PROBABILITY FLOW ODE-DRIVEN DIFFUSION POSTERIOR SAMPLING FOR LARGE SCALE INVERSE PROBLEM WITH TWO-STAGE MEASUREMENT ALIGNMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion-based posterior sampling combines learned diffusion priors with measurement likelihoods to solve inverse problems. However, existing methods typically evaluate measurement guidance using a predicted clean sample or posterior mean rather than the current diffusion state, which can introduce biased likelihood gradients under noisy observations and nonlinear forward operators. They also often require hundreds or thousands of reverse-time steps, resulting in expensive likelihood evaluations. We propose probability flow ODE posterior sampling (PF-ODE), a deterministic framework that applies measurement likelihood guidance at the current diffusion state directly into the sampler, eliminating the reliance on posterior mean estimate. This high-order probability flow ODE solver through Strang splitting works in guided sampling while reduces discretization error and likelihood evaluations. For challenging nonlinear inverse problems, we further introduce a two-stage guidance strategy: phase-based guidance first mitigates cycle skipping and establishes a robust initialization, followed by proximal optimization using full measurements to enforce data consistency and recover fine-scale structures. Experiments on noisy image inpainting, sparse-view CT, and seismic full-waveform inversion demonstrate improved reconstruction quality with fewer likelihood evaluations than existing diffusion posterior sampling methods.

1 INTRODUCTION

Inverse problems, defined as recovering an unknown signal from noisy, incomplete, or indirect measurements, lie at the heart of computational imaging, spanning tasks as diverse as low-dose computed tomography (CT), image super-resolution, and seismic full-waveform inversion (FWI). Formally, one observes $y = \mathcal{A}(x) + n$, where the forward operator $\mathcal{A}$ may be linear, as in a projection or masking operator, or strongly nonlinear, as in a wave-equation solver, and the goal is to reconstruct x, which could be medical or geophysical images. Because such problems are typically ill-posed, their solution depends critically on prior knowledge of the underlying signal. Classical reconstruction methods encode this prior through hand-crafted regularizers, such as penalized weighted least squares for low-dose CT (Wang et al., 2006), whereas supervised deep networks instead learn a direct measurement-to-image mapping (Jin et al., 2017; Yang et al., 2023). While effective, the latter requires large paired datasets and generalize poorly when the forward operator, sampling pattern, or noise level departs from the training distribution.

Diffusion and score-based generative models (Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021a) have recently emerged as promising methods for learning expressive priors that circumvent this limitation. By learning the score of the data distribution through denoising score matching, they decouple the generative prior from the measurement process, so that a single pretrained model can be reused across many tasks in a zero-shot fashion (Yang et al., 2023; Daras et al., 2024). This property has given rise to *diffusion posterior sampling*, which draws samples from the posterior $p(x|y) \propto p(x) \cdot p(y|x)$ by augmenting the reverse-time diffusion with a measurement-likelihood term. Building on early score-based reconstruction for medical imaging (Song et al., 2022), methods such as DPS (Chung et al., 2023a), PIGDM (Song et al., 2023), bi-level guided diffusion (Askari et al., 2024), and decoupled data-consistency schemes (Li et al., 2025) have achieved strong results across a broad range of imaging inverse problems (Chung et al., 2025).

Despite this progress, two obstacles continue to limit the fidelity and practicality of existing approaches. First, the guidance signal is misaligned with the diffusion state. The likelihood $p(y|x_t)$ at a noisy intermediate state x_t is intractable, so most methods approximate it by first computing

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.

the posterior mean estimate $\hat{x}_0 = \mathbb{E}[x_0|x_t]$ via Tweedie's formula and then evaluating the forward operator on this predicted clean sample, $\log p(y|x_t) \approx \log p(y|\hat{x}_0)$. Early in the reverse process, however, $\hat{x}_0$ is a coarse and unreliable estimate; when the observation is noisy and $\mathcal{A}$ is nonlinear, the resulting likelihood gradient becomes biased and the reconstruction loses fine structural detail. Second, sampling is computationally expensive. Stochastic reverse-time samplers typically require hundreds to thousands of steps, and every step demands at least one forward—and often one adjoint-evaluation of $\mathcal{A}$. When $\mathcal{A}$ is a large-scale PDE solver, as in sparse-view CT (Shen et al., 2025) or FWI (Taufik & Alkhalifah, 2025; Xie et al., 2025), these repeated evaluations dominate the total cost and make high-resolution reconstruction impractical. Deterministic probability-flow ODE (PF-ODE) samplers (Song et al., 2021b; Zhang & Chen, 2022; Jiang et al., 2024; Jayaram et al., 2025) can shorten the trajectory, but injecting measurement guidance into the ODE while taking large time steps easily destabilizes integration.

