DIFFOPS: EFFICIENTLY SOLVING DIFFUSION INVERSE PROBLEMS WITH OPERATOR-AWARE OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models provide expressive priors that can improve reconstruction from incomplete or noisy measurements for ill-posed inverse problems. However, existing methods often apply a uniform measurement-correction procedure across diverse inverse problems, despite substantial differences in measurement models and numerical requirements. To address this limitation, we introduce DiffOps, a flexible framework that integrates operator-aware optimization into diffusion-based reconstruction. In this framework, diffusion first provides learned prior information to guide reconstruction where measurements alone cannot fully determine the signal, after which numerical corrections exploit the available operator structure to enforce measurement consistency. As a result, classical operator-selection principles can be integrated directly within diffusion iterations according to the numerical demands of each specific problem. Through experiments on both image restoration and scientific imaging tasks, we demonstrate that DiffOps achieves competitive or superior reconstruction quality and robustness while requiring fewer diffusion-network evaluations.

1 INTRODUCTION

Diffusion models provide implicit learned priors for signal reconstruction in ill-posed inverse problems across diverse acquisition modalities. A pretrained diffusion prior can be reused across different measurement operators without task-specific retraining (Song et al., 2021, Wu et al., 2024b). Existing approaches have combined diffusion with different correction strategies, including spectral updates, nullspace projections, and approximate likelihood guidance. These methods have demonstrated their effectiveness in image restoration by combining learned priors with a specified measurement operator and observation model (Kawar et al., 2022, Wang et al., 2023, Chung et al., 2023a). However, measurement operators often have distinct physical and numerical properties across applications, whereas many existing methods rely on a common correction procedure across different inverse problems.

As diffusion-based solvers are applied to an increasingly broad range of scientific inverse problems, they encounter forward models with distinct structures and numerical challenges (Zheng et al., 2025b, Chung et al., 2023b, Feng et al., 2023, Zheng et al., 2025a). A single diffusion-based correction strategy is therefore unlikely to suit all classes of problems. Classical optimization addresses this issue by choosing algorithms according to the physical and numerical properties of each problem (Rice, 1976, Nocedal & Wright, 2006). This motivates bringing the same operator-aware principle into diffusion-based reconstruction.

Recent decoupled approaches offer an alternative framework for diffusion-based reconstruction by separating the denoising step from measurement correction, without requiring the diffusion step to follow the time-noise trajectory (Zhang et al., 2025, Zhu et al., 2023, Chen et al., 2026, Wu et al., 2024a), which makes operator-aware correction practical. Specifically, this separation allows each correction step to be tailored to the measurement operator and the noise model. Building on this observation, we introduce DiffOps, an operator-aware framework for diffusion inverse problems. At each noise level, the diffusion model provides a reference estimate that initializes and anchors the numerical correction, while the measurement model and available operator information determine the appropriate update. Spectral structure, derivative availability, and nonlinearity guide the solver selection, while the noise model determines the data-fidelity term (Nocedal & Wright, 2006). After

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

sampling, the pretrained diffusion prior provides the regularization for this local problem while the outer denoising process remains decoupled from the local optimization for flexible jumps across diffusion timesteps. The regularization weight controls adherence to the diffusion estimate, and the solver budget controls the numerical effort devoted to measurement correction.

We evaluate DiffOps on synthetic operators with FFHQ benchmarks and scientific imaging problems spanning linear and nonlinear operators under realistic measurement noise. DiffOps achieves competitive or improved reconstruction quality and robustness while requiring fewer diffusion-network evaluations. More broadly, our results show that diffusion-based inverse problem reconstruction does not need to rely on a single numerical procedure across all inverse problems. A learned diffusion prior can instead serve purely as the prior, with corrections adapted to each measurement operator. This aligns modern diffusion-based reconstruction with classical algorithm-selection principles.

Our contributions are summarized as follows:

1. **Operator-aware formulation.** We extend the decoupled reconstruction perspective by using problem-dependent numerical solver selection, grounded in the measurement model and operator structure, within a shared iteration pipeline.
2. **Flexible numerical correction.** Based on analytical analysis and numerical experiments, we accommodate different solvers according to the local requirements and illustrate where diffusion helps in these iterative updates.
3. **Evaluation across operator classes.** We demonstrate competitive or improved reconstruction quality and robustness with fewer diffusion-network function evaluations (NFE), on synthetic benchmarks and scientific imaging problems spanning different types of operators.

2 Preliminaries and Related Work

Solving inverse problems with diffusion inference. For the inverse problem of recovering a clean signal x from measurements $y \sim p(y \mid x)$ via inference under the posterior $p(x \mid y)$, Bayes' rule gives $p(x \mid y) = p(y \mid x)p(x)/p(y)$. The acquisition model specifies the likelihood for a candidate x and defines the data-fidelity term

$$\mathcal{L}(x) = -\log p(y \mid x). \tag{1}$$

With a pretrained diffusion prior, posterior sampling requires the score of the noise-level posterior $p_t(x_t \mid y)$. Here, $x_t = x + \sigma_t \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, I)$, is the Gaussian-corrupted signal, and p_t denotes its prior marginal density. Under this diffusion process, Bayes' rule can be reformulated at the score level as

$$\nabla_{x_t} \log p_t(x_t \mid y) = \nabla_{x_t} \log p_t(x_t) + \nabla_{x_t} \log p_t(y \mid x_t). \tag{2}$$

The pretrained diffusion model provides an estimate of the prior score function (Song & Ermon, 2019, Ho et al., 2020, Song et al., 2021, Karras et al., 2022). However, the noise-level likelihood $p_t(y \mid x_t)$ should, in theory, be calculated by averaging the measurement likelihood over the conditional distribution $p(x_0 \mid x_t)$ (Chung et al., 2023a). Evaluating this average and its logarithmic gradient with respect to x_t is generally intractable (Daras et al., 2024).

