-DPS avoids this failure mode: it maintains meaningful diversity while reducing artifacts and improving perceptual quality. In the experiments conducted during the rebuttal, our goal was to demonstrate clearly that EquiReg preserves meaningful diversity, reflecting the posterior uncertainty, rather than unstructured or unconstrained variability.

## E IMPLEMENTATION DETAILS FOR IMAGE RESTORATION TASKS

**Experimental Setup.** We evaluate EquiReg on a variety of linear and nonlinear restoration tasks for natural images. We fix sets of 100 images from FFHQ and ImageNet as our validation sets. All images are normalized from  $[0, 1]$ . For the majority of experiments, we use noise level  $\sigma_y = 0.05$ .

(we indicate  $\sigma_y$  in our tables). For linear inverse problems, we consider (1) box inpainting, (2) random inpainting, (3) Gaussian deblur, (4) motion deblur, and (5) super-resolution. We apply a random  $128 \times 128$  pixel box for box inpainting, and a 70% random mask for random inpainting. For Gaussian and motion deblur, we use kernels of size  $61 \times 61$ , with standard deviations of 3.0 and 0.5, respectively. For super-resolution, we downscale images by a factor of 4 using a bicubic resizer. For nonlinear inverse problems, we consider (1) phase retrieval, (2) nonlinear deblur, and (3) high dynamic range (HDR). We use an oversampling rate of 2.0 for phase retrieval, and due to instability of the task, we generate four independent reconstructions and take the best result (as also done in DPS (Chung et al., 2023), DAPS (Zhang et al., 2025a), and DiffStateGrad (Zirvi et al., 2025)). We use the default setting from (Tran et al., 2021) for nonlinear deblur, and a scale factor of 2 for HDR.

**Hyperparameters.** Our method introduces a single hyperparameter  $\lambda_t$  that controls the amount of regularization applied. Below we include a table detailing the use of this hyperparameter in the main experiments (Table 12). For majority of experiments, we keep  $\lambda_t$  constant throughout iterations. For all unscaled experiments, we employ early stopping, setting  $\lambda_t = 0$  for the last 10% of sampling.

Table 12: Equivariance regularization weight  $\lambda_t$  used in main experiments.

Method	Box Inpainting	Random Inpainting	Gaussian Deblur	Motion Deblur	Super-resolution ( $\times 4$ )
<i>FFHQ</i> $256 \times 256$					
Equi-PSLD	0.05	0.05	0.03	0.03	0.02
EquiCon-PSLD	0.01	0.01	0.01	0.01	0.01
Equi-ReSample	0.03	0.05	0.02	0.02	0.05
EquiCon-ReSample	0.001	0.001	0.001	0.001	0.001
Equi-DPS	0.0001	0.001	0.001	0.001	0.1
<i>ImageNet</i> $256 \times 256$					
EquiCon-PSLD	0.0015	0.05	0.06	0.07	0.001

**PSLD.** We integrate EquiReg into PSLD by simply adding an additional gradient update step using our regularization term (Algorithms 2 and 3).

In our experiments, we use the official PSLD implementation from Rout et al. (2023), running with its default settings to reproduce the baseline results. We note that in our code, we do not square the norm when computing the gradient, aligning with PSLD’s implementation.

---

**Algorithm 2** Equi-PSLD for Image Restoration Tasks

---

```

Require:  $T, \mathbf{y}, \{\eta_t\}_{t=1}^T, \{\gamma_t\}_{t=1}^T, \{\tilde{\sigma}_t\}_{t=1}^T$ 
Require:  $\mathcal{E}, \mathcal{D}, \mathcal{A}\mathbf{x}_0^*, \mathcal{A}, \mathbf{s}_\theta, \mathbf{T}_g$  and  $\mathbf{S}_g, \{\lambda_t\}_{t=1}^T$ 
1:  $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T - 1$  to 0 do
3:    $\hat{\mathbf{s}} \leftarrow \mathbf{s}_\theta(\mathbf{z}_t, t)$ 
4:    $\mathbf{z}_{0|t} \leftarrow \frac{1}{\sqrt{\alpha_t}}(\mathbf{z}_t + (1 - \bar{\alpha}_t)\hat{\mathbf{s}})$ 
5:    $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
6:    $\mathbf{z}'_{t-1} \leftarrow \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{z}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{z}_{0|t} + \tilde{\sigma}_t \boldsymbol{\epsilon}$ 
7:    $\mathbf{z}''_{t-1} \leftarrow \mathbf{z}'_{t-1} - \eta_t \nabla_{\mathbf{z}_t} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{z}_{0|t}))\|_2^2$ 
8:    $\mathbf{z}_{t-1} \leftarrow \mathbf{z}''_{t-1} - \gamma_t \nabla_{\mathbf{z}_t} \|\mathbf{z}_{0|t} - \mathcal{E}(\mathcal{A}^T \mathcal{A}\mathbf{x}_0^* + (\mathbf{I} - \mathcal{A}^T \mathcal{A})\mathcal{D}(\mathbf{z}_{0|t}))\|_2^2$ 
9:    $\mathbf{z}_{t-1} \leftarrow \mathbf{z}_{t-1} - \lambda_t \nabla_{\mathbf{z}_t} \|\mathbf{S}_g(\mathcal{D}(\mathbf{z}_{0|t})) - \mathcal{D}(\mathbf{T}_g(\mathbf{z}_{0|t}))\|_2^2$ 
10: end for
11: return  $\mathcal{D}(\mathbf{z}_{0|t})$ 

```

---

1566

1567

**Algorithm 3** EquiCon-PSLD for Image Restoration Tasks

1568

1569

**Require:**  $T, \mathbf{y}, \{\eta_t\}_{t=1}^T, \{\gamma_t\}_{t=1}^T, \{\tilde{\sigma}_t\}_{t=1}^T$ 

1570

**Require:**  $\mathcal{E}, \mathcal{D}, \mathcal{A}\mathbf{x}_0^*, \mathcal{A}, \mathbf{s}_\theta, \mathbf{T}_g$  and  $\mathbf{S}_g, \{\lambda_t\}_{t=1}^T$ 

1571

1:  $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 

1572

2: **for**  $t = T - 1$  **to** 0 **do**

1573

3:  $\hat{\mathbf{s}} \leftarrow \mathbf{s}_\theta(\mathbf{z}_t, t)$ 

1574

4:  $\mathbf{z}_{0|t} \leftarrow \frac{1}{\sqrt{\alpha_t}}(\mathbf{z}_t + (1 - \bar{\alpha}_t)\hat{\mathbf{s}})$ 

1575

5:  $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 

1576

6:  $\mathbf{z}'_{t-1} \leftarrow \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{z}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\mathbf{z}_{0|t} + \tilde{\sigma}_t\boldsymbol{\epsilon}$ 

1577

7:  $\mathbf{z}''_{t-1} \leftarrow \mathbf{z}'_{t-1} - \eta_t \nabla_{\mathbf{z}_t} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{z}_{0|t}))\|_2^2$ 

1578

8:  $\mathbf{z}_{t-1} \leftarrow \mathbf{z}''_{t-1} - \gamma_t \nabla_{\mathbf{z}_t} \|\mathbf{z}_{0|t} - \mathcal{E}(\mathcal{A}^T \mathcal{A}\mathbf{x}_0^* + (\mathbf{I} - \mathcal{A}^T \mathcal{A})\mathcal{D}(\mathbf{z}_{0|t}))\|_2^2$ 

1579

9:  $\mathbf{z}_{t-1} \leftarrow \mathbf{z}_{t-1} - \lambda_t \nabla_{\mathbf{z}_t} \|\mathbf{z}_{0|t} - \mathcal{E}(\mathbf{S}_g^{-1}(\mathcal{D}(\mathbf{T}_g(\mathbf{z}_{0|t}))))\|_2^2$ 

1580

10: **end for**

1581

11: **return**  $\mathcal{D}(\mathbf{z}_{0|t})$ 

1582

1583

1584

1585

1586

1587

1588

**ReSample.** We integrate EquiReg into ReSample by adding our regularization term into the hard data consistency step (Algorithms 4 and 5). We note that the ReSample algorithm employs a two-stage approach; initially, it performs pixel-space optimization, and later it performs latent-space optimization. We apply EquiReg in the latent-space optimization stage.