We address both issues with two strategies: Firstly, we introduce a two-stage measurement-guidance strategy tailored to more complex nonlinear operators. The first stage relies on measurement phase information to produce a robust posterior sample that mitigates cycle skipping and supplies an accurate initialization (Thaker et al., 2025); the second stage solves a proximal optimization problem using the full measurement to reinforce data consistency and restore fine detail (Wu et al., 2024). Also we utilized a Stang-splitting based probability flow ODE posterior sampling, a deterministic framework for large-scale inverse problems working efficiently in guided sampling. Rather than routing guidance through the posterior mean estimate, the proposed PF-ODE applies measurement-aligned likelihood guidance directly at the current diffusion state and folds it into the PF-ODE dynamics (Zhang et al., 2025b), removing a principal source of gradient bias.

We evaluate our method on noisy image inpainting, sparse-view CT reconstruction, and seismic full-waveform inversion. Across all three tasks it consistently improves reconstruction quality while requiring fewer likelihood evaluations than existing diffusion posterior sampling methods. Our main contributions are summarized as follows:

- **State-aligned ODE guidance.** We formulate a deterministic posterior sampler that incorporates measurement-aligned likelihood guidance at the *current* diffusion state directly within the PF-ODE, eliminating the bias introduced by posterior-mean-based guidance under noisy and nonlinear forward operators.
- **Two-stage measurement guidance.** We propose a phase-then-proximal guidance strategy that first exploits phase-only information to avoid cycle skipping and then enforces full data consistency through proximal optimization, jointly improving robustness and fine-detail recovery.
- **Broad empirical validation.** Our method demonstrates consistent gains in both reconstruction quality and computational efficiency across linear and nonlinear inverse problems spanning natural images, medical CT, and geophysical FWI.

2 Related Works

2.1 Diffusion Priors for Imaging Inverse Problems

Diffusion models (Ho et al., 2020; Song et al., 2021a) have emerged as powerful tools for ill-posed imaging inverse problems, such as inpainting and super-resolution. Recent methodologies use a single *unconditional* pretrained prior and inject measurement information only at sampling time, avoiding retraining for new tasks. RePaint (Lugmayr et al., 2022) enforces the known region by repeated resampling for inpainting. For *linear* problems, DDRM (Kawar et al., 2022) and DDNM (Wang et al., 2022) exploit an SVD / range–null-space decomposition of the forward operator to guarantee data consistency in closed form. To handle general operators that might be nonlinear and noisy, DPS (Chung et al., 2023a) approximates the intractable posterior score via a Gaussian likelihood centered at the Tweedie denoised estimate, ΠGDM (Song et al., 2023) uses a pseudoinverse-guided approximation, and MCG (Chung et al., 2022) augments the guidance with a manifold-constrained correction to keep iterates on the data manifold. Together these results establish the guidance-based sampling framework on which our method builds.

2.2 Diffusion for Sparse-View CT and Seismic Inversion

Beyond natural images, diffusion priors have been extended to scientific and medical inverse problems, whose forward operators are governed by physics principles. Song et al. (2022) solve sparse-view and limited-angle CT reconstruction by alternating unconditional score-based updates with data-consistency projections through the Radon transform. DiffusionMBIR (Chung et al., 2023b)

couples a pretrained 2D diffusion prior with a model-based total-variation regularizer to reconstruct 3D volumes from a 2D prior, and physics-based adjoint operators (forward and back projections) are incorporated directly into the guidance term. Seismic full waveform inversion (FWI) is even more complex: it is highly nonlinear and non-convex, and gradient-based solvers are prone to cycle skipping and convergence to spurious local minima (Virieux et al., 2009). Learned priors have therefore been used to regularize the velocity-model space, initially through GAN (Gayon-Lombardo et al., 2003) and CNN parameterizations, and more recently through diffusion priors that constrain the solution to geological structures while offering a route to uncertainty quantification.