A practical approximation replaces this conditional average with the measurement likelihood evaluated at a single clean-signal estimate. Under Gaussian corruption, one choice is a learned denoiser that approximates the conditional mean given by the Tweedie identity (Chung et al., 2023a):

$$\hat{x}_0(x_t) = D_\theta(x_t, \sigma_t) \approx \mathbb{E}[x_0 \mid x_t]. \tag{3}$$

The estimate $\hat{x}_0(x_t)$ can also be obtained by numerically integrating the probability-flow ODE from x_t (Zhang et al., 2025, Lu et al., 2025). DiffOps switches to this ODE estimate below a noise level $\bar{\sigma}$ (Appendix A.1).

Thus, we approximate the noisy likelihood $p_t(y \mid x_t)$ by $p(y \mid \mathcal{A}(\hat{x}_0(x_t)))$ for a given forward operator $\mathcal{A}$ and reformulate the data-fidelity term as $\mathcal{L}(x) = -\log p(y \mid \mathcal{A}(x))$. This gives the approximate conditional score:

$$\nabla_{x_t} \log p_t(x_t \mid y) \approx \nabla_{x_t} \log p_t(x_t) - \nabla_{x_t} \mathcal{L}\big(\hat{x}_0(x_t)\big). \tag{4}$$

This means that the measurement gradient is propagated through the estimate $\hat{x}_0(x_t)$ after denoising.

Separating the prior update from the measurement correction. Variable-splitting and decoupled annealing methods directly incorporate the measurement-consistency term into clean-signal space during each correction step, together with the prior term (Zhu et al., 2023, Wu et al., 2024a, Zhang et al., 2025). Given the denoised estimate $\hat{x}_0(x_t)$ at a fixed noise level, a generic correction step combines the measurement gradient with the conditional prior score:

$$x^{(j+1)} = x^{(j)} + \eta_t \left[\nabla_x \log p\big(x^{(j)} \mid x_t\big) - \nabla_x \mathcal{L}\big(x^{(j)}\big) \right] + \xi^{(j)}. \tag{5}$$

Here, η_t is the step size, and $\xi^{(j)}$ is an optional stochastic noise term. The basic premise of these decoupled structures is to retain the time-marginal distribution $p_t(x_t \mid y)$ of the diffusion at each level t. Under an isotropic Gaussian approximation centered at $\hat{x}_0(x_t)$, the prior score becomes proportional to $-\big(x^{(j)} - \hat{x}_0(x_t)\big)$. Sampling-based methods (Wu et al., 2024a, Zhang et al., 2025) can potentially recover variability by applying an MCMC step, thereby exploring the local posterior. In contrast, optimization-based methods (Xu et al., 2025, Zhu et al., 2023) seek the minimizer of the objective function:

$$x_0 \approx \arg\min_x \left\{ \mathcal{L}(x) + \frac{\lambda_t}{2} \|x - \hat{x}_0(x_t)\|^2 \right\}, \tag{6}$$

in which $\lambda_t > 0$ controls adherence to the initial estimate from prior information relative to the measurement loss. Fully decoupled formulations go further by relaxing the time-marginal consistency imposed by the annealed flow (Chen et al., 2026) and treating measurement correction as a standalone numerical problem.

The same separation of the prior and data terms has been applied across a range of reconstruction problems (Venkatakrishnan et al., 2013), with the numerical form of the correction varying across them. Denoiser-based splitting, from plug-and-play onward, uses FFT, conjugate-gradient, or proximal updates (Zhang et al., 2022, Kamilov et al., 2023), and ensemble-based diffusion methods use forward evaluations to approximate black-box models (Zheng et al., 2025a, Wang et al., 2025). Benchmark comparisons across various scientific problems (Zheng et al., 2025b) reveal that a general diffusion inverse problem method effective on one class of forward models can lose accuracy on another.

Algorithm selection and numerical structure. For a fixed prior estimate, the local update in Eq. 6 has the classical form of Tikhonov regularization, pulling the solution toward the reference state $\hat{x}_0$ (Engl et al., 1996, Hansen, 1998). Accordingly, solving this regularized inverse problem with a classical numerical algorithm relies on the structure of the forward model under the operator $\mathcal{A}$ and the likelihood update (Rice, 1976, Nocedal & Wright, 2006). As such, the local optimal solution for a linear Gaussian model with noise scale σ_y satisfies

$$\big(\mathcal{A}^*\mathcal{A} + \lambda_t \sigma_y^2 I\big) x_0^* = \mathcal{A}^* y + \lambda_t \sigma_y^2 \hat{x}_0, \tag{7}$$

where $\mathcal{A}^*$ is the adjoint, and increasing λ_t improves the conditioning while shifting the weight between the measurements and the reference estimate.

The structure of Eq. 7 provides general guidance for algorithm selection: when the normal operator has a tractable diagonalization, it can be solved spectrally using a Wiener filter (Wiener, 1949, Hansen, 1998); conjugate gradients are useful when only matrix-free forward and adjoint products are available (Shewchuk, 1994), making the method computationally practical for large linear systems (Nocedal & Wright, 2006). For nonlinear operators, the local linearization varies with the current iterate, so the analysis depends on the local geometry. The Jacobian $J(x) = \partial \mathcal{A}(x)/\partial x$ describes the local sensitivity, while $J^\mathsf{T} J/\sigma_y^2$ provides an approximate Hessian (Nocedal & Wright, 2006). When direct derivatives are unavailable, such as for a black-box forward model, derivative-free methods use forward evaluations to construct a tractable local approximation (Iglesias et al., 2013, Larson et al., 2019).

As diffusion annealing changes the reference estimate and regularization weight, each noise level presents a distinct numerical problem. The solver therefore follows the forward operator and its numerical structure rather than a generic prescription. The numerical method and computational budget can then be chosen according to the operator and the requirements at that noise level.

3 METHOD

Built on the decoupled formulation (Zhang et al., 2025, Chen et al., 2026), we use Eq. 6 as the core correction rule to define a standalone operator-aware optimization strategy as shown in the pipeline in Fig. 1. At each diffusion level, the denoiser reference and regularization weight are held fixed during the measurement-correction solve. The regularization scale depends on the diffusion prior and the current noise level, and on whether noise corrupts the range space of $\mathcal{A}$.