1592

1593

In our experiments, we use the official ReSample implementation from Song et al. (2023a), running with its default settings to reproduce the baseline results.

1594

1595

1596

1597

1598

1599

1600

**Algorithm 4** Equi-ReSample for Image Restoration Tasks

1601

1602

**Require:** Measurements  $\mathbf{y}$ ,  $\mathcal{A}(\cdot)$ , Encoder  $\mathcal{E}(\cdot)$ , Decoder  $\mathcal{D}(\cdot)$ , Score function  $\mathbf{s}_\theta(\cdot, t)$ , Pretrained LDM Parameters  $\beta_t, \bar{\alpha}_t, \eta, \delta$ , Hyperparameter  $\gamma$  to control  $\sigma_t^2$ , Time steps to perform resample  $C$ ,  $\mathbf{T}_g$  and  $\mathbf{S}_g$ ,  $\{\lambda_t\}_{t=1}^T$

1603

1:  $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 

▷ Initial noise vector

1604

2: **for**  $t = T - 1, \dots, 0$  **do**

1605

3:  $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 

1606

4:  $\hat{\mathbf{e}}_{t+1} = \mathbf{s}_\theta(\mathbf{z}_{t+1}, t+1)$ 

▷ Compute the score

1607

5:  $\hat{\mathbf{z}}_0(\mathbf{z}_{t+1}) = \frac{1}{\sqrt{\bar{\alpha}_{t+1}}}(\mathbf{z}_{t+1} - \sqrt{1 - \bar{\alpha}_{t+1}}\hat{\mathbf{e}}_{t+1})$ ▷ Predict  $\hat{\mathbf{z}}_0$  using Tweedie's formula

1608

6:  $\mathbf{z}'_t = \sqrt{\bar{\alpha}_t}\hat{\mathbf{z}}_0(\mathbf{z}_{t+1}) + \sqrt{1 - \bar{\alpha}_t - \eta\delta^2}\hat{\mathbf{e}}_{t+1} + \eta\delta\boldsymbol{\epsilon}_t$ 

▷ Unconditional DDIM step

1609

7: **if**  $t \in C$  **then**

▷ ReSample time step

1610

8: Initialize  $\hat{\mathbf{z}}_0(\mathbf{y})$  with  $\hat{\mathbf{z}}_0(\mathbf{z}_{t+1})$ 

1611

9: **for** each step in gradient descent **do**

1612

10:  $\mathbf{g} \leftarrow \nabla_{\hat{\mathbf{z}}_0(\mathbf{y})} \frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\hat{\mathbf{z}}_0(\mathbf{y})))\|_2^2 + \lambda_t \nabla_{\hat{\mathbf{z}}_0(\mathbf{y})} \|\mathbf{S}_g(\mathcal{D}(\hat{\mathbf{z}}_0(\mathbf{y}))) - \mathcal{D}(\mathbf{T}_g(\hat{\mathbf{z}}_0(\mathbf{y})))\|_2^2$ 

1613

11: Update  $\hat{\mathbf{z}}_0(\mathbf{y})$  using gradient  $\mathbf{g}$ 

1614

12: **end for**

1615

13:  $\mathbf{z}_t = \text{StochasticResample}(\hat{\mathbf{z}}_0(\mathbf{y}), \mathbf{z}'_t, \gamma)$ ▷ Map back to  $t$ 

1616

14: **else**

▷ Unconditional sampling if not resampling

1617

15:  $\mathbf{z}_t = \mathbf{z}'_t$ 

1618

16: **end if**

1619

17: **end for**18:  $\mathbf{x}_0 = \mathcal{D}(\mathbf{z}_0)$ 

▷ Output reconstructed image

19: **return**  $\mathbf{x}_0$

1620

**Algorithm 5** EquiCon-ReSample for Image Restoration Tasks

---

1621 **Require:** Measurements  $\mathbf{y}$ ,  $\mathcal{A}(\cdot)$ , Encoder  $\mathcal{E}(\cdot)$ , Decoder  $\mathcal{D}(\cdot)$ , Score function  $\mathbf{s}_\theta(\cdot, t)$ , Pretrained LDM  
 1622 Parameters  $\beta_t, \bar{\alpha}_t, \eta, \delta$ , Hyperparameter  $\gamma$  to control  $\sigma_t^2$ , Time steps to perform resample  $C$ ,  $\mathbf{T}_g$  and  $\mathbf{S}_g$ ,  
 1623  $\{\lambda_t\}_{t=1}^T$

1624 1:  $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Initial noise vector

1625 2: **for**  $t = T-1, \dots, 0$  **do**

1626 3:    $\epsilon_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

1627 4:    $\hat{\mathbf{e}}_{t+1} = \mathbf{s}_\theta(\mathbf{z}_{t+1}, t+1)$  ▷ Compute the score

1628 5:    $\hat{\mathbf{z}}_0(\mathbf{z}_{t+1}) = \frac{1}{\sqrt{\bar{\alpha}_{t+1}}}(\mathbf{z}_{t+1} - \sqrt{1-\bar{\alpha}_{t+1}}\hat{\mathbf{e}}_{t+1})$  ▷ Predict  $\hat{\mathbf{z}}_0$  using Tweedie's formula

1629 6:    $\mathbf{z}'_t = \sqrt{\bar{\alpha}_t}\hat{\mathbf{z}}_0(\mathbf{z}_{t+1}) + \sqrt{1-\bar{\alpha}_t - \eta\delta^2}\hat{\mathbf{e}}_{t+1} + \eta\delta\epsilon_1$  ▷ Unconditional DDIM step

1630 7:   **if**  $t \in C$  **then** ▷ ReSample time step

1631 8:     Initialize  $\hat{\mathbf{z}}_0(\mathbf{y})$  with  $\hat{\mathbf{z}}_0(\mathbf{z}_{t+1})$

1632 9:     **for** each step in gradient descent **do** ▷ Compute the score

1633 10:        $\mathbf{g} \leftarrow \nabla_{\hat{\mathbf{z}}_0(\mathbf{y})} \frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\hat{\mathbf{z}}_0(\mathbf{y})))\|_2^2 + \lambda_t \nabla_{\hat{\mathbf{z}}_0(\mathbf{y})} \|\hat{\mathbf{z}}_0(\mathbf{y}) - \mathcal{E}(S_g^{-1}(\mathcal{D}(\mathbf{T}_g(\hat{\mathbf{z}}_0(\mathbf{y}))))\|_2^2$

1634 11:       Update  $\hat{\mathbf{z}}_0(\mathbf{y})$  using gradient  $\mathbf{g}$

1635 12:     **end for** ▷ Map back to  $t$

1636 13:      $\mathbf{z}_t = \text{StochasticResample}(\hat{\mathbf{z}}_0(\mathbf{y}), \mathbf{z}'_t, \gamma)$  ▷ Map back to  $t$

1637 14:   **else** ▷ Unconditional sampling if not resampling

1638 15:      $\mathbf{z}_t = \mathbf{z}'_t$

1639 16:   **end if**

1640 17: **end for** ▷ Output reconstructed image

17: **return**  $\mathbf{x}_0$

---

1641

1642

1643

1644

**DPS.** Similar to PSLD, we integrate EquiReg into DPS by simply adding an additional gradient update step using our regularization term (Algorithm 6).

1645

1646

1647

1648

In our experiments, we use the official DPS implementation from Chung et al. (2023), running with its default settings to reproduce the baseline results.

