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, unstable 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 $\boldsymbol{x}^* \in \mathbb{R}^d$ from undersampled noisy measurements:

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

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

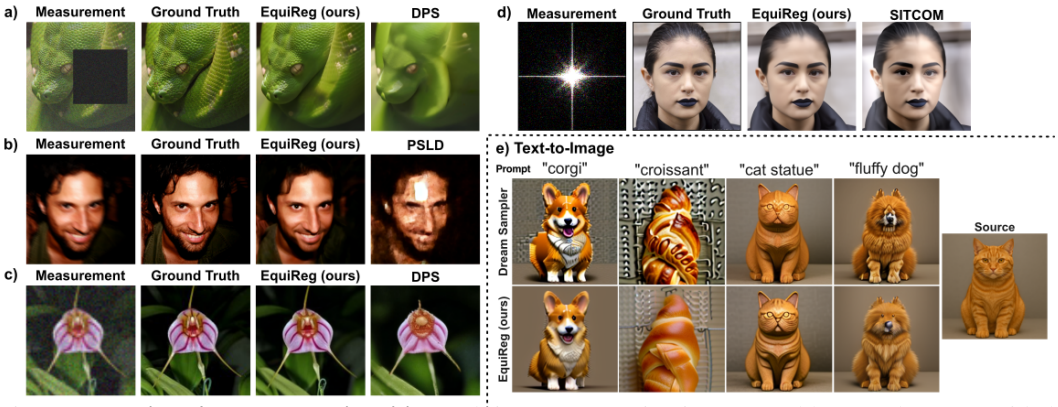

Figure 1: **EquiReg's broad applicability.** a-d) image restoration inverse problems and e) text-guided image generation, resulting in artifact reduction and more realistic generation. Here, EquiReg refers to our regularization being applied to the diffusion sampling method on the same row.

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

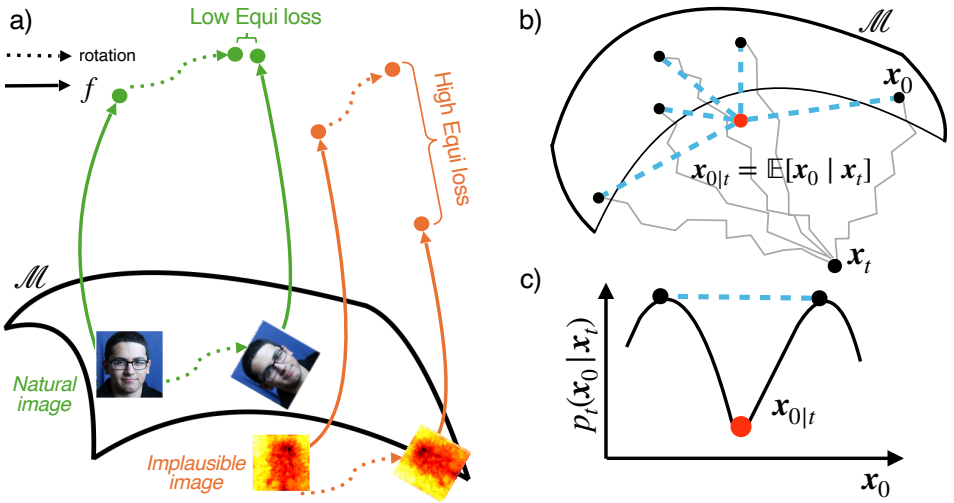

Figure 2: **Equivariance Regularized (EquiReg) diffusion for inverse problems.** a) EquiReg uses a stand-alone function whose equivariance error is lower for on-manifold and higher for off-manifold data. EquiReg regularizes the posterior sampling trajectory for improved performance. b-c) Illustration of off-manifold posterior expectation challenge via manifold and distribution interpretation.