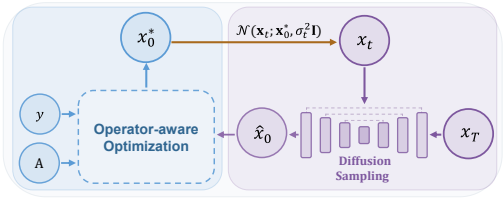

Figure 1: Overview of DiffOps. The diffusion stage supplies the reference $\hat{x}_0$; operator-aware optimization corrects it into x_0^* with y and $\mathcal{A}$.

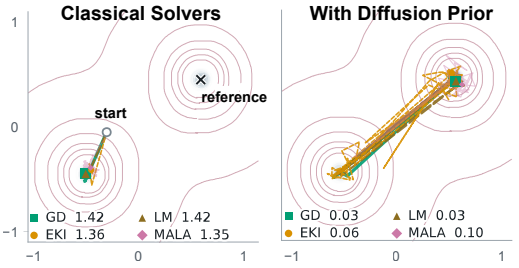

Figure 2: Classical solvers and the same solvers with diffusion on a two-dimensional nonconvex problem; numbers show the distance to the true point.

3.1 THE LOCAL DATA UPDATE GIVEN A DIFFUSION PRIOR

Diffusion regularization. At diffusion level t, the local correction optimizes x by taking the denoised estimate $\hat{x}_0(x_t)$ as the starting point after diffusion sampling from the noisy state x_t, followed by operator-aware optimization under measurement y and operator $\mathcal{A}$, given by:

$$F_t(x) = \mathcal{L}(x) + \frac{\lambda_t}{2}\|x - \hat{x}_0(x_t)\|^2. \tag{8}$$

Diffusion therefore provides two forms of regularization. First, implicit regularization arises from the warm start at $\hat{x}_0(x_t)$, which supplies the components that the measurements cannot determine. Second, diffusion provides explicit regularization through a local constraint centered on $\hat{x}_0(x_t)$, which limits the measurement noise that the correction can inject. We view this term as a local constraint that keeps the corrected solution in the neighborhood of $\hat{x}_0(x_t)$, rather than merely as a heuristic approximation to the score of the diffusion prior. Accordingly, the coefficient λ_t determines the relative contribution of this local regularization constraint with respect to the prescribed data-fidelity term. Given $\hat{x}_0(x_t)$, the corrected estimate x_0^* is the output of the selected solver on this objective and need not be its exact minimizer, because intermediate outputs are re-noised before the next stage and exact minimization matters mainly at the final noise level (Appendix A.1). We therefore use a small, fixed number of solver steps at every level, rather than solving each subproblem to convergence.

Uncertainty at each noise level. At fixed x_t, following the Gaussian approximations of $p(x_0 \mid x_t)$ in prior work (Song et al., 2023, Zhang et al., 2025), we characterize the sampling uncertainty as $\mathcal{N}\big(\hat{x}_0(x_t), \sigma_t^2 I/\beta\big)$, where β is assumed constant across all noise levels t. Accordingly, the negative log-density, $\beta\|x - \hat{x}_0(x_t)\|^2/(2\sigma_t^2)$, induces the local explicit regularization, yielding

$$\lambda_t = \frac{\beta}{\sigma_t^2}. \tag{9}$$

Here, β is calibrated from the diffusion prior and held fixed throughout the diffusion annealing process.

For a Gaussian inverse problem, the normalized data-fidelity term is given by $\mathcal{L}(x) = \|\mathcal{A}(x) - y\|^2/(2\sigma_y^2)$. Combining this with the regularization scale λ_t yields the relative weighting $\beta\sigma_y^2/\sigma_t^2$. Consequently, a larger β reflects lower uncertainty in the diffusion reference and more strongly constrains the iterate toward the denoiser estimate. In practice, β depends on the diffusion prior and on whether noise corrupts the range space of $\mathcal{A}$: noiseless measurements permit $\beta = 0$, whereas noisy ones need $\beta > 0$ to avoid fitting the noise.

3.2 OPERATOR-AWARE OPTIMIZATION WITH DIFFUSION ITERATIONS

Local iteration update. At each diffusion level, the warm-start estimate $\hat{x}_0(x_t)$ and the weight λ_t jointly determine the measurement correction. According to Eq. 9, the explicit regularization weight varies with the diffusion noise level, thereby permitting larger measurement-driven corrections at noisier stages while more strongly anchoring the estimate to the diffusion reference at lower-noise stages. Moreover, even under the same diffusion-scheduler annealing step, the inner optimization at each σ_t may employ different correction methods depending on the requirements of the specific inverse problem. Consequently, the measurement operator determines how Eq. 8 is solved locally.

The explicit regularization term therefore contributes $\lambda_t I$ to the local curvature. If L_{data} bounds the smoothness of the data-fidelity term, a conservative first-order update for the inner solver in Eq. 5 is given by

$$x^{(j+1)} = x^{(j)} - \eta_t \left[\nabla \mathcal{L}(x^{(j)}) + \lambda_t \big(x^{(j)} - \hat{x}_0(x_t) \big) \right], \qquad \eta_t \leq \frac{1}{L_{\mathrm{data}} + \lambda_t}. \tag{10}$$

Thus, the diffusion schedule governs not only the strength of the reference constraint but also the numerical scale of the correction step. As the noise level decreases and λ_t increases, the correction becomes more strongly anchored to $\hat{x}_0(x_t)$, while the admissible first-order step size correspondingly decreases. For other solver families, the same λ_t enters through shifted linear systems, proximal updates, or local curvature models.