1649

1650

1651

1652

**Algorithm 6** Equi-DPS for Image Restoration Tasks

---

1653 **Require:**  $T, \mathbf{y}, \{\zeta_t\}_{t=1}^T, \{\tilde{\sigma}_t\}_{t=1}^T, \mathbf{s}_\theta, \mathcal{E}, \mathbf{T}_g$  and  $\mathbf{S}_g, \{\lambda_t\}_{t=1}^T$

1654 1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

1655 2: **for**  $t = T-1$  **to** 0 **do**

1656 3:    $\hat{\mathbf{s}} \leftarrow \mathbf{s}_\theta(\mathbf{x}_t, t)$

1657 4:    $\mathbf{x}_{0|t} \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t + (1-\bar{\alpha}_t)\hat{\mathbf{s}})$

1658 5:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

1659 6:    $\mathbf{x}'_{t-1} \leftarrow \frac{\sqrt{\bar{\alpha}_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}\beta_t}}{1-\bar{\alpha}_t}\mathbf{x}_{0|t} + \tilde{\sigma}_t\epsilon$

1660 7:    $\mathbf{x}_{t-1} \leftarrow \mathbf{x}'_{t-1} - \zeta_t \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\mathbf{x}_{0|t})\|_2^2$

1661 8:    $\mathbf{x}_{t-1} \leftarrow \mathbf{x}_{t-1} - \lambda_t \nabla_{\mathbf{x}_t} \|\mathbf{S}_g(\mathcal{E}(\mathbf{x}_{0|t})) - \mathcal{E}(\mathbf{T}_g(\mathbf{x}_{0|t}))\|_2^2$

1662 9: **end for**

1663 10: **return**  $\mathbf{x}_0$

---

1664

1665

1666

1667

1668

**SITCOM.** We augment the original SITCOM algorithm by introducing an additional equivariant refinement stage at each reverse diffusion step. After completing the standard measurement and backward-consistency gradient updates, we perform a second optimization over the equivariance loss, enforcing consistency between  $\mathcal{E}(\mathbf{T}_g(\mathbf{v}))$  and  $\mathbf{T}_g(\mathcal{E}(\mathbf{v}))$  (Algorithm 7).

1669

1670

1671

1672

1673

In our experiments, we use the official SITCOM implementation from Alkhouri et al. (2025), running with its default settings to reproduce the baseline results.

---

1674 **Algorithm 7** Equi-SITCOM for Image Restoration Tasks

1675

1676 **Require:** Measurements  $\mathbf{y}$ , forward operator  $\mathcal{A}(\cdot)$ , pre-trained DM  $\epsilon_\theta(\cdot, \cdot)$ , diffusion steps  $N$ , schedule  $\bar{\alpha}_i$ ,  
1677 measurement gradient steps  $K$ , equivariant gradient steps  $K_{\text{equi}}$ , stop  $\delta$ , lr  $\gamma$ , reg.  $\lambda$ .

1678 **Ensure:** Restored image  $\hat{\mathbf{x}}$ .

1679 1: **Initialize**  $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $\Delta t = \lfloor \frac{T}{N} \rfloor$ . ▷ Reducing diffusion sampling steps

1680 2: **for**  $i = N, N-1, \dots, 1$  **do** ▷ Init for closeness (C3)

1681 3:    $\mathbf{v}_i^{(0)} \leftarrow \mathbf{x}_i$  ▷ Adam on measurement/backward consistency (C1, C2)

1682 4:   **for**  $k = 1, \dots, K$  **do** ▷ Adam on measurement/backward consistency (C1, C2)

1683 5:      $\mathbf{v}_i^{(k)} \leftarrow \mathbf{v}_i^{(k-1)} - \gamma \nabla_{\mathbf{v}_i} \left[ \left\| \mathcal{A} \left( \frac{1}{\sqrt{\bar{\alpha}_i}} (\mathbf{v}_i - \sqrt{1 - \bar{\alpha}_i} \epsilon_\theta(\mathbf{v}_i, i\Delta t)) \right) - \mathbf{y} \right\|_2^2 + \lambda \|\mathbf{x}_i - \mathbf{v}_i\|_2^2 \right] \Big|_{\mathbf{v}_i = \mathbf{v}_i^{(k-1)}}$

1684 6:     **if**  $\left\| \mathcal{A} \left( \frac{1}{\sqrt{\bar{\alpha}_i}} (\mathbf{v}_i^{(k)} - \sqrt{1 - \bar{\alpha}_i} \epsilon_\theta(\mathbf{v}_i^{(k)}, i\Delta t)) \right) - \mathbf{y} \right\|_2^2 < \delta^2$  **then** ▷ Prevent noise overfitting

1685 7:       **break**

1686 8:     **end if**

1687 9:   **end for**

1688 10:    $\mathbf{v}_i^{(0)} \leftarrow \mathbf{v}_i^{(k)}$  ▷ Initialize to optimized  $\mathbf{v}_i$

1689 11:   **for**  $k = 1, \dots, K_{\text{equi}}$  **do**

1690 12:      $\mathbf{v}_i^{(k)} \leftarrow \mathbf{v}_i^{(k-1)} - \gamma \nabla_{\mathbf{v}_i} \left[ \left\| \mathcal{E}(T_g(\mathbf{v}_i^{(k)})) - T_g(\mathcal{E}(\mathbf{v}_i^{(k)})) \right\|_2^2 \right] \Big|_{\mathbf{v}_i = \mathbf{v}_i^{(k-1)}}$

1691 13:     **if**  $\left\| \mathcal{E}(T_g(\mathbf{v}_i^{(k)})) - T_g(\mathcal{E}(\mathbf{v}_i^{(k)})) \right\|_2^2 < \delta^2$  **then**

1692 14:       **break**

1693 15:     **end if**

1694 16:   **end for**

1695 17:    $\hat{\mathbf{v}}_i \leftarrow \mathbf{v}_i^{(k)}$  ▷ Backward diffusion consistency (C2)

1696 18:    $\hat{\mathbf{x}}_0' \leftarrow \frac{1}{\sqrt{\bar{\alpha}_i}} [\hat{\mathbf{v}}_i - \sqrt{1 - \bar{\alpha}_i} \epsilon_\theta(\hat{\mathbf{v}}_i, i\Delta t)]$  ▷ Backward consistency (C2)

1697 19:    $\mathbf{x}_{i-1} \leftarrow \sqrt{\bar{\alpha}_{i-1}} \hat{\mathbf{x}}_0' + \sqrt{1 - \bar{\alpha}_{i-1}} \boldsymbol{\eta}_i$ ,  $\boldsymbol{\eta}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Forward consistency (C3)

1698 20: **end for**

1699 21: **return**  $\hat{\mathbf{x}} = \mathbf{x}_0$

---

1700

1701

1702 **F EXPERIMENT SETUP FOR PDE RECONSTRUCTIONS**

1703

1704 **Helmholtz equation.** The Helmholtz equation represents wave propagation in heterogeneous  
1705 media:

1706

$$\nabla^2 u(x) + k^2 u(x) = a(x), \quad x \in (0, 1)^2, \quad (6)$$

1707 with  $k = 1$  and  $u|_{\partial\Omega} = 0$ . Coefficient fields  $a(x)$  are generated according to  $a \sim \mathcal{N}(0, (-\Delta + 9\mathbf{I})^2)$ .  
1708 We note that this system has reflection equivariance along  $x_1 = \frac{1}{2}, x_2 = \frac{1}{2}, x_1 = x_2$  and rotation  
1709 equivariance by  $\frac{\pi}{2}, \pi, \frac{3\pi}{2}$ .