solution maximizes the posterior distribution $p(\boldsymbol{x}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})$, where $p(\boldsymbol{y}|\boldsymbol{x})$ is the likelihood of the measurements and $p(\boldsymbol{x})$ 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 $\boldsymbol{x}_0 \sim p_{\text{data}} \to \boldsymbol{x}_t \to \boldsymbol{x}_T \sim p_T(\boldsymbol{x}) \approx \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$ and a reverse denoising process parameterized by a neural network score $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)$ (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 $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{y})$, decomposed into $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) + \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{y}|\boldsymbol{x}_t)$. This introduces challenges, as the likelihood score $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{y}|\boldsymbol{x}_t)$ is only computationally tractable when $t = 0$. To handle the likelihood for $t > 0$, many state-of-the-art methods approximate the posterior $p_t(\boldsymbol{x}_0|\boldsymbol{x}_t)$ 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 2c), 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 $\boldsymbol{x}_0$, it may lie off the data manifold even when individual samples remain on it (Figure 2b). 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.

Our contributions. We propose *Equivariance Regularized* (EquiReg) diffusion, an equivariance-based regularization framework for solving inverse problems with diffusion models (Figure 2a). 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 MPE functions 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 EquiReg as it cannot discriminate between on- and off-manifold samples. Instead, a separate non-equivariant architecture can be used to train to derive EquiReg for regularization.

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 improving 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 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 ℓ_2 error of FunDPS (Mammadov et al., 2024; 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; Kadkhodaie 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 via a reverse noising process. The forward noising process transforms the data sample $\boldsymbol{x}_0 \sim p_{\text{data}}$ via a series of additive noise into an approximately Gaussian distribution ($p_{\text{data}} \to p_t \to \mathcal{N}(0, I)$ as $t \to \infty$), described by the stochastic differential equation (SDE) $\mathrm{d}\boldsymbol{x} = -\frac{\beta_t}{2}\boldsymbol{x}\mathrm{d}t + \sqrt{\beta_t}\mathrm{d}\boldsymbol{w}$, where $\boldsymbol{w}$ is a standard Wiener process, and the drift and diffusion coefficients are parameterized by a monotonically increasing noise scheduler $\beta_t \in (0, 1)$ in time t (Ho et al., 2020). Reversing the forward diffusion process is described by (Anderson, 1982)

$$\mathrm{d}\boldsymbol{x} = [-\tfrac{\beta_t}{2}\boldsymbol{x} - \beta_t \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)]\,\mathrm{d}t + \sqrt{\beta_t}\mathrm{d}\bar{\boldsymbol{w}} \tag{2}$$

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

$$\mathrm{d}\boldsymbol{x} = [-\tfrac{\beta_t}{2}\boldsymbol{x}\mathrm{d}t - \beta_t(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) + \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{y}|\boldsymbol{x}_t))]\mathrm{d}t + \sqrt{\beta_t}\mathrm{d}\bar{\boldsymbol{w}} \tag{3}$$

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

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

Equivariance. Equivariance is a strategy for incorporating symmetries into deep learning (Bronstein et al., 2021). Prior work has applied equivariance to graph networks (Satorras et al., 2021), convolutional networks (Cohen and Welling, 2016; Romero and Lohit, 2022), Lie groups for modelling dynamical systems (Finzi et al., 2020), and diffusion models (Wang et al., 2024) with applications in molecular generation (Hoogeboom et al., 2022; Cornet et al., 2024), autonomous driving (Chen et al., 2023b), robotics (Brehmer et al., 2023), crystal structure prediction (Jiao et al., 2023), and audio inverse problems (Moliner et al., 2023). Equivariance guidance has also been used to improve temporal consistency in video generation (Daras et al., 2024). The benefits of equivariance as a prior to solve inverse problems (Scanvic et al., 2025) are theoretically supported in compressed sensing (Tachella et al., 2023). An equivariant function respects symmetries under group transformations, i.e.,