Operator-aware solver selection. Classical solvers can be selected according to the analytical and numerical structure of the forward operator. For diagonalizable linear operators, Wiener filtering provides a closed-form spectral solution. When explicit access to $\mathcal{A}$ and $\mathcal{A}^*$ is available, matrix-free conjugate gradient (CG) methods can be applied in a Krylov subspace. For more general operators, gradient descent (GD) provides general-purpose first-order updates combined with the diffusion prior, whereas available Jacobian or curvature information can be used by higher-order methods, including quasi-Newton (L-BFGS) (Liu & Nocedal, 1989) and Gauss–Newton (GN)/Levenberg–Marquardt (LM) (Levenberg, 1944, Marquardt, 1963) algorithms. When the forward model is treated as a black box, derivative-free or ensemble methods, such as ensemble Kalman inversion (EKI), can provide effective approximations based on empirical means and covariances. Moreover, operator-splitting frameworks such as the alternating direction method of multipliers (ADMM) (Boyd et al., 2011) can separate the data-fidelity term from the regularizer through alternating proximal updates. Non-Gaussian observation models, such as Poisson noise, can keep their exact likelihood through the Poisson negative log-likelihood, with expectation-maximization (EM) iterations (Richardson, 1972, Shepp & Vardi, 1982) used for optimization rather than Gaussian approximations. These approaches are different numerical solutions of the same fixed-reference correction objective; explicit solver derivations are provided in Appendix A.2.

The diffusion reference supplies prior information when measurement fitting alone is insufficient. Figure 2 illustrates this in a two-dimensional nonconvex problem with symmetric measurement ambiguity: classical runs remain in the basin opposite the ground truth, whereas annealed runs converge near it. Appendix B.3 examines additional initializations and seeds, distinguishing basin selection from regularization of noisy or weakly constrained estimates. For unconstrained linear Gaussian subproblems, unpreconditioned CG initialized at $\hat{x}_0(x_t)$ preserves the reference nullspace component at every well-defined iterate in exact arithmetic. This result holds even with finite inner iterations and $\lambda_t = 0$, thereby making the role of initialization evident. Inner iterations refine the current subproblem, whereas diffusion evaluations update the reference.

3.3 NONLINEAR OPERATORS

For a nonlinear forward operator $\mathcal{A}$, the local measurement sensitivity varies with the current iterate. Under a Gaussian likelihood, first-order updates use $J(x)^{\mathsf{T}}(\mathcal{A}(x) - y)/\sigma_y^2$, whereas Gauss–Newton updates additionally incorporate the curvature model $J(x)^{\mathsf{T}}J(x)/\sigma_y^2 + \lambda_t I$. This directional scaling can help resolve poorly conditioned local subproblems, though its effectiveness depends on the linearization accuracy and the solver budget. Appendix B.2 provides a detailed analysis of this distinction.

Global ambiguity is common in phase retrieval problems, where improving a local fit does not guarantee converging to the desired solution basin. Consequently, evaluating the success rate of

reaching the target basin relative to the reference is important. For this reason, our framework adopts a mixed-solver strategy that moves from global exploration at high noise levels to local refinement as the noise decreases; in practice, the switch is made at a fixed noise level. Pointwise clipping in high-dynamic-range (HDR) imaging presents a distinct challenge, as saturated coordinates yield zero local measurement gradients for nonlinear problems. When minimizing an uncoupled objective initialized at the reference point, both gradient and diagonal Gauss–Newton steps leave these saturated coordinates unchanged during the inner solve. Alternatively, ensemble methods can capture a measurement response via finite perturbations that cross the clipping threshold into the unsaturated regime.

Together, these cases show that the correction mechanism should reflect the operator's available information and constraints while preserving a unified reference-based interface. The solver choice resolves the local numerical subproblem, whereas the prior guides the reconstruction in directions where the measurements provide insufficient information.

4 EXPERIMENTS

We evaluate DiffOps in two complementary settings. Image-restoration benchmarks test the operator-aware correction with the same FFHQ pretrained prior (Karras et al., 2019), and scientific imaging tasks test it under domain-specific acquisition physics and task-specific priors (Zheng et al., 2025b).

Datasets and priors. The image-restoration benchmark uses 100 held-out FFHQ images at 256×256 (Karras et al., 2019), and a single pretrained unconditional FFHQ diffusion model is used for all operators without task-specific retraining. The scientific benchmarks cover four tasks: 96 fastMRI knee slices (Zbontar et al., 2018), 100 synthetic black-hole images per observation-coverage level, 100 fluorescence microscopy images per point-spread function (PSF) and photon level, and 100 permittivity fields for inverse scattering. Each scientific task uses its own pretrained diffusion prior. Dataset preprocessing, forward-model configurations, and normalization are described in Appendices C and E.

Measurement generation. The FFHQ operators are box inpainting, $4\times$ bicubic downsampling, Gaussian and motion blur, pointwise saturation, network-based nonlinear blur (Tran et al., 2021), and Fourier phase retrieval, following established protocols in a single run (Chung et al., 2023a, Zhang et al., 2025). Gaussian measurement noise has $\sigma_y = 0.05$ on $[-1, 1]$, and phase-retrieval noise is specified on Fourier magnitudes in $[0, 1]$. MRI uses a multi-coil model with fourfold acceleration ($R = 4$) on both simulated k-space and raw scanner data. Black-hole imaging (Chael et al., 2018, Event Horizon Telescope Collaboration, 2019) uses gain-invariant closure phases and log closure amplitudes at 3%, 10%, and 100% observation coverage. Poisson deconvolution uses about 10 and 100 expected photons at full intensity on FFHQ and microscopy images. Inverse scattering uses a dense Green's-function operator with additive complex Gaussian noise (Zheng et al., 2025b).

Baselines and configurations. On FFHQ, we compare DiffOps with DPS, DAPS, and DAPS++ (Chung et al., 2023a, Zhang et al., 2025, Chen et al., 2026), and with DDRM, DDNM, DiffPIR, and DCDP (Kawar et al., 2022, Wang et al., 2023, Zhu et al., 2023, Li et al., 2024). On the scientific tasks, we compare with DPS, DAPS, PnP-DM, and RED-diff (Wu et al., 2024a, Mardani et al., 2024). All baselines use their official implementations and recommended task-specific hyperparameters, and controlled experiments share the same measurement noise realizations. The selected solvers, regularization weights, inner iteration budgets, and annealing schedules are listed in Appendix F. All DiffOps experiments run on an NVIDIA A10 GPU.