1710 **Navier-Stokes equations.** Following the methodology of (Li et al., 2020), we model the time  
1711 evolution of a vorticity field,  $u(x, t)$ , governed by:

1712

$$\partial_t u(x, t) + \mathbf{w}(x, t) \cdot \nabla u(x, t) = \nu \Delta u(x, t) + f(x), \quad x \in (0, 1)^2, t \in (0, T], \quad (7)$$

1713

$$\nabla \cdot \mathbf{w}(x, t) = 0, \quad x \in (0, 1)^2, t \in [0, T], \quad (8)$$

1714

$$u(x, 0) = a(x), \quad x \in (0, 1)^2, \quad (9)$$

1715 where  $\mathbf{w}$  is the velocity field;  $\nu = \frac{1}{1000}$ , viscosity; and  $f$ , a fixed forcing term. The initial condition  
1716  $a(x)$  is drawn from  $\mathcal{N}(0, 7^{3/2}(-\Delta + 49\mathbf{I})^{-5/2})$  under periodic boundary conditions. The forcing  
1717 term is  $f(x) = 0.1 (\sin(2\pi(x_1 + x_2)) + \cos(2\pi(x_1 + x_2)))$ . We borrow the dataset from (Huang  
1718 et al., 2024). We note that this system has a reflection symmetry along the  $x_1 = x_2$  axis.

1719 **Implementation details.** EquiReg, as a regularizer for diffusion posterior sampling, can be adapted  
1720 to many inverse solvers in a plug-and-play manner. For PDE experiments, we use the same model  
1721 weights and configurations as FunDPS (Yao et al., 2025). Error rates are calculated using the  $L^2$   
1722 relative error between the predicted and true solutions, averaged on 100 randomly selected test  
1723 samples. We provide the information on the EquiReg scaling weights in Table 13.

Table 13: **EquiReg loss used in PDE experiments.**

		Helmholtz		Navier-Stokes		
		Forward	Inverse	Forward	Inverse	
		EquiReg Norm Type	MSE	L2	MSE	L2
1733	EquiReg Weight $\lambda$		100	100	100	1000

## G THEORETICAL ANALYSIS

### G.1 SUMMARY OF THE ANALYSIS

The theoretical framework presented in this paper is intended to motivate and guide the design of effective regularizers. This perspective, grounded in optimal transport theory (Ferreira and Valencia-Guevara, 2018), serves as an intuitive interpretation of the dynamics and motivates the design of regularization strategies such as EquiReg. We note that whether diffusion models follow exact Wasserstein dynamics still remains an open problem (Zheng et al., 2025).

**Proposition G.1.** *Let  $\rho(\mathbf{x}, t)$  be the distribution of  $\mathbf{x}_{T-t}$  driven by the ideal reverse dynamics (eq. (3)). Then, the evolution of  $\rho$  follows the Wasserstein-2 gradient flow associated with minimizing functional  $\Phi(\rho, t)$  defined as  $\beta_{T-t} \int [\rho \phi(\mathbf{x}, t) + \frac{1}{2} \rho \log \rho] d\mathbf{x}$ , where  $\phi(\mathbf{x}, t) = -(\log p_{T-t}(\mathbf{x}|\mathbf{y}) + \frac{1}{4} \|\mathbf{x}\|^2)$ .*

The dynamics of  $\rho$  remain the same if we replace  $\phi(\mathbf{x}, t)$  with  $\phi_C(\mathbf{x}, t) := \phi(\mathbf{x}, t) - C(t)$  for arbitrary temporal function  $C(t)$ . Without loss of generality, we assume  $\phi_C(\mathbf{x}, t) < 0$  for all  $\mathbf{x}$  and  $t$ . In practice, the density function  $p_{T-t}$  is not available and thus  $\phi_C(\mathbf{x}, t)$  is approximated as  $\hat{\phi}$  with  $p_{T-t}(\mathbf{x}_{T-t}|\mathbf{y}) \approx \tilde{C} p_{T-t}(\mathbf{x}_{T-t}) p(\mathbf{y}|\mathbb{E}[\mathbf{x}_0|\mathbf{x}_{T-t}])$  where  $\tilde{C}$  only depends on  $\mathbf{y}$ .

Because the conditional expectation  $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_{T-t}]$  is a linear combination of all candidate  $\mathbf{x}_0$ , the approximation remains relatively accurate when  $T-t$  is small (i.e.,  $\mathbf{x}_{T-t}$  stays close to the data manifold under low noise) but may incur high error for larger  $T-t$ , as shown in Figure 2b. To mitigate this, we reweight the contributions to the first term of  $\Phi$ , down-weighting unreliable estimates, and amplifying the reliable ones. The resulting reweighted functional is

$$\tilde{\Phi}(\rho, t) = \beta_{T-t} \left[ Z_t^{-1} \int \rho(\mathbf{x}) \hat{\phi}_c(\mathbf{x}, t) e^{\frac{\mathcal{R}(\mathbf{x})}{\hat{\phi}_c(\mathbf{x}, t)}} d\mathbf{x} + \frac{1}{2} \int \rho(\mathbf{x}) \log \rho(\mathbf{x}) d\mathbf{x} \right], \quad (10)$$

where  $Z_t = \int e^{\frac{\mathcal{R}(\mathbf{x})}{\hat{\phi}_c(\mathbf{x}, t)}} d\mathbf{x}$  is the normalizing factor, and  $\mathcal{R}(\mathbf{x})$  is a positive regularization that is nearly zero near the data manifold and much larger elsewhere. Intuitively, since  $\hat{\phi}_C < 0$ , the weight is nearly one for  $\mathbf{x}$  near the data manifold and much smaller elsewhere.

**Proposition G.2.** *(Informal) The evolution of  $\rho$ , the probability distribution of  $\mathbf{x}_{T-t}$  driven by the practical and regularized reverse dynamics (eq. (11)), is an approximation of the Wasserstein-2 gradient flow associated with minimizing  $\tilde{\Phi}$ .*

$$d\mathbf{x} = \left[ -\frac{\beta_t}{2} \mathbf{x} dt - \beta_t \nabla_{\mathbf{x}_t} (\log p_t(\mathbf{x}_t) + \log \int p(\mathbf{y}|\mathbf{x}_0) \tilde{p}_t(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0 - \mathcal{R}(\mathbf{x}_t)) \right] dt + \sqrt{\beta_t} d\bar{\mathbf{w}} \quad (11)$$

### G.2 PRELIMINARY AND NOTATIONS

We first remind the readers of gradient flow under the Wasserstein-2 metric and introduce the notations related to the diffusion model.

**Wasserstein Gradient Flow** Let  $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{+\infty\}$  be a functional of probability distributions. The Wasserstein gradient flow of  $\mathcal{F}$  is characterized by the minimizing movement scheme (also known as JKO scheme) introduced by (Jordan et al., 1998). For a fixed time step  $\tau > 0$ , the sequence  $(\rho_k)_{k \in \mathbb{N}}$  of probability densities is defined recursively by:

$$\rho_{k+1} \in \arg \min_{\rho \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ \frac{1}{2\tau} W_2^2(\rho, \rho_k) + \mathcal{F}(\rho) \right\},$$

where  $W_2$  denotes the 2-Wasserstein distance, and each  $\rho_k$  is a probability density representing the distribution at time  $t = k\tau$ . In the limit  $\tau \rightarrow 0$ , this discrete-time scheme recovers the continuous-time gradient flow of  $\mathcal{F}$  under the  $W_2$  metric.

1782 **Diffusion Model** A diffusion model defines a forward stochastic process  $(\mathbf{x}_t)_{t \in [0, T]}$  governed by  
 1783 the Itô SDE:

$$d\mathbf{x}_t = f(\mathbf{x}_t, t) dt + \sqrt{\beta_t} d\mathbf{w}_t, \quad (12)$$

1784 where  $\mathbf{w}_t$  is standard Brownian motion,  $\beta_t > 0$  is a time-dependent variance schedule, and  $f(\mathbf{x}, t)$   
 1785 is a drift term. For instance,  $f \equiv 0$  for a variance-exploding SDE and  $f(\mathbf{x}, t) = -\frac{\beta_t}{2}\mathbf{x}$  for a  
 1786 variance-preserving SDE defined in (Song et al., 2021). In this work, we carry out our analysis under  
 1787 a more general setting.