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

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

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

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

3 EquiReg: Equivariance Regularized Diffusion

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

This paper addresses the propagation error introduced by the approximation of posterior $p_t(\boldsymbol{x}_0|\boldsymbol{x}_t)$ by incorporating an explicit regularization term. The proposed framework is general and can be

applied as plug-in on a wide range of pixel and latent-space diffusion models. Given $p_t(\boldsymbol{y}|\boldsymbol{x}_t) = \int p(\boldsymbol{y}|\boldsymbol{x}_0)p_t(\boldsymbol{x}_0|\boldsymbol{x}_t)\mathrm{d}\boldsymbol{x}_0$, let $\tilde{p}_t(\boldsymbol{x}_0|\boldsymbol{x}_t)$ denote an approximation of the posterior to make the likelihood tractable. We formulate the regularized reverse diffusion dynamics as

$$\mathrm{d}\boldsymbol{x} = [-\tfrac{\beta_t}{2}\boldsymbol{x}\mathrm{d}t - \beta_t\nabla_{\boldsymbol{x}_t}(\log p_t(\boldsymbol{x}_t) + \log \textstyle\int p(\boldsymbol{y}|\boldsymbol{x}_0)\tilde{p}_t(\boldsymbol{x}_0|\boldsymbol{x}_t)\mathrm{d}\boldsymbol{x}_0 - \mathcal{R}(\boldsymbol{x}_t))]\mathrm{d}t + \sqrt{\beta_t}\mathrm{d}\bar{\boldsymbol{w}}, \quad (4)$$

where $\mathcal{R}(\boldsymbol{x}_t)$ is the regularizer. Applying this to DPS (Chung et al., 2023), as an example, takes the form in Algorithm 1). This formulation raises two questions: i) how to design the regularizer, and ii) how to interpret the role of $\mathcal{R}$ in regularizing conditional diffusion models and its impact on the sampling trajectory. We gain insight into the desirable properties of an optimal regularizer by reinterpreting the reverse conditional diffusion process as a time-inhomogeneous Wasserstein gradient flow (Ferreira and Valencia-Guevara, 2018) (see Propositions H.1 and H.2 in Appendix).

Algorithm 1 Equi-DPS for Inverse Problems.

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

This connection suggests that an effective regularizer, while applied locally to each sample in each iteration, can be interpreted as penalizing trajectories deviating from the data manifold, i.e., exhibiting high error for off- or far-manifold samples, and low error for on- or near-manifold ones. This motivates designing a regularizer that corrects the entire functional being minimized globally, in contrast to prior works that focus only on locally reducing likelihood error. We focus on equivariance as a global property that enforces geometric symmetries and guides the diffusion process toward the data manifold. To realize this idea, we seek functions that exhibit approximate equivariance and discriminate on- from off-manifold samples.

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

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

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

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

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. 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).

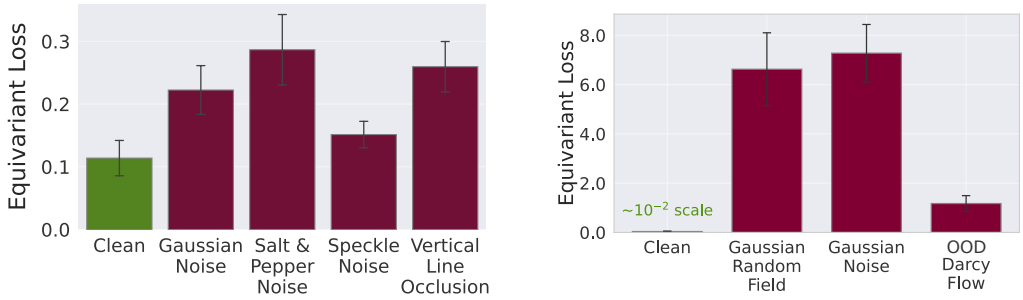

(a) Equivariance from augmentation. (b) Equivariance from data symmetries.