Evaluation. We evaluate FFHQ restorations using PSNR, SSIM, LPIPS, and FID. We measure phase-retrieval success as the fraction of images with PSNR exceeding 20 dB after resolving the 180° ambiguity, from one run per image (seed 0). Scientific benchmarks combine standard image-quality metrics with domain-specific measurement fidelity, and complete metric specifications are provided in Appendices C and E. Task-specific measures address acquisition ambiguities and resolution, separating image resemblance from agreement with measurements. NFE counts diffusion-network evaluations and excludes inner operator evaluations. Reported runtimes use each method's native hardware and are not a common-hardware timing comparison. The vertical axis in each panel of

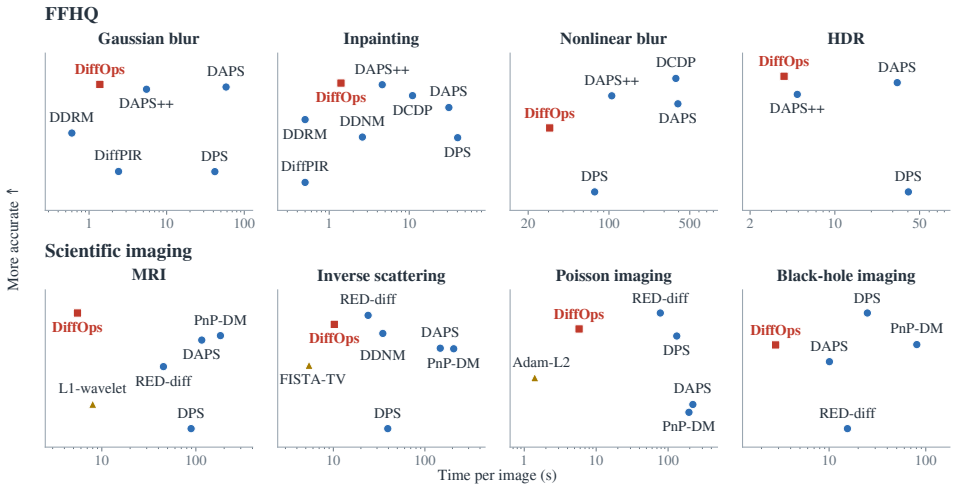

Figure 3: Reconstruction quality and runtime across eight inverse problems.

Fig. 3 averages that task's quality metrics over the test images and subtracts their standard deviation, so higher positions indicate stronger and more consistent quality.

Overall results. DiffOps achieves strong reconstruction quality across various inverse problems by choosing the applicable solver and achieves reasonable perceptual metrics or measurement fidelity with fewer diffusion functions than other benchmarks. On FFHQ, it uses few diffusion-network evaluations and gives the favorable quality–runtime trade-off shown in Fig. 3. The pretrained prior supplies image statistics that the measurements lack, while the numerical solver enforces the observation constraints at low cost. This benefit is more pronounced in scientific imaging, where each problem has its own forward and noise model. Across these problems, DiffOps often achieves a favorable balance of image quality, measurement fidelity, and computation.

Operator-aware optimization on synthetic operators. Here we compare DiffOps results on synthetic operators using FFHQ. Tab. 1 shows that DiffOps achieves better perceptual quality on Gaussian blur and HDR, usually with fewer NFEs than other benchmarks. For each problem, the pretrained FFHQ prior stays fixed under the same annealing schedule, while the inner solver changes to adapt to the different problem. For linear problems, solvers that converge to the same local optimum produce similar results, as shown by the solver ablations in Appendix D. At the same time, the results show that high perceptual quality does not require a dense diffusion sampling schedule, although more sampling steps can still improve detail for some samplers, as illustrated in Appendix D.

Table 1: Quantitative comparison across FFHQ inverse problems, reporting reconstruction quality and diffusion-network cost. **Bold** means best; underlined means second best.

	Gauss. blur		4× SR		box inpaint.		motion blur		HDR		nonlin. blur		
method	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	NFE
DAPS++	.170	50.9	**.170**	**43.8**	**.130**	**40.4**	**.135**	38.5	.171	44.4	.180	48.8	100
DAPS	.179	52.6	.192	55.3	.176	49.7	.137	**38.1**	.163	40.8	.191	57.5	1000
DPS	.285	73.2	.357	80.9	.259	75.4	.327	76.8	.284	77.4	.281	74.2	1000
DDRM	.218	78.2	.191	66.3	.210	79.2	—	—	—	—	—	—	20
DDNM	.216	57.8	.180	64.5	.146	50.5	—	—	—	—	—	—	100
DiffPIR	.293	62.7	.269	55.0	.344	65.6	.257	50.0	—	—	—	—	100
DCDP	.169	54.4	.335	92.0	.202	48.7	.187	59.3	—	—	**.160**	**44.8**	—
ours	**.158**	**45.6**	.175	47.5	.140	42.0	.184	50.7	**.156**	**35.9**	.209	54.4	50
data step	*spectral*		*gradient*		*ensemble*		*proximal*		*ensemble*		*GN×2*		

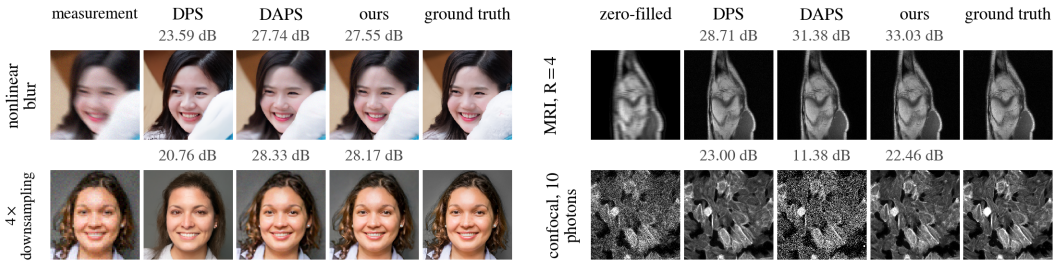

Figure 4: Qualitative reconstruction across four measurement problems, with DiffOps using numerical corrections suited to their structure.