1788 **Assumption G.1.** *The drift term is a gradient field,  $f(\mathbf{x}, t) = \nabla h(\mathbf{x}, t)$  for a scalar function  $h$ .*

1791 This process progressively transforms an initial data distribution  $\mathbf{x}_0 \sim p_0$  into a tractable reference  
 1792 distribution (e.g., approximately a Gaussian  $\mathcal{N}(0, I)$ ) at time  $T$ .

1793 Sampling is performed by simulating the *reverse-time SDE*:

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - \beta_t \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)] dt + \sqrt{\beta_t} d\bar{\mathbf{w}}_t, \quad (13)$$

1794 where  $p_t$  is the marginal density of  $\mathbf{x}_t$ , and  $\bar{\mathbf{w}}_t$  is a standard Brownian motion in reverse time.

1795 In practice, the score function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$  is approximated by a neural network  $s_{\theta}(\mathbf{x}, t)$  trained to  
 1796 estimate the score of the forward process. For *conditional sampling*, where we sample  $\mathbf{x}_0$  given some  
 1797 observed variable  $y$ , the score is replaced by  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}|y)$  and decomposed as

$$\nabla_{\mathbf{x}} \log p_t(\mathbf{x}|y) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) + \nabla_{\mathbf{x}} \log p_t(y|\mathbf{x}), \quad (14)$$

1802 based on Bayes' rule.

1803 To simplify notation in the sequel, we perform a time reparameterization  $t = T - t'$ , so that the  
 1804 reverse process is written as a forward SDE over  $t \in [0, T]$ :

$$d\mathbf{x}_t = -[f(\mathbf{x}_t, T - t) - \beta_{T-t} [\nabla_{\mathbf{x}} \log p_{T-t}(\mathbf{x}_t) + \nabla_{\mathbf{x}} \log p_t(y|\mathbf{x}_t)]] dt + \sqrt{\beta_{T-t}} d\mathbf{w}_t, \quad (15)$$

1805 This form describes the generative process as evolving forward from  $t = 0$  to  $t = T$ , matching the  
 1806 usual direction of analysis in gradient flow frameworks.

### 1809 G.3 PROOF OF PROPOSITION G.1

1811 In this work, we consider Wasserstein gradient flow under the setting where the functional  $\mathcal{F}$  depends  
 1812 on time.

1814 **Lemma G.1.** *Consider a time-dependent functional  $\mathcal{F}(\rho, t) = \int \rho(\mathbf{x})V(\mathbf{x}, t)d\mathbf{x} + \int \alpha(t)\rho \log \rho d\mathbf{x}$ .  
 1815 Then the particle description of Wasserstein-2 gradient flow associated with this functional derived  
 1816 by JKO scheme is*

$$d\mathbf{x}_t = -\nabla V(\mathbf{x}_t, t)dt + \sqrt{2\alpha(t)}d\mathbf{w}_t. \quad (16)$$

1818 *Proof.* Consider the following optimization

$$\min_{\rho'} \mathcal{F}(\rho', t + \Delta t) - \mathcal{F}(\rho, t) + \frac{1}{2\Delta t} W_2^2(\rho, \rho'), \quad (17)$$

1822 where the change of density is restricted to the Liouville equation

$$\partial_t \rho = -\nabla \cdot (\rho v(\mathbf{x}, t)), \text{ and } \rho'(x) = \rho(x) - \Delta t \nabla \cdot (\rho(\mathbf{x})v(\mathbf{x})) + o(\Delta t). \quad (18)$$

1825 Using the static formulation of  $W_2$  distance, we have

$$W_2^2(\rho, \rho') = \int \rho(\mathbf{x}) \|\mathbf{x} - T^*(\mathbf{x})\|^2 d\mathbf{x} = \Delta t^2 \int \rho(\mathbf{x}) \|v^*(\mathbf{x})\|^2 d\mathbf{x}, \quad (19)$$

1827 where  $T^*(\mathbf{x})$  is the optimal transport map, and  $v^*(\mathbf{x})$  is the associated optimal velocity field.

1829 Thus, we can rewrite the eq. (17) as

$$\inf_v \mathcal{F}(\rho, t) - \Delta t \int \nabla \cdot (\rho(\mathbf{x})v(\mathbf{x})) \frac{\delta \mathcal{F}(\rho, t)}{\delta \rho}(\mathbf{x}) d\mathbf{x} + \Delta t \int [\rho(\mathbf{x})\partial_t V(\mathbf{x}, t) + \dot{\alpha}(t)\rho \log \rho] d\mathbf{x} - \mathcal{F}(\rho, t) + \frac{\Delta t}{2} \int \rho(\mathbf{x}) \|v(\mathbf{x})\|^2 d\mathbf{x}, \quad (20)$$

$$- \mathcal{F}(\rho, t) + \frac{\Delta t}{2} \int \rho(\mathbf{x}) \|v(\mathbf{x})\|^2 d\mathbf{x}, \quad (21)$$

1836 which simplifies to  
 1837

$$1838 \min_v \int \rho(\mathbf{x}) \left\langle v(\mathbf{x}), \nabla \frac{\delta \mathcal{F}(\rho, t)}{\delta \rho}(\mathbf{x}) \right\rangle d\mathbf{x} + \frac{1}{2} \int \rho(\mathbf{x}) \|v(\mathbf{x})\|^2 d\mathbf{x}, \quad (22)$$

1840 since the last term in the first line of (20) does not depend on  $v$ . and further to  
 1841

$$1842 \min_v \int \rho(\mathbf{x}) \left\| v(\mathbf{x}) + \nabla \frac{\delta \mathcal{F}(\rho, t)}{\delta \rho}(\mathbf{x}) \right\|^2 d\mathbf{x}. \quad (23)$$

1844 From the optimality condition of the above problem, we obtain  
 1845

$$1846 \quad 1847 v(\mathbf{x}, t) = -\nabla \frac{\delta \mathcal{F}(\rho, t)}{\delta \rho}(\mathbf{x}) = -(\nabla V(\mathbf{x}, t) + \alpha(t) \nabla \log \rho(\mathbf{x}, t)). \quad (24)$$

1848 We note that By Hörmander's theorem, a smooth density  $\rho(\mathbf{x}, t)$  exists for  $t > 0$ , ensuring that the  
 1849 above  $v$  is well-defined. The corresponding evolution of probability density is  
 1850

$$1851 \quad \partial_t \rho(\mathbf{x}, t) = -\nabla \cdot (\rho(\mathbf{x}, t) v(\mathbf{x}, t)) \quad (25)$$

$$1852 \quad 1853 = \nabla \cdot (\rho(\mathbf{x}, t) (\nabla V(\mathbf{x}, t) + \alpha(t) \frac{\nabla \rho(\mathbf{x}, t)}{\rho})) \quad (26)$$

$$1854 \quad 1855 = -\nabla \cdot (\rho(\mathbf{x}, t) (-\nabla V(\mathbf{x}, t)) + \alpha(t) \Delta \rho(\mathbf{x}, t)), \quad (27)$$

1856 which is exactly the Fokker-Planck equation describing the evolution of the probability density  
 1857 describing the particles following

$$1858 \quad d\mathbf{x}_t = -\nabla V(\mathbf{x}_t, t) dt + \sqrt{2\alpha(t)} d\mathbf{w}_t. \quad (28)$$

1860  $\square$

1862 Now we come back to Proposition G.1. From eq. (15) we know that choosing  
 1863

$$1864 \quad V(\mathbf{x}, t) = h(\mathbf{x}, T-t) - \beta_{T-t} [\log p_{T-t}(\mathbf{x}) + \log p_{T-t}(\mathbf{y}|\mathbf{x})] \text{ and } \alpha(t) = \frac{\beta_{T-t}}{2} \quad (29)$$

1865 in Lemma G.1 completes the proof, where  $h$  is defined in Assumption G.1.