2.3 PF-ODE Sampling and Measurement Alignment

The probability-flow ODE (PF-ODE) (Song et al., 2021b) defines a deterministic sampler that shares its marginal distributions with the reverse-time SDE, and admits fast high-order integrators such as DDIM (Song et al., 2021a), DPM-Solver++ (Lu et al., 2025), and the EDM samplers (Karras et al., 2022). Applying measurement guidance along a deterministic ODE trajectory, however, is subtle: a likelihood gradient derived for the stochastic path is not generally consistent with the ODE marginals, so directly transplanting SDE-style guidance can introduce a mismatch between the sampled trajectory and the data constraint. This has motivated a family of *measurement-alignment* or hard data-consistency schemes that alternate prior and likelihood updates rather than summing their gradients. DDS (Chung et al., 2024) alternates ODE prior steps with conjugate-gradient data-consistency solves; ODDPS (Shen et al., 2025) decouples the likelihood and prior updates on the ODE using physics-based adjoint operators; and DAPS (Zhang et al., 2025a) decouples consecutive states via an Langevin correction to improve stability under strong measurement guidance. Our work follows this measurement-alignment perspective, but differs in that it exploits low frequeuncy measurement phase information in the first stage, and solves a proximal optimization problem using full data in the second stage.

3 Background

3.1 Diffusion Models

Denoising diffusion probabilistic models (DDPMs) define a fixed forward noising process and a learned reverse process (Ho et al., 2020). Given $\mathbf{x}_0 \sim p(\mathbf{x}_0)$, the forward transition is

$$q(\mathbf{x}_t \mid \mathbf{x}_{t-1}) = \mathcal{N}\big(\mathbf{x}_t;\ \sqrt{1-\beta_t}\,\mathbf{x}_{t-1},\ \beta_t\mathbf{I}\big), \tag{1}$$

with variance schedule $\{\beta_t\}_{t=1}^T$. Defining $\alpha_t = 1-\beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, its marginal is

$$q(\mathbf{x}_t \mid \mathbf{x}_0) = \mathcal{N}\big(\mathbf{x}_t;\ \sqrt{\bar{\alpha}_t}\,\mathbf{x}_0,\ (1-\bar{\alpha}_t)\mathbf{I}\big), \tag{2}$$

or $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. As $t \to T$, $\mathbf{x}_T$ approaches Gaussian noise.

The learned reverse chain is

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^{T} p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t), \tag{3}$$

$$p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t) = \mathcal{N}\big(\mathbf{x}_{t-1};\ \boldsymbol{\mu}_\theta(\mathbf{x}_t, t),\ \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)\big). \tag{4}$$

DDPMs commonly predict the injected noise with $\boldsymbol{\epsilon}_\theta$ using

$$\mathcal{L}(\theta) = \mathbb{E}_{t,\mathbf{x}_0,\boldsymbol{\epsilon}}\Big[\big\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}, t)\big\|_2^2\Big]. \tag{5}$$

3.2 Sampling

Starting from $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, ancestral DDPM sampling iterates

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\Big(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\,\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)\Big) + \sigma_t\mathbf{z}, \qquad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{6}$$

where σ_t controls stochasticity. Equivalently, the score $s_\theta(\mathbf{x}_t, t) \approx \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ satisfies (Song et al., 2021b)

$$s_\theta(\mathbf{x}_t, t) \approx -\frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)}{\sqrt{1-\bar{\alpha}_t}}, \tag{7}$$

which yields the Langevin-type form

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\big(\mathbf{x}_t + \beta_t\,\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)\big) + \sigma_t\mathbf{z}. \tag{8}$$

Tweedie's formula provides the clean posterior mean (Efron, 2011),

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \frac{1}{\sqrt{\bar{\alpha}_t}}\big(\mathbf{x}_t + (1-\bar{\alpha}_t)s_\theta(\mathbf{x}_t, t)\big). \tag{9}$$

3.3 Diffusion Posterior Sampling for Inverse Problems

We consider measurements

$$\mathbf{y} = A(\mathbf{x}_0) + \mathbf{w}, \qquad \mathbf{x}_0 \sim p(\mathbf{x}_0), \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}), \tag{10}$$

with possibly nonlinear A. Posterior sampling replaces the prior score by

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y}) = \underbrace{\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)}_{\text{prior score}} + \underbrace{\nabla_{\mathbf{x}_t} \log p(\mathbf{y} \mid \mathbf{x}_t)}_{\text{likelihood correction}}. \tag{11}$$

The likelihood term is generally intractable because it marginalizes over $p(\mathbf{x}_0 \mid \mathbf{x}_t)$. DPS (Chung et al., 2023a) approximates it using the Tweedie estimate, $p(\mathbf{y} \mid \mathbf{x}_t) \approx p(\mathbf{y} \mid \hat{\mathbf{x}}_0(\mathbf{x}_t))$. Under Gaussian measurement noise,

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} \mid \mathbf{x}_t) \approx -\frac{1}{2\sigma_y^2} \nabla_{\mathbf{x}_t} \big\| \mathbf{y} - A\big(\hat{\mathbf{x}}_0(\mathbf{x}_t)\big) \big\|_2^2, \tag{12}$$

requiring gradients through the denoiser and making the guidance sensitive to errors in $\hat{\mathbf{x}}_0$.