Nonlinear forward models. Comparisons on nonlinear forward models demonstrate the advantage of adapting the solver to each problem, even at different diffusion rendering stages, particularly when the optimization landscape is nonconvex. As shown in Fig. 4, DiffOps gives visually comparable results on nonlinear deblurring with a Gauss–Newton solver while requiring fewer NFE.

For Fourier phase retrieval, the main difficulty is global ambiguity. When gradient-based corrections become trapped in a local optimum, subsequent diffusion regularization provides limited guidance and may lead to reconstruction failure. We therefore use a mixed strategy that incorporates Douglas–Rachford (DR) splitting to enforce magnitude correction under the diffusion prior. However, DR alone can lose important relative-phase constraints, producing solutions that miss details present in the reference image, as shown in Fig. 5. DiffOps combines the complementary strengths of both methods: DR provides broader correction at high noise levels, while gradient-based updates provide local guidance and recover fine details at lower noise levels. Under the same NFE budget in Tab. 2, our method yields the fewest failed reconstructions and the strongest reconstruction quality among all compared methods. With the diffusion budget fixed, this ablation suggests that staging complementary corrections matters more than relying on either solver alone. Additional phase-retrieval ablations are provided in Appendices C and D.

Table 2: Quantitative comparison on phase retrieval with benchmark methods and different inner solvers. **Bold** indicates the best within each NFE budget.

NFE	Method	Success ↑	PSNR ↑	LPIPS ↓	s/img ↓
1000	DPS	13%	12.10	0.641	69.6
	DAPS	68%	24.24	0.301	46.8
	DAPS++	79%	26.70	0.248	72.3
	Ours (DR → GD)	**94%**	**29.38**	**0.166**	**38.9**
100	Ours (DR → GD)	**74%**	**24.66**	**0.289**	3.92
	Ours (DR)	68%	20.85	0.402	**3.03**
	Ours (GD)	52%	21.40	0.377	5.71
	Ours (GN)	61%	22.18	0.363	64.6

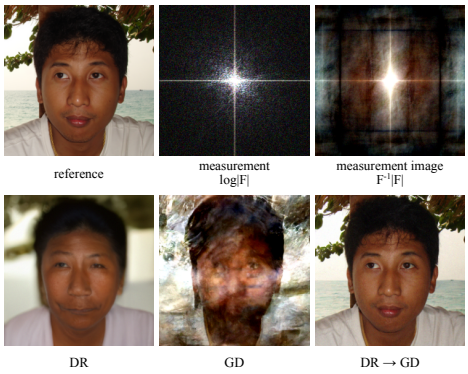

Figure 5: Qualitative reconstruction results on the phase retrieval problem with different inner solvers.

Operator-aware optimization on scientific problems. The scientific benchmarks include common imaging fidelity measures and verify whether the operator-aware correction carries over to image restoration when the forward physics and noise differ substantially. Tab. 3 shows that the benefits follow the structure of each imaging system.

In multi-coil MRI, spatially varying coil sensitivities rule out a direct Fourier-domain solve, but matrix-free Krylov iterations enforce consistency with the subsampled k-space data at low cost. DiffOps therefore fits this setting, applies a few conjugate-gradient (CG) updates to enforce consistency with the scanner data, and closely matches the acquired measurements. Appendix E gives the full comparison and the sensitivity to β under simulated and raw measurements. As noted, scanner

k-space measurements have an estimated noise level of $\sigma_y \approx 0.04$, but we adjust β accordingly to enable comparison with other benchmarks.

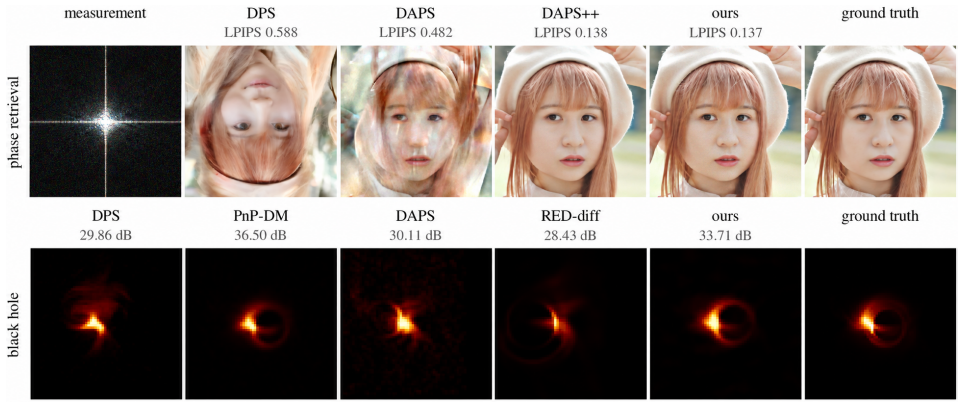

Figure 6: Reconstructions for two nonconvex inverse problems: Fourier phase retrieval and black-hole imaging from interferometric closure measurements.