#### 1868 G.4 DETAILED VERSION OF PROPOSITION G.2

1869 In practice, one does not have access to  $\log p_t(\mathbf{y}|\mathbf{x}_t)$  which appears in the reverse SDE. The most  
 1870 popular approach is do the following approximation,  
 1871

$$1872 \quad p_t(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0) p(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0 = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)} [p(\mathbf{y}|\mathbf{x}_0)] \approx p(\mathbf{y}) \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t], \quad (30)$$

1874 which can be interpreted as exchanging two operations, the conditional expectation and the measure-  
 1875 ment  $p(\mathbf{y}|\cdot)$ .

1876 As discussed in the main text, since the conditional expectation is a linear combination over all  
 1877 possible values of  $\mathbf{x}_0$ , it may fall outside the data manifold, resulting in physically invalid samples.  
 1878 One of the central challenges in diffusion-based inverse sampling is guiding the sampling trajectory,  
 1879 generated by the reverse SDE dynamics, toward the data manifold. A common strategy is to  
 1880 incorporate regularization into the reverse SDE to encourage manifold adherence. In this work,  
 1881 building on the perspective of Wasserstein gradient flow as outlined above, we provide a novel  
 1882 interpretation of the role played by such regularization terms.

1883 We show that the regularizer serves to reweight the contribution of different regions in the calculation  
 1884 of the underlying functional being minimized,  $\Phi(\rho, t)$  defined in Proposition G.1. Specifically, it  
 1885 amplifies the influence of regions where the density estimate is reliable (typically near the data  
 1886 manifold), while down-weighting regions with poor approximation quality of based on eq. (30), often  
 1887 corresponding to off-manifold samples.

1888 Following from what we have shown in the main text,  $\Phi(\rho, t)$  has the form of  $\beta_{T-t} \int [\rho \phi(\mathbf{x}, t) +$   
 1889  $\frac{1}{2} \rho \log \rho] d\mathbf{x}$  for a function  $\phi(\mathbf{x}, t)$ , which can be derived by (29). The  $\log p_t(\mathbf{y}|\mathbf{x})$  term in (29) or

1890  $\nabla \log p_t(\mathbf{y}|\mathbf{x})$  term in (28), equivalently, is computed based on approximation (30). We denote the  
 1891 corresponding approximation of  $\phi(\mathbf{x}, t)$  as  $\hat{\phi}(\mathbf{x}, t)$ . As discussed in the main text, we can assume  
 1892 without loss of generality that  $\phi(\mathbf{x}, t) < 0$  and  $\hat{\phi}(\mathbf{x}, t) < 0$ . We have  
 1893

$$1894 \hat{\Phi}(\rho, t) = \beta_{T-t} \left[ \int_{\mathbf{x} \in N(\mathcal{M})} \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) d\mathbf{x} + \int_{\mathbf{x} \notin N(\mathcal{M})} \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) d\mathbf{x} + \frac{1}{2} \int \rho \log \rho d\mathbf{x} \right], \quad (31)$$

1895 where  $N(\mathcal{M})$  denotes a neighborhood of the data manifold  $\mathcal{M}$ . Intuitively, we aim to focus on the  
 1896 contribution from regions near  $\mathcal{M}$ , which corresponds to the first term, while down-weighting the  
 1897 influence of points farther away, where the approximation tends to be unreliable. For instance, we  
 1898 can introduce two positive weights  $A \gg B$  and adopt the modified functional  
 1899

$$1900 \tilde{\Phi}(\rho, t) = \beta_{T-t} \left[ A \int_{\mathbf{x} \in N(\mathcal{M})} \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) d\mathbf{x} + B \int_{\mathbf{x} \notin N(\mathcal{M})} \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) d\mathbf{x} + \frac{1}{2} \int \rho \log \rho d\mathbf{x} \right]. \quad (32)$$

1901 In this work, we further generalize this idea and consider a continuous weight function,  
 1902

$$1903 \tilde{\Phi}(\rho, t) = \beta_{T-t} \left[ \int \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) \lambda(\mathbf{x}) d\mathbf{x} + \frac{1}{2} \int \rho \log \rho d\mathbf{x} \right], \quad (33)$$

1904 where the non-negative weight  $\lambda(\mathbf{x})$  is large for  $\mathbf{x} \in N(\mathcal{M})$  and small elsewhere.  
 1905

1906 In practice, a nonnegative regularization function  $\mathcal{R}(\mathbf{x})$  is introduced, ideally being nearly zero for  $\mathbf{x}$   
 1907 near the data manifold and much larger elsewhere. We consider the following modified functional  
 1908 with weight function  $\lambda(\mathbf{x}, t) := e^{\frac{\mathcal{R}(\mathbf{x})}{\hat{\phi}(\mathbf{x}, t)}}$ ,  
 1909

$$1910 \tilde{\Phi}(\rho, t) = \beta_{T-t} \left[ \int \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) e^{\frac{\mathcal{R}(\mathbf{x})}{\hat{\phi}(\mathbf{x}, t)}} d\mathbf{x} + \frac{1}{2} \int \rho(\mathbf{x}) \log \rho(\mathbf{x}) d\mathbf{x} \right]. \quad (34)$$

1911 Note that  $\hat{\phi} < 0$ , we have that  
 1912

$$1913 \mathcal{R}(\mathbf{x}) \approx \begin{cases} 0, & \mathbf{x} \in N(\mathcal{M}) \\ \gg 1, & \mathbf{x} \text{ far away from } N(\mathcal{M}) \end{cases} \Rightarrow \lambda(\mathbf{x}, t) \approx \begin{cases} 1, & \mathbf{x} \in N(\mathcal{M}) \\ 0, & \mathbf{x} \text{ far away from } N(\mathcal{M}) \end{cases}.$$

1914 Next, we consider practical algorithms based on this reweighted functional. In practice, we only have  
 1915 the score function instead of the function value of  $\log p_{T-t}(\mathbf{x})$ . Thus, the Wasserstein gradient flow  
 1916 associated with (34) is intractable since we cannot evaluate the weight function. We consider the  
 1917 following approximation based on  $e^\delta \approx 1 + \delta$  when  $\delta$  is sufficiently small,  
 1918

$$1919 \tilde{\Phi}(\rho, t) \approx \beta_{T-t} \left[ \int \rho(\mathbf{x}) \hat{\phi}(\mathbf{x}, t) \left( 1 + \frac{\mathcal{R}(\mathbf{x})}{\hat{\phi}(\mathbf{x}, t)} \right) d\mathbf{x} + \frac{1}{2} \int \rho(\mathbf{x}) \log \rho(\mathbf{x}) d\mathbf{x} \right] \quad (35)$$

$$1920 = \beta_{T-t} \left[ \int \rho(\mathbf{x}) (\hat{\phi}(\mathbf{x}, t) + \mathcal{R}(\mathbf{x})) d\mathbf{x} + \frac{1}{2} \int \rho(\mathbf{x}) \log \rho(\mathbf{x}) d\mathbf{x} \right]. \quad (36)$$

1921 By Lemma G.1, the dynamics of  $\mathbf{x}$  driven by the Wasserstein gradient flow associated with the  
 1922 approximated functional above is  
 1923

$$1924 d\mathbf{x} = [-f(\mathbf{x}, T-t) - \beta_{T-t} \nabla_{\mathbf{x}} (\log p_{T-t}(\mathbf{x}) + \log \hat{p}_{T-t}(\mathbf{y}|\mathbf{x}) + \mathcal{R}(\mathbf{x}))] dt + \sqrt{\beta_{T-t}} d\bar{w}. \quad (37)$$

1925 This completes the proof.  
 1926

1927 **Remark 1.** Since  $\hat{\phi} < 0$ , and  $e^A \geq 1 + A$  for any  $A \in \mathbb{R}$ , the dynamics derived by the approximated  
 1928 functional in (36) is evolving to minimize an upper bound of the reweighted functional  $\tilde{\Phi}$ .  
 1929

## 1930 H ADDITIONAL BACKGROUND INFORMATION

1931 **Solving inverse problems with deep learning prior to diffusion models.** Earlier works (Metzler  
 1932 et al., 2016; Romano et al., 2017; Zhang et al., 2017; Metzler et al., 2017) used deep neural networks  
 1933 as denoisers to solve inverse problems. Furthermore, deep generative models such as variational  
 1934 autoencoders (VAEs) (Kingma, 2013), and generative adversarial networks (GANs) (Goodfellow  
 1935 et al., 2014) have been used to solve inverse problems.