4 Methodology

Recovering an oscillatory signal from its measurement naturally separates into two sub problems due to their very different characters. The first is *kinematic*: when the noise level is high, events in the predicted and observed data are misaligned in time, and this results in a coarse and highly nonlinear discrepancy that squared-ℓ_2 fidelity cannot resolve without cycle-skipping. The second is *dynamic*: once events line up in time with less noise, what remains is a mismatch in amplitude and waveform shape. We therefore synchronize measurement guidance with the diffusion noise level: phase-robust guidance corrects kinematics at high noise, followed by a data-proximal solve that enforces full measurement consistency at low noise. We first define the proximal update and then the two-stage schedule.

4.1 Measurement Alignment with a Data-Proximal Solution

DPS (Sec. 3.3) uses a single gradient correction and differentiates through the denoiser. Instead, at each reverse step we compute a data-consistent estimate from the Tweedie mean $\hat{\mathbf{x}}_0$:

$$\tilde{\mathbf{x}}_0 = \arg\min_{\mathbf{x}} \ \big\| \mathbf{y} - \mathcal{A}(\mathbf{x}) \big\|_2^2 + \rho_t \big\| \mathbf{x} - \hat{\mathbf{x}}_0 \big\|_2^2, \tag{13}$$

where

$$\rho_t \triangleq \lambda \frac{\sigma_y^2}{\bar{\sigma}_t^2}, \qquad \bar{\sigma}_t^2 := \frac{1 - \bar{\alpha}_t}{\bar{\alpha}_t}. \tag{14}$$

Thus early noisy steps favor the prior and later steps increasingly enforce the measurements. For linear $\mathcal{A}(\mathbf{x}) = A\mathbf{x}$, Eq. 13 gives

$$\big(A^\top A + \rho_t \mathbf{I}\big)\, \tilde{\mathbf{x}}_0 = A^\top \mathbf{y} + \rho_t \hat{\mathbf{x}}_0, \tag{15}$$

solved with a few warm-started conjugate-gradient iterations; nonlinear operators use a short Gauss–Newton solve initialized at $\hat{\mathbf{x}}_0$.

The solution re-enters the sampler through

$$\hat{\boldsymbol{\epsilon}} = \frac{1}{\sqrt{1 - \bar{\alpha}_t}} \big(\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\, \tilde{\mathbf{x}}_0\big), \tag{16}$$

and the DDIM-style update

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\, \tilde{\mathbf{x}}_0 + \sqrt{1 - \bar{\alpha}_{t-1}} \left(\sqrt{1 - \gamma}\, \hat{\boldsymbol{\epsilon}} + \sqrt{\gamma}\, \boldsymbol{\epsilon}_t\right), \tag{17}$$

where $\gamma \in [0, 1]$ controls stochasticity. Unlike the DPS correction, the proximal step solves data consistency to the desired tolerance without differentiating through the denoiser.

4.2 Two-Stage Guidance: Phase Alignment via Cross-correlation and Full-Data Alignment via the Proximal Solution

For oscillatory measurements, raw squared-ℓ_2 fidelity is highly non-convex under traveltime errors and can cause cycle skipping. We therefore switch at diffusion time $t^\star$ from phase-robust alignment to full-data proximal alignment.