Photon-limited fluorescence microscopy shows the importance of matching the statistical observation model. With an expectation-maximization (EM) step that respects Poisson statistics and nonnegativity, DiffOps lets the diffusion prior recover cellular structure from low photon counts. The confocal comparison shows that preserving spatial organization and minimizing pixelwise error can favor different methods under the same sparse measurements. The FFHQ Poisson experiments further suggest that likelihood-matched updates preserve structural detail and perceptual quality. In black-hole

Table 3: Quantitative comparison on scientific imaging tasks, reporting reconstruction quality and measurement consistency. **Bold**: best; underlined: second best.

problem	metric	DPS	DAPS	PnP-DM	RED-diff	**ours**
MRI, R=4 raw	PSNR/SSIM/misfit	25.69/.549/35.08	**28.77**/**.691**/32.12	27.88/.680/32.27	28.74/.626/31.59	28.02/.661/**31.58**
black hole, 10%	aligned/blur/χ^2_{cp}	**25.40**/**32.69**/9.6	23.23/29.24/2.1	24.55/30.88/1.5	22.89/28.05/2.2	23.90/30.10/**1.3**
microscopy (confocal), 10 ph.	PSNR/SSIM	22.44/.673	13.10/.295	14.43/.203	**23.50**/.661	22.25/**.703**
FFHQ Poisson, 100 ph.	PSNR/SSIM/LPIPS	25.30/.681/.242	19.50/.591/.405	24.70/.728/.299	**27.98**/**.800**/.213	26.40/.736/**.202**
inverse scattering	PSNR/SSIM/rel. L_2	32.01/.850/.076	34.88/.955/.054	34.62/.969/.056	**36.57**/.981/**.045**	35.81/**.989**/.049

imaging with very-long-baseline interferometry, DiffOps achieves the closest physical consistency, bringing the χ^2 error closest to 1. We also show that our method reproduces favorable visual results with a nonlinear solver. The quasi-Newton (L-BFGS) solver fits the nonconvex closure phases and amplitudes while keeping the solution anchored to the learned prior. Inverse scattering likewise favors different methods depending on whether spatial structure or aggregate field error matters more, as shown in Appendix E.4. Together, these tasks show how DiffOps pairs pretrained priors with operator-aware solvers and adjusts the effort spent on measurement enforcement and diffusion sampling to each problem.

5 Conclusion

We introduce DiffOps, a framework that integrates pretrained diffusion priors with classical, operator-aware optimization. Its numerical correction adapts to the measurement operator and noise model, while the diffusion prior supplies information where measurements alone are insufficient. Across image restoration and scientific imaging tasks, DiffOps improves reconstruction quality, recovery reliability, or measurement consistency on most tasks while often requiring fewer diffusion evaluations. Overall, our method extends diffusion-based inverse problem solving by adapting the numerical correction strategy to the requirements of each specific problem. This flexibility may broaden the use of pretrained diffusion priors across a wider range of inverse problems.

AI USE STATEMENT

Generative AI tools assisted with drafting and copy-editing the manuscript, searching the literature and checking the bibliography, and developing analysis and plotting code. The reported measurements come from the runs and evaluators described in the paper, not from generative AI. The authors verified the derivations, claims, citations, code outputs, and results and take responsibility for the final manuscript.

ETHICS STATEMENT

This work uses public datasets and previously released observations under their original terms. It collects no new human-subject data. A reconstruction may appear detailed even where measurements provide little information, because the learned prior can supply that detail. This matters when interpreting clinical or scientific images. We distinguish measurement-supported information from prior-supplied content and report data-fit quantities alongside image-quality metrics where available.

REPRODUCIBILITY STATEMENT

Appendix A describes the full outer loop and solver catalogue. Appendix B provides the nullspace-preservation analysis, local nonlinear analyses, and synthetic illustrations supporting Section 3. Appendices C and E specify the data, noise, and metric conventions, and Appendix F lists the per-problem configurations. The supplementary material includes the code.

REFERENCES

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. doi: 10.1561/2200000016. URL https://doi.org/10.1561/2200000016.

Andrew A. Chael, Michael D. Johnson, Katherine L. Bouman, Lindy L. Blackburn, Kazunori Akiyama, and Ramesh Narayan. Interferometric imaging directly with closure phases and closure amplitudes. *The Astrophysical Journal*, 857(1):23, 2018. doi: 10.3847/1538-4357/aab6a8. URL https://doi.org/10.3847/1538-4357/aab6a8.

Hao Chen, Renzheng Zhang, and Scott S Howard. Daps++: Rethinking diffusion inverse problems with decoupled posterior annealing. In *European Conference on Computer Vision*, pp. 223–240. Springer, 2026. doi: 10.1007/978-3-032-37023-5_13. URL https://doi.org/10.1007/978-3-032-37023-5_13.

Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023a. doi: 10.48550/arXiv.2209.14687. URL https://arxiv.org/abs/2209.14687.

Hyungjin Chung, Dohoon Ryu, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye. Solving 3d inverse problems using pre-trained 2d diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023b. doi: 10.1109/CVPR52729.2023.02159. URL https://doi.org/10.1109/CVPR52729.2023.02159.

Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G. Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024. doi: 10.48550/arXiv.2410.00083. URL https://arxiv.org/abs/2410.00083.

Heinz W. Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, 1996. doi: 10.1007/978-94-009-1740-8. URL https://doi.org/10.1007/978-94-009-1740-8.

Event Horizon Telescope Collaboration. First M87 Event Horizon Telescope results. IV. imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019. doi: 10.3847/2041-8213/ab0e85. URL https://doi.org/10.3847/2041-8213/ab0e85.

Berthy T. Feng, Jamie Smith, Michael Rubinstein, Huiwen Chang, Katherine L. Bouman, and William T. Freeman. Score-based diffusion models as principled priors for inverse imaging. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. doi: 10.1109/ICCV51070.2023.00965. URL https://doi.org/10.1109/ICCV51070.2023.00965.

Per Christian Hansen. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. Monographs on Mathematical Modeling and Computation. SIAM, 1998. doi: 10.1137/1.9780898719697. URL https://doi.org/10.1137/1.9780898719697.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020. doi: 10.48550/arXiv.2006.11239. URL https://arxiv.org/abs/2006.11239.

Marco A. Iglesias, Kody J. H. Law, and Andrew M. Stuart. Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013. doi: 10.1088/0266-5611/29/4/045001. URL https://doi.org/10.1088/0266-5611/29/4/045001.

Ulugbek S. Kamilov, Charles A. Bouman, Gregery T. Buzzard, and Brendt Wohlberg. Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 40(1):85–97, 2023. doi: 10.1109/MSP.2022.3199595. URL https://doi.org/10.1109/MSP.2022.3199595.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. doi: 10.1109/CVPR.2019.00453. URL https://doi.org/10.1109/CVPR.2019.00453. arXiv:1812.04948; FFHQ dataset.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, volume 35, 2022. doi: 10.52202/068431-1926. URL https://doi.org/10.52202/068431-1926.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022. doi: 10.52202/068431-1714. URL https://doi.org/10.52202/068431-1714.