1944 et al., 2014) were employed. Notable applications include compressed sensing (Bora et al., 2017) and  
 1945 MRI (Jalal et al., 2021).

1946 **Applications on diffusion models to solve inverse problems.** Most popular applications include  
 1947 image restoration (Chung et al., 2023; 2022b; Kawar et al., 2022; Lugmayr et al., 2022; Saharia  
 1948 et al., 2022; Song et al., 2023a; Rout et al., 2023; Zhu et al., 2023; ?; Zirvi et al., 2025), medical  
 1949 imaging (Song et al., 2022; Chung and Ye, 2022; Chung et al., 2022a; Hung et al., 2023; Dorjsembe  
 1950 et al., 2024; Li et al., 2024; Kazerouni et al., 2023; Bian et al., 2024), and solving partial differential  
 1951 equations (PDEs) (Isakov, 2006; Huang et al., 2024; Shysheya et al., 2024; Liu et al., 2023; Li et al.,  
 1952 2025; Baldassari et al., 2023; Mammadov et al., 2024a; Yao et al., 2025). On the methodology side,  
 1953 there has been numerous advancements (Chung et al., 2023; 2022b; Kawar et al., 2022; Lugmayr  
 1954 et al., 2022; Saharia et al., 2022; Song et al., 2023a; Rout et al., 2023; Zhu et al., 2023; ?; Zirvi et al.,  
 1955 2025; Song et al., 2022; Chung and Ye, 2022; Chung et al., 2022a; Hung et al., 2023; Dorjsembe  
 1956 et al., 2024; Li et al., 2024; Kazerouni et al., 2023; Bian et al., 2024; Huang et al., 2024; Shysheya  
 1957 et al., 2024; Mammadov et al., 2024b; Cardoso et al., 2024).

1958 **Resources for Definition 3.2 on vanishing-error autoencoders.** Manifold constrained distribution-  
 1959 dependent equivariance error uses the notion of *vanishing-error autoencoders* (Shao et al., 2018;  
 1960 Anders et al., 2020; He et al., 2024) (Definition H.1), also known as an asymptotically-trained  
 1961 autoencoder (Anders et al., 2020) or a perfect autoencoder (He et al., 2024). Vanishing-error  
 1962 autoencoders have previously been employed by diffusion-based inverse solvers to preserve the  
 1963 diffusion process on the manifold (He et al., 2024).

1964 **Definition H.1 (Vanishing-Error Autoencoder).** A vanishing-error autoencoder under the manifold  
 1965  $\mathcal{M}$  with encoder  $\mathcal{E} : \mathcal{X} \rightarrow \mathcal{Z}$  and decoder  $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{X}$  with  $\mathcal{Z} = \mathbb{R}^k$  where  $k < d$ , has zero  
 1966 reconstruction error under the support of the data distribution  $\mathcal{X}$ , i.e.,  $\forall \mathbf{x} \in \mathcal{X} \subset \mathcal{M}$ ,  $\mathbf{x} = \mathcal{D}(\mathcal{E}(\mathbf{x}))$ .  
 1967 It follows that the decoder is surjective on the data manifold,  $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{M}$  (He et al., 2024), and the  
 1968 encoder-decoder composition forms an identity map, i.e.,  $\forall \mathbf{z} \in \mathcal{M}$ ,  $\mathbf{z} = \mathcal{E}(\mathcal{D}(\mathbf{z}))$ .

1969 **Equivariance.** Let  $\mathbf{z} \in \mathbb{R}^d$  and  $\mathbf{x} = f(\mathbf{z}) \in \mathbb{R}^d$ . For rotation and reflection equivariance, the  
 1970 transformations  $T_g$  and  $S_g$  can be defined by a rotation matrix  $\mathbf{R} \in \mathbb{R}^{d \times d}$ ; then, a function  $f$  with  
 1971 the rotation equivariant property would satisfy  $\mathbf{R}\mathbf{x} = f(\mathbf{R}\mathbf{z})$ . For translation equivariance, the  
 1972 transformations would be  $T_g(\mathbf{z}) = \mathbf{z} + g$  and  $S_g(\mathbf{x}) = \mathbf{x} + g$ , where  $g \in \mathbb{R}^d$ . Hence, for a translation  
 1973 equivariance function  $f$ , we would have  $\mathbf{x} + g = f(\mathbf{z} + g)$ . For the case where the output dimension  
 1974 is larger than the input,  $f : \mathbb{R}^k \rightarrow \mathbb{R}^d$  with  $d > k$ , translation equivariance can be defined up to a  
 1975 discrete scale, i.e.,  $T_g(\mathbf{z}) = \mathbf{z} + g$  and  $S_g(\mathbf{x}) = T_{sg}(\mathbf{z})$  where  $s = d/k$ . The equivariance properties  
 1976 of translation, rotation, and reflections, combined, are referred to as E(3) symmetries. Without  
 1977 reflections, the symmetries form a Euclidean group SE(3) (Thomas et al., 2018; Fuchs et al., 2020).

1978 E(3), SE(3), and SO(3) are important symmetry groups in 3D Euclidean space, with well-established  
 1979 applications in physics and chemistry, computer vision, and reinforcement learning (Cohen and  
 1980 Welling, 2016; Thomas et al., 2018; Hoogeboom et al., 2022; Xu et al., 2024; Park et al., 2025).  
 1981 Finally, our contributions are complementary to, and can be combined with, the growing literature on  
 1982 meta-learning and automatic symmetry discovery to learn symmetry groups and their actions directly  
 1983 from data (Zhou et al.; Quessard et al., 2020; Dehmamy et al., 2021; Mohapatra et al., 2025).

1984 **Data manifold hypothesis.** Let data  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$  be in an ambient space of dimension  $d$  with  
 1985 support  $\mathcal{X}$  distribution. We assume that data are sampled from a low-dimensional manifold  $\mathcal{M}$  (Cay-  
 1986 ton et al., 2005; Ma and Fu, 2012) embedded in a high-dimensional space (Assumption H.1). This  
 1987 hypothesis is popular in machine learning (Bordt et al., 2023), and has been studied mathematically  
 1988 in the literature (Narayanan and Mitter, 2010; Bortoli, 2022). Moreover, empirical evidence in image  
 1989 processing supports the manifold hypothesis (Weinberger and Saul, 2006; Fefferman et al., 2016),  
 1990 and diffusion-based solvers assume this property (He et al., 2024; Chung et al., 2022b; 2023).

1991 **Assumption H.1 (Manifold Hypothesis).** Let  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$  be a data sample. The support  $\mathcal{X}$  of the  
 1992 data distribution lies on a  $k$  dimensional manifold  $\mathcal{M}$  within an ambient space  $\mathbb{R}^d$  where  $k \ll d$ .