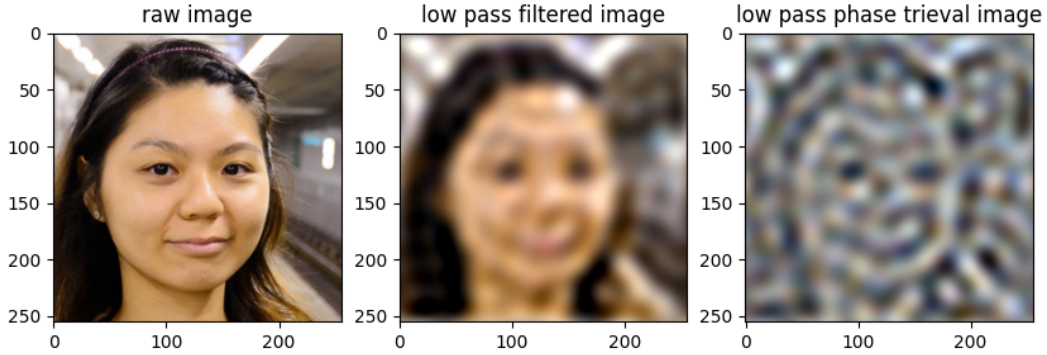


Figure 1: **Low-pass filtering and phase retrieval.** Comparison of the original image, its low-pass version, and a phase-retrieval reconstruction using Fourier phase only.

Stage 1: phase alignment ($t > t^\star$). At high noise, we replace the raw residual with a phase-robust misfit built from dynamic time warping (DTW) or only low frequency component (Figure 1). Especillay DTW estimates a monotonic bounded-strain shift field between $\mathcal{A}(\hat{\mathbf{x}}_0)$ and $\mathbf{y}$,

$$\mathbf{u}^\star = \arg\min_{\mathbf{u}\in\mathcal{U}} \sum_i \big(\mathcal{A}(\hat{\mathbf{x}}_0)[i] - \mathbf{y}[i + u_i]\big)^2, \tag{18}$$

solved by dynamic programming (Hale, 2013; Ma & Hale, 2013). Warping the observation, $\tilde{\mathbf{y}} = \mathcal{W}_{\mathbf{u}^\star}\mathbf{y}$, defines

$$f_{\text{phase}}(\mathbf{x}) = \tfrac{1}{2}\big\|\mathcal{A}(\mathbf{x}) - \tilde{\mathbf{y}}\big\|_2^2, \tag{19}$$

which largely removes traveltime mismatch. We then apply

$$\mathbf{x}_{t-1} = \mathbf{x}'_{t-1} - \zeta_t \nabla_{\mathbf{x}_t} f_{\text{phase}}\big(\hat{\mathbf{x}}_0(\mathbf{x}_t)\big), \qquad \zeta_t = \frac{\zeta}{\big\|A(\hat{\mathbf{x}}_0) - \tilde{\mathbf{y}}\big\|_2}, \tag{20}$$

holding $\mathbf{u}^\star$ fixed during differentiation. This stage primarily corrects the cycle-skip-prone kinematics.

Stage 2: full-data alignment ($t \le t^\star$). After phase alignment, we switch to Eq. 13, solve for $\tilde{\mathbf{x}}_0$, form Eq. 16, and apply Eq. 17. The now well-aligned full-waveform residual enables accurate amplitude and fine-structure recovery.

Stage transition. The switch $t^\star$ may be fixed or triggered when $\|\mathbf{u}^\star\|_\infty$ falls below a tolerance. This coarse-to-fine strategy follows the principle of robust kinematic alignment before amplitude fitting (Warner & Guasch, 2016).

4.3 Probability Flow ODE for Diffusion Models with Strang Splitting

The discrete diffusion process approaches the variance-preserving SDE

$$d\mathbf{x}_t = -\tfrac{1}{2}\beta(t)\,\mathbf{x}_t\,dt + \sqrt{\beta(t)}\,d\mathbf{w}_t, \tag{21}$$

whose deterministic marginal-preserving probability-flow ODE (PF-ODE) (Song et al., 2021c), integrated from $t = T$ to 0, is

$$d\mathbf{x}_t = \tfrac{1}{2}\beta(t)\big[\mathbf{x}_t + \nabla_{\mathbf{x}_t}\log p_t(\mathbf{x}_t)\big]dt = \tfrac{1}{2}\beta(t)\big[\mathbf{x}_t + s_\theta(\mathbf{x}_t, t)\big]dt. \tag{22}$$

Because this flow is deterministic, high-order integration permits much larger steps than ancestral sampling; Recent techniques can accelerate unguided sampling by applying high-order numerical methods to the sampling process when viewed as differential equations. However, the same techniques have been discovered that they did not work for guided sampling (Wizadwongsa & Suwajanakorn, 2023). We slightly improve Strang Splitting work and use in our experiments. We split the PF-ODE into linear drift and nonlinear score terms,

$$L(\mathbf{x}) = \tfrac{1}{2}\beta(t)\,\mathbf{x}, \qquad N(\mathbf{x}, t) = \tfrac{1}{2}\beta(t)\,s_\theta(\mathbf{x}, t), \tag{23}$$

and use the symmetric Strang composition

$$\Phi_{\Delta t} = \Phi^L_{\Delta t/2} \circ \Phi^N_{\Delta t} \circ \Phi^L_{\Delta t/2}. \tag{24}$$