Jeffrey Larson, Matt Menickelly, and Stefan M. Wild. Derivative-free optimization methods. *Acta Numerica*, 28:287–404, 2019. doi: 10.1017/S0962492919000060. URL https://doi.org/10.1017/S0962492919000060.

Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944. doi: 10.1090/qam/10666. URL https://doi.org/10.1090/qam/10666.

Xiang Li, Soo Min Kwon, Shijun Liang, Ismail R. Alkhouri, Saiprasad Ravishankar, and Qing Qu. Decoupled data consistency with diffusion purification for image restoration. *arXiv preprint arXiv:2403.06054*, 2024. doi: 10.48550/arXiv.2403.06054. URL https://arxiv.org/abs/2403.06054.

Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989. doi: 10.1007/BF01589116. URL https://doi.org/10.1007/BF01589116.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 22(4):730–751, 2025. doi: 10.1007/s11633-025-1562-4. URL https://doi.org/10.1007/s11633-025-1562-4.

Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. In *International Conference on Learning Representations*, 2024. doi: 10.48550/arXiv.2305.04391. URL https://arxiv.org/abs/2305.04391.

Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963. doi: 10.1137/0111030. URL https://doi.org/10.1137/0111030.

Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, 2nd edition, 2006. doi: 10.1007/978-0-387-40065-5. URL https://doi.org/10.1007/978-0-387-40065-5.

John R. Rice. The algorithm selection problem. *Advances in Computers*, 15:65–118, 1976. doi: 10.1016/S0065-2458(08)60520-3. URL https://doi.org/10.1016/S0065-2458(08)60520-3.

William Hadley Richardson. Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62(1):55–59, 1972. doi: 10.1364/JOSA.62.000055. URL https://doi.org/10.1364/JOSA.62.000055.

Lawrence A. Shepp and Yehuda Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, 1982. doi: 10.1109/TMI.1982.4307558. URL https://doi.org/10.1109/TMI.1982.4307558.

Jonathan Richard Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. Technical report, Carnegie Mellon University, August 1994. URL https://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf.

Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=9_gsMA8MRKQ.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32, 2019. doi: 10.48550/arXiv.1907.05600. URL https://arxiv.org/abs/1907.05600.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. doi: 10.48550/arXiv.2011.13456. URL https://arxiv.org/abs/2011.13456.

Phong Tran, Anh Tuan Tran, Quynh Phung, and Minh Hoai. Explore image deblurring via encoded blur kernel space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11956–11965, 2021. doi: 10.1109/CVPR46437.2021.01178. URL https://doi.org/10.1109/CVPR46437.2021.01178.

Singanallur V. Venkatakrishnan, Charles A. Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *IEEE Global Conference on Signal and Information Processing*, pp. 945–948, 2013. doi: 10.1109/GlobalSIP.2013.6737048. URL https://doi.org/10.1109/GlobalSIP.2013.6737048.

Austin Wang, Hongkai Zheng, Zihui Wu, Ricardo Baptista, Daniel Z. Huang, and Yisong Yue. Ensemble Kalman sampling and diffusion prior in tandem: A split Gibbs framework. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=3DfCxd0yxo.

Yinhui Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations*, 2023. doi: 10.48550/arXiv.2212.00490. URL https://arxiv.org/abs/2212.00490.

Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. MIT Press, 1949. doi: 10.7551/mitpress/2946.001.0001. URL https://doi.org/10.7551/mitpress/2946.001.0001.

Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine L. Bouman. Principled probabilistic imaging using diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems*, 2024a. doi: 10.52202/079017-3759. URL https://doi.org/10.52202/079017-3759.

Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine L Bouman. Principled probabilistic imaging using diffusion models as plug-and-play priors. *Advances in Neural Information Processing Systems*, 37:118389–118427, 2024b. doi: 10.52202/079017-3759. URL https://doi.org/10.52202/079017-3759.

Tongda Xu, Xiyan Cai, Xinjie Zhang, Xingtong Ge, Dailan He, Ming Sun, Jingjing Liu, Ya-Qin Zhang, Jian Li, and Yan Wang. Rethinking diffusion posterior sampling: From conditional score estimator to maximizing a posterior. In *International Conference on Learning Representations*, volume 2025, pp. 100101–100124, 2025. doi: 10.48550/arXiv.2501.18913. URL https://openreview.net/forum?id=GcvLoqOoXL.

Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastMRI: An open dataset and benchmarks for accelerated MRI. *arXiv preprint arXiv:1811.08839*, 2018. doi: 10.48550/arXiv.1811.08839. URL https://arxiv.org/abs/1811.08839.

Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20895–20905, 2025. doi: 10.1109/CVPR52734.2025.01946. URL https://doi.org/10.1109/CVPR52734.2025.01946.

Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6360–6376, 2022. doi: 10.1109/TPAMI.2021.3088914. URL https://doi.org/10.1109/TPAMI.2021.3088914.

Hongkai Zheng, Wenda Chu, Austin Wang, Nikola Kovachki, Ricardo Baptista, and Yisong Yue. Ensemble kalman diffusion guidance: A derivative-free method for inverse problems. *Transactions on Machine Learning Research*, 2025a. doi: 10.48550/arXiv.2409.20175. URL https://arxiv.org/abs/2409.20175.

Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy T. Feng, Caifeng Zou, Yu Sun, Nikola Kovachki, Zachary E. Ross, Katherine L. Bouman, and Yisong Yue. Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. In *International Conference on Learning Representations*, 2025b. doi: 10.48550/arXiv.2503.11043. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/e297068f5ec7f483a017948715b591fb-Abstract-Conference.html.

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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1219–1229, 2023. doi: 10.1109/CVPRW59228.2023.00129. URL https://doi.org/10.1109/CVPRW59228.2023.00129.