Figure 3: **MPE function examples.**

Figure 3a 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:

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

where $\boldsymbol{x}_{0|t}$ and $\boldsymbol{z}_{0|t}$ are function of $\boldsymbol{x}_t$ and $\boldsymbol{z}_t$, 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 3b 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 $\|S_g(\text{FNO}(\boldsymbol{x}_{0|t})) - \text{FNO}(T_g(\boldsymbol{x}_{0|t}))\|_2^2$ on in-distribution as opposed to out-of-distribution data, with reflection as the group action.

Table 1: **EquiReg improves SITCOM under reduced measurement consistency steps (K_{meas}).** We reduce K_{meas} and add an equal amount of EquiReg steps (K_{EquiReg}). Evaluated on motion deblur for FFHQ sampled with 50 DDIM steps.

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

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 other methods including MCG (Chung et al., 2022b), MPGD-AE (He et al., 2024), and DiffStateGrad (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 256×256 validation dataset (Karras et al., 2021), and b) the ImageNet 256×256 validation dataset (Deng et al., 2009). For pixel-based experiments, we use i) the pre-trained model from (Chung et al., 2023)

Table 2: **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.

		Super Resolution				Gaussian Blur			
Method	**Period**	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

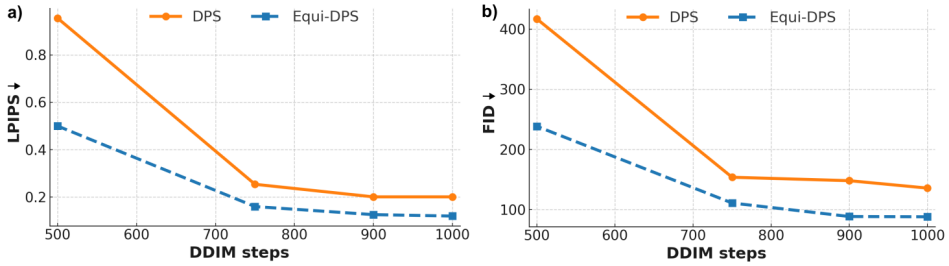

Figure 4: **Advantages of EquiReg under reduced DDIM steps.** Super-resolution on FFHQ.

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.

We evaluate EquiReg on a variety of linear and nonlinear restoration tasks for natural images (see Appendix F for task details). We adopt the pre-trained encoder-decoder $\mathcal{E}$-$\mathcal{D}$ 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.

First, we show that adding EquiReg optimization steps consistently enables SITCOM to achieve superior performance with significantly faster runtime using fewer measurement consistency steps (Table 1). 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 4). We also show that EquiReg is able to preserve performance when applied with lower frequency (Table 2).

Table 3: **EquiReg for ReSample on linear and nonlinear tasks.** FFHQ 256×256 with $\sigma_y = 0.01$.

Task	**Method**	**LPIPS↓**	**FID↓**	**PSNR↑**	**SSIM↑**
Linear					
Gaussian deblur	ReSample	0.253	55.65	27.78	0.757
	Equi-ReSample	0.197	64.86	**29.08**	**0.825**
	EquiCon-ReSample	**0.156**	**54.72**	28.18	0.777
Motion deblur	ReSample	0.160	40.14	30.55	0.854
	Equi-ReSample	0.120	46.28	**30.92**	**0.870**
	EquiCon-ReSample	**0.078**	**37.61**	30.73	0.860
Super-res. (×4)	ReSample	0.204	40.46	28.02	0.790
	Equi-ReSample	**0.098**	43.56	**29.74**	**0.849**
	EquiCon-ReSample	0.112	**40.38**	28.27	0.801
Box inpainting	ReSample	0.198	108.30	19.91	0.807
	Equi-ReSample	**0.150**	**59.69**	**22.56**	**0.832**
	EquiCon-ReSample	0.171	110.70	21.04	0.815
Random inpainting	ReSample	0.115	36.12	31.27	0.892
	Equi-ReSample	**0.047**	29.88	**31.47**	**0.908**
	EquiCon-ReSample	**0.047**	**28.81**	31.21	0.904
Nonlinear					
HDR	ReSample	0.190	**49.06**	**24.88**	**0.819**
	Equi-ReSample	**0.133**	49.52	24.71	0.815
	EquiCon-ReSample	0.135	49.98	24.67	0.817
Phase retrieval	ReSample	0.237	97.86	27.61	0.750
	Equi-ReSample	**0.155**	**85.22**	**28.16**	0.770
	EquiCon-ReSample	0.159	88.75	28.11	**0.774**
Nonlinear deblur	ReSample	0.188	56.06	29.54	0.842
	Equi-ReSample	0.128	55.09	29.45	0.840
	EquiCon-ReSample	**0.125**	**54.62**	**29.55**	**0.843**