This Strang operator-splitting formulation separates the diffusion dynamics from the guidance term to enable efficient numerical integration. Freezing β over a substep gives the exact linear half-flow

$$\Phi^L_{\Delta t/2}(\mathbf{x}) = \exp\big(\tfrac{1}{4}\beta\,\Delta t\big)\,\mathbf{x}, \tag{25}$$

while the nonlinear flow is integrated with the midpoint method. The resulting symmetric scheme is second-order accurate and confines network evaluations to the nonlinear substep. In our experiment, This method enables 10–50 PF-ODE steps instead of the origin 1000 SDE steps. The sampling procedure is summarized in Algorithm 1 (Appendix A).

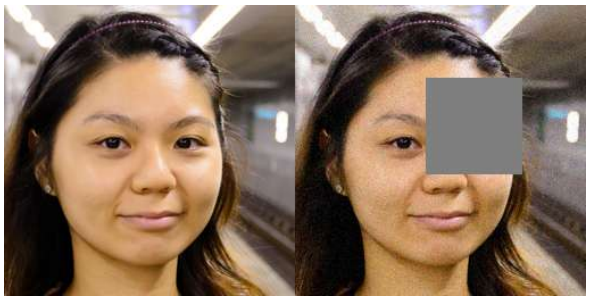


Figure 2: **the image from FFHQ dataset before (left) and after inpainting (right) which has been applied an element-wise multiplication with a mask plus random noise.** .

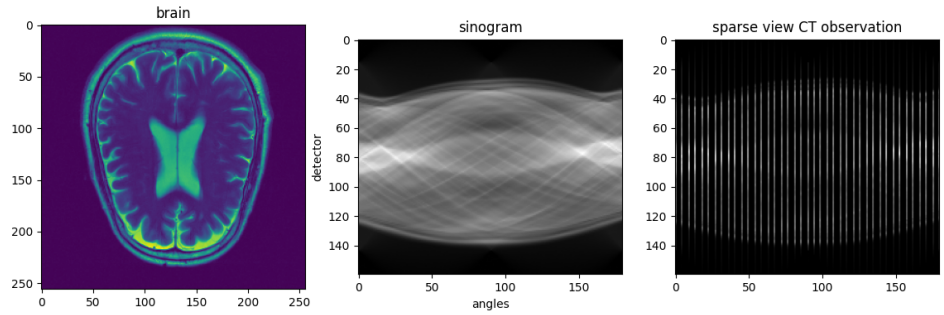


Figure 3: **Measurement process for sparse-view CT.** The original brain image is projected by the Radon transform to produce a full-view sinogram (180 views). A binary view-selection mask is then applied to generate the sparse-view measurements used for reconstruction.

5 Experiments

This section evaluates the proposed efficient PFODE with two measurement alignment on natural image restoration, sparse-view CT reconstruction, and seismic full waveform inversion (FWI). Experiments are designed to answer two questions: (1) whether the proposed Strang-splitting PF-ODE (SS-PFODE) with two-stage measurement alignment guidance improves posterior sampling compared with posterior-mean guidance; and (3) whether the proposed two-stage guidance strategy generalizes from linear imaging problems to large-scale nonlinear inverse problems.

5.1 Experimental Setup

Natural image restoration. Image inpainting (Figure 2) and posterior sampling experiments are conducted on the FFHQ dataset. Images are resized to 256×256. Random irregular masks covering 20–60% of pixels are generated following the standard free-form masking protocol. The same pre-trained unconditional diffusion model is used for all competing methods. We compare DPS-SDE, DPS-DDIM, DPS-PFODE sampling, and the proposed SS-PFODE using identical denoisers and noise schedules. Reconstruction quality is evaluated using PSNR, SSIM, and LPIPS. The results are summed up in Table 1.