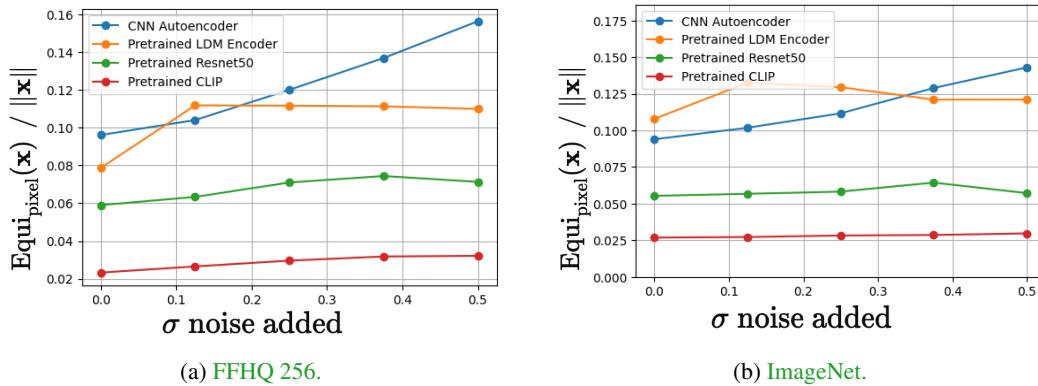
## 1993 I ADDITIONAL EXPERIMENTS ON MPE FUNCTIONS

1994 We compare several networks and show that MPE consistently emerges across them: as Gaussian  
 1995 noise is added to natural images, the equivariance loss systematically increases. We examine both (i)  
 1996 the emergence of MPE properties in different functions (neural networks) and (ii) the effect of using

1998 these functions within EquiReg on identical inverse problem settings. Specifically, for each dataset  
 1999 (FFHQ 256 and ImageNet), we consider four MPE function classes: (1) the pre-trained encoder  
 2000 of the latent diffusion model (LDM) used in our main experiments (Rombach et al., 2022), (2) a  
 2001 CNN autoencoder trained on the corresponding training distribution (FFHQ or ImageNet) with flip  
 2002 (FFHQ) or rotation (ImageNet) augmentations, (3) a pre-trained ResNet-50 (He et al., 2016), and (4)  
 2003 a pre-trained CLIP encoder (Radford et al., 2021). For each network, we evaluate equivariance loss  
 2004 under the relevant symmetry (flip for FFHQ, rotation for ImageNet) as Gaussian noise is added to  
 2005 100 natural images at increasing noise levels.

2006 Our results show that all four networks exhibit clear MPE behavior; their equivariance error increases  
 2007 as the noise level of the input grows. At the same time, the strength of the MPE property varies across  
 2008 architectures. Notably, the CNN autoencoder trained on the true data distribution shows the strongest  
 2009 MPE behavior, with equivariance error rising most sharply as images are corrupted, in line with  
 2010 our systematic guidelines for constructing MPE functions (Section 3). This is precisely the regime  
 2011 where the training distribution of the function matches the distribution of the inverse problem (e.g.,  
 2012 training on ImageNet train and evaluating on ImageNet test). In contrast, the LDM encoder exhibits  
 2013 the weakest MPE signal among the four, while ResNet-50 and CLIP fall between these extremes.  
 2014 These trends are visualized in Figures 15 and 20.

2015 We then apply each of these MPE functions within the EquiReg framework on the same inverse  
 2016 problem configurations: two datasets (FFHQ 256 and ImageNet), two diffusion-based solvers (DPS  
 2017 and SITCOM), and two tasks (super-resolution and motion deblurring). Across all settings and all  
 2018 MPE choices, EquiReg consistently improves reconstruction quality relative to the corresponding  
 2019 baseline without regularization (“None”). Tables 14 and 15 summarize these results. Taken together,  
 2020 these experiments demonstrate that (a) MPE properties naturally emerge in widely used pre-trained  
 2021 networks, making EquiReg easy to deploy in practice, and (b) EquiReg is robust across a range of  
 2022 MPE functions, including cases where the MPE property is relatively weak. Importantly, our main  
 2023 results use the LDM encoder which is the weakest MPE function in this ablation, suggesting that even  
 2024 stronger empirical gains are achievable using other MPE functions, such as the CNN autoencoder.  
 2025 We leave this as a future area of exploration.



2039 Figure 20: **Equivariance error vs.  $\sigma$  noise added.** As more noise is added, equivariance error,  
 2040 computed with all MPE functions, increases.

## J COMPUTING RESOURCES

2045 We conduct experiments on two NVIDIA GeForce RTX 4090 GPUs with 24 GB of VRAM. We note  
 2046 that we use pre-trained models and perform inference, so not much compute is required.

## K ASSETS

2048 We use the publicly available code from **PSLD** (<https://github.com/LituRout/PSLD>),  
 2049 **ReSample** (<https://github.com/soominkwon/resample>), **DPS** (<https://github.com>).

2052 Table 14: DPS superresolution with  $\lambda = 0.01$  using different MPE functions.  
2053

2054 (a) FFHQ 256.

MPE function	PSNR	SSIM	LPIPS	FID
None	23.160 (1.923)	0.657 (0.072)	0.193 (0.057)	129.528
LDM Encoder (FFHQ)	26.581 (2.457)	0.773 (0.044)	0.120 (0.030)	87.437
CNN Autoencoder (FFHQ)	26.866 (1.943)	0.771 (0.044)	0.116 (0.029)	85.352
Pretrained ResNet50	26.873 (1.941)	0.771 (0.044)	0.116 (0.029)	85.138
Pretrained CLIP	26.860 (1.942)	0.771 (0.044)	0.116 (0.029)	85.495

2062 (b) ImageNet.  
2063

MPE function	PSNR	SSIM	LPIPS	FID
None	19.727 (4.292)	0.407 (0.180)	0.541 (0.182)	446.829
LDM Encoder (ImageNet)	22.200 (4.295)	0.568 (0.146)	0.384 (0.130)	311.636
CNN Autoencoder (ImageNet)	22.178 (4.294)	0.568 (0.148)	0.375 (0.125)	312.530
Pretrained ResNet50	22.176 (4.290)	0.568 (0.148)	0.375 (0.125)	314.590
Pretrained CLIP	22.177 (4.293)	0.568 (0.148)	0.376 (0.125)	313.468

2071 Table 15: SITCOM motion deblurring on FFHQ 256 with  $\lambda = 0.05$  using different MPE functions.  
2072

MPE function	PSNR	SSIM	LPIPS
None	27.670 (1.343)	0.790 (0.031)	0.221 (0.040)
LDM Encoder (FFHQ)	28.357 (1.379)	0.806 (0.031)	0.200 (0.036)
CNN Autoencoder (FFHQ)	28.852 (1.376)	0.819 (0.044)	0.193 (0.033)
Pretrained ResNet50	28.682 (1.388)	0.811 (0.036)	0.198 (0.036)

2080 com/DPS2022/diffusion-posterior-sampling), and SITCOM (<https://github.com/sjames40/SITCOM>).  
20812083 

## L BROADER IMPACTS

2084  
2085 On the positive side, high-fidelity image restoration can improve downstream tasks in medical  
2086 imaging, remote-sensing and environmental monitoring (e.g., denoising satellite observations to  
2087 track pollution or deforestation). Likewise, accelerated PDE-solving via learned diffusion priors  
2088 may enable faster, more accurate simulations for climate modeling, fluid-dynamics research, and  
2089 engineering design. On the other hand, robust reconstruction methods could be misappropriated for  
2090 privacy-invasive surveillance or to create deceptive imagery. We emphasize that our method does not  
2091 amplify these existing risks.  
20922093 

## M RESPONSIBLE RELEASE

2094  
2095 Our approach uses only publicly available datasets and standard pre-trained diffusion models, intro-  
2096 ducing no novel dual-use or privacy risks. Consequently, no additional safeguards are required.  
2097