Highlighting the benefits of EquiReg for latent diffusion models, Table 3, Table 4a, and Table 5 show that EquiReg consistently improves 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).

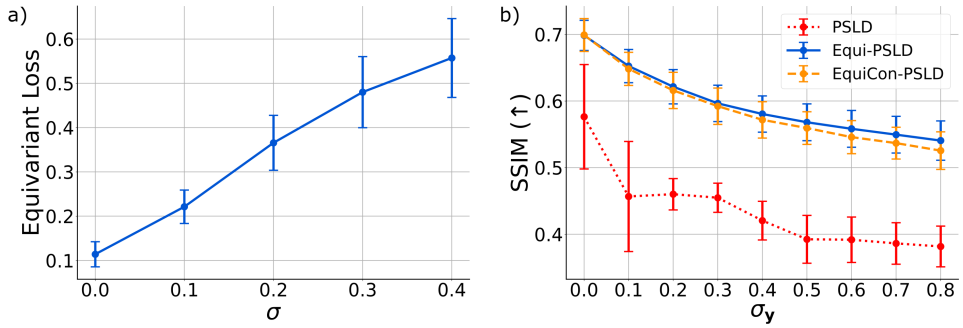

Figure 5: **EquiReg is effective across a range of measurement noise levels.** a) Equivariance error computed over a pre-trained decoder on increasingly noisy inputs. b) EquiReg performance computed over a range of measurement noise levels on the FFHQ dataset.

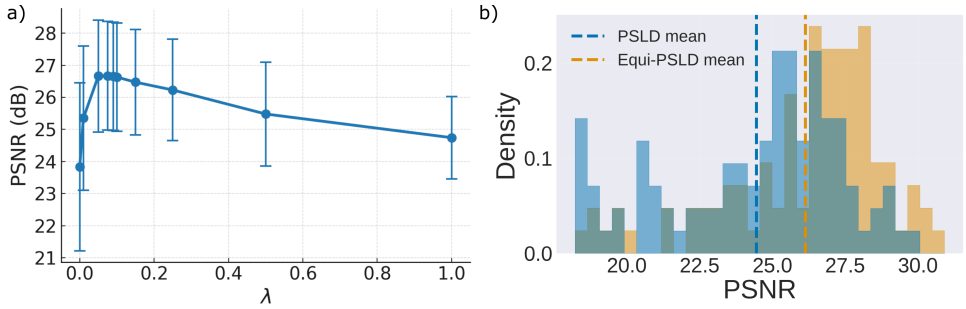

Figure 6: **Robustness of EquiReg, demonstrated on PSLD.** a) EquiReg is robust to the choice of λ_t. b) EquiReg reduces failure cases and enhances reconstruction fidelity for super-resolution on FFHQ.

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 λ_t (Figure 6a, see Appendix B for details).

Solving PDEs from sparse observations. EquiReg is evaluated on two important PDE problems: the Helmholtz and Navier-Stokes equations (see Appendix G). 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.