Sparse-view CT reconstruction. Sparse-view CT experiments are performed on the fastMRI dataset by treating reconstructed slices as attenuation images (Figure 3). Forward projection, adjoint backprojection, and filtered backprojection (FBP) are implemented using the All Scale Tomographic Reconstruction Antwerp (ASTRA) toolbox with a parallel-beam geometry. The forward operator consists of the Radon transform, while the likelihood gradient uses the exact adjoint operator instead of numerical differentiation. FBP is used as the classical reconstruction baseline and also provides initialization for comparison. Our implementation follows the standard ASTRA projection pipeline with strip projectors, parallel geometry, and GPU acceleration when available. (Van Aarle et al., 2015) The implementation directly employs ASTRA forward projection, adjoint back projection, and FBP routines.

Seismic full-waveform inversion. Seismic full-waveform inversion (FWI) is an advanced geophysical imaging technique used to build high-resolution subsurface models, such as velocity and density, by minimizing differences between recorded seismograms and numerically simulated wavefields. Conventional FWI is vulnerable to cycle skipping when simulated and observed data are misaligned by more than roughly half a wavelength, causing convergence to local minima. Time-matching

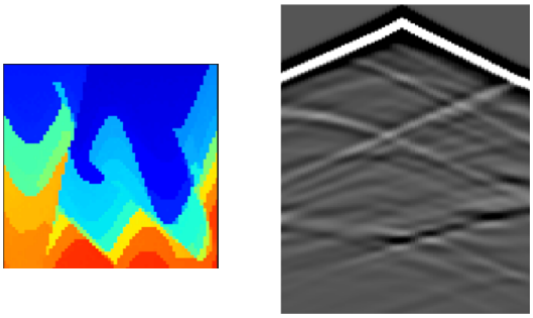


Figure 4: **Velocity model and seismic shot gather data for the left velocity model.** The right seismic shot gather data is generated by finite difference acoustic modeling based on the left true velocity model.



Figure 5: (a) DPS-SDE; (b) DPS-DDIM; (c) DPS-ΠGDM; (d) DPS-PFODE; (e) proposed method with phase-retrieval guidance; (f) proposed method with low-pass guidance; (g) ground truth. DPS-SDE uses 1000 time steps; all other methods use 50 time steps.

FWI (Wang et al., 2019) reduces this sensitivity by replacing strict sample-by-sample waveform comparison, such as the L2 norm, with nonlinear alignment methods, including Dynamic Time Warping (DTW).

For large-scale deepwater projects in the Gulf of Mexico (GoM), seismic datasets may range from hundreds of terabytes to multiple petabytes. Production FWI requires large numbers of forward and adjoint wave-propagation simulations that solve 3D acoustic or elastic wave equations over thousands of shots and many inversion iterations. Consequently, a multi-parameter production FWI workflow can require several weeks to 3–6 months for computation, quality control, and finalization.

The nonlinear FWI inverse problem we aimed at is evaluated on the OpenFWI dataset (Figure 4) which is an open-source collection of large-scale benchmark datasets designed for machine learning-driven seismic full waveform inversion(Deng et al., 2021). Acoustic wave propagation is solved using a finite-difference wave-equation simulator with sparse acquisition geometry and band-limited Ricker wavelets. During the first stage of posterior sampling, only phase-aligned measurements are used to suppress cycle skipping; the second stage switches to full-waveform proximal optimization after the estimated velocity model enters the attraction basin of the solution. We compare against classical FWI baseline, TV-regularized FWI (Inada et al., 2025), RED-DiffEq (Shan et al., 2025), DiffPIR-based FWI (Zhu et al., 2023), and naive DPS.

5.2 Experimental Results

Natural image restoration. Table 1 compares image inpainting performance on FFHQ (Figure 5). Conventional DPS using DDPM requires hundreds of likelihood evaluations to recover fine image structures. DDIM substantially reduces sampling cost but often introduces over-smoothed textures because measurement guidance is still computed from the posterior mean estimate. PF-ODE further accelerates inference; however, large integration steps accumulate numerical error and amplify guidance instability. In contrast, the proposed Strang-splitting PFODE preserves facial details while requiring only a small number of likelihood evaluations.

Table 1: FFHQ comparison using identical denoisers and noise schedules. Reconstruction quality in Figure 5 is evaluated using PSNR, SSIM, and LPIPS.

Method	PSNR↑	SSIM↑	LPIPS↓
DPS-SDE	25.4865	0.9817	0.1284
DPS-DDIM	20.9991	0.9485	0.2373
DPS-ΠGDM	25.8369	0.9809	**0.0468**
DPS-PFODE	25.7094	0.9797	0.1010
SS-PFODE (Phase Retrieval)	25.7096	0.9817	0.1096
SS-PFODE (Low-pass Filter)	**25.9250**	**0.9828**	0.1125

Table 2: Sparse-view CT comparison using identical denoisers and noise schedules. Reconstruction quality in Figure 6 is evaluated using PSNR, SSIM, and LPIPS.

Method	PSNR↑	SSIM↑	LPIPS↓
DPS-SDE	29.1236	0.9325	0.2486
DPS-DDIM	18.9090	0.5190	0.2895
DPS-ΠGDM	22.1792	0.7381	0.5950
DPS-PFODE	26.6980	0.8533	0.2253
SS-PFODE (Phase Retrieval)	**34.2491**	**0.9793**	**0.1752**
SS-PFODE (Low-pass Filter)	33.3135	0.9745	0.1963

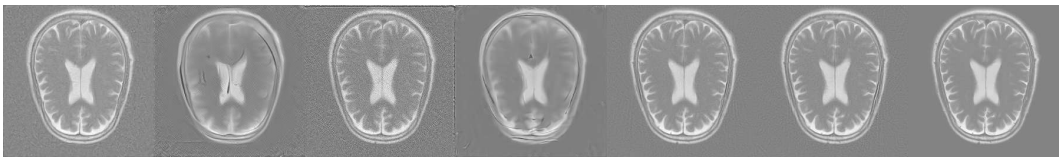


Figure 6: (a) DPS-SDE; (b) DPS-DDIM; (c) DPS-ΠGDM; (d) DPS-PFODE; (e) proposed method with phase-retrieval guidance; (f) proposed method with low-pass guidance; (g) ground truth. DPS-SDE uses 1000 time steps; all other methods use 100 time steps.

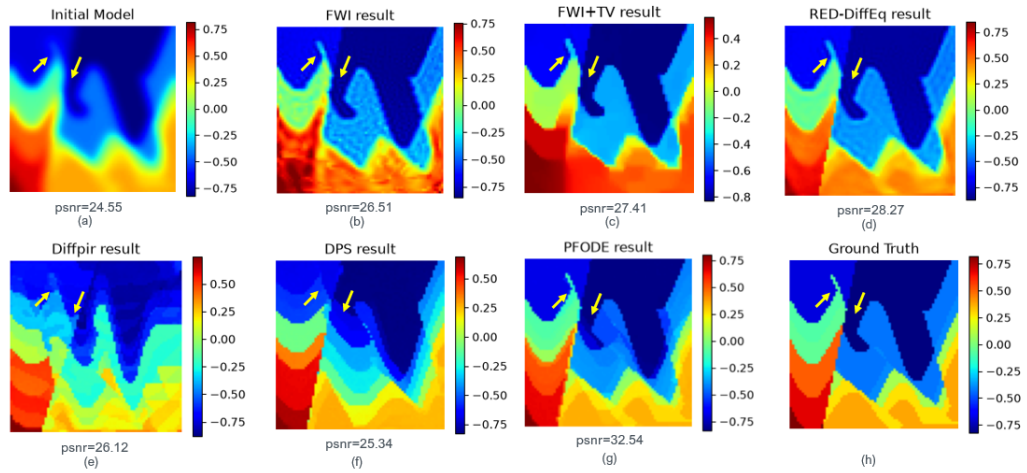


Figure 7: Comparison between initial model(a), classic FWI result (b), FWI result using Total variation regularization (c), RED-DiffEq result(d), Diffpir result(e), naive DPS result(f), proposed PFODE result (h) and Ground Truth; It also show the corresponding PSNR values for each result: 24.55 for initial model (a), 26.51 for classic FWI result (b), 27.41 for FWI+TV (c), 28.27 for RED-DiffEq (d), 26.12 for Diffpir (e), 25.34 for naive DPS (f), 32.54 for FWI with proposed method (g).

Sparse-view CT reconstruction. Figure 6 demonstrates sparse-view CT reconstruction results. Classical FBP suffers from severe streaking artifacts because of the limited number of projection views. DPS suppresses most streak artifacts but tends to smooth high-frequency anatomical structures. The proposed SS-PFODE with two stage measurement alignment produces cleaner reconstructions with better preservation of fine edges while requiring substantially fewer forward and adjoint operator evaluations. Quantitative comparisons consistently show higher PSNR and SSIM together with lower reconstruction error. The reconstructions are evaluated in Table 2.

Seismic full-waveform inversion. As shown in Figure 7