Table 6: **Solving PDEs from sparse observations.**

	Steps (N)	Helmholtz		Navier-Stokes	
		Forward	Inverse	Forward	Inverse
DiffusionPDE	2000	12.64%	19.07%	3.78%	9.63%
FunDPS	500	2.13%	17.16%	3.32%	8.48%
Equi-FunDPS (ours)	500	**2.12%**	**15.91%**	**3.06%**	**7.84%**

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

We integrate EquiReg into the state-of-the-art FunDPS framework (Mammadov et al., 2024; Yao et al., 2025), where we compute the Equi loss with respect to equivariance learned by an FNO trained on the corresponding inverse problem. We use reflection symmetry (i.e., flipping along the $y = x$ axis), and observe no significant performance difference when using other transformations such as rotations or alternating flips. Equi-FunDPS improves performance (Table 6), measured by relative ℓ_2 loss, across various tasks, especially in inverse problems where a strong data prior is critical.

Table 4: **EquiReg for diffusion models on FFHQ**. 256×256 with $\sigma_{\boldsymbol{y}} = 0.05$.

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↑
PSLD	0.357	106.2	22.87	**0.322**	**84.62**	24.25	0.313	89.72	24.51	0.158	43.02	24.22	0.246	49.77	29.05
Equi-PSLD	0.344	94.09	**24.42**	0.338	99.14	24.83	0.289	90.88	**26.32**	0.098	**31.54**	24.19	**0.188**	41.61	**30.43**
EquiCon-PSLD	**0.320**	**83.18**	24.38	**0.322**	89.87	**25.14**	**0.277**	**79.39**	26.14	**0.092**	35.07	**24.26**	0.204	**40.75**	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↑
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)	**0.114**	**48.76**	**26.32**	**0.094**	**41.71**	**28.23**	**0.120**	**51.00**	**27.15**	**0.099**	40.47	23.39	**0.068**	33.65	**32.16**
DiffStateGrad-DPS	0.128	52.73	26.29	0.118	50.14	27.61	0.186	73.02	24.65	0.114	47.53	**24.10**	0.107	49.42	30.15
MCG	0.340	101.2	6.72	0.702	310.5	6.72	0.520	87.64	20.05	0.309	**40.11**	19.97	0.286	**29.26**	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.

Table 5: **EquiReg for latent diffusion models on ImageNet**. 256×256 with $\sigma_{\boldsymbol{y}} = 0.05$.

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	**214.5**	**22.01**	**196.3**	**22.69**	**198.5**	**22.34**	**137.6**	**19.25**	**65.14**	**27.03**

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

5 Conclusion

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

Limitations and future work. EquiReg's effectiveness depends on the quality of the pre-trained backbone diffusion. EquiReg is a plug-and-play regularization framework that can be applied to a variety of guidance-based diffusion models; thus, it does not directly address the approximations of the underlying diffusion models, but instead regularizes for improved performance. Also, since EquiReg is a regularization mechanism, it improves performance precisely in regimes where baseline methods degrade or fail. Hence, one cannot expect EquiReg to improve the performance of a diffusion model beyond the capability of a regularization framework. Finally, applying EquiReg requires task-specific design choices: selecting an appropriate symmetry group and identifying suitable MPE functions for the problem at hand. While we presented two systematic approaches to construct MPE functions for imaging and PDEs, the process of identifying MPE functions varies across applications and represents an important area for methodological development. This task-specific design also makes EquiReg broadly adaptable across diverse domains beyond the considered applications.

REFERENCES

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

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

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

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

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

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

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

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

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

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

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

Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. doi: 10.1162/NECO_a_00142.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Abbas Mammadov, Julius Berner, Kamyar Azizzadenesheli, Jong Chul Ye, and Anima Anandkumar. Diffusion-based inverse solver on function spaces with applications to pdes. *Machine Learning and the Physical Sciences Workshop at NeurIPS*, 2024. URL https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_253.pdf.

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

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

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

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

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

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

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

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

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

Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-Guided Diffusion Models for Inverse Problems. 2023b. URL https://openreview.net/forum?id=9_gsMA8MRKQ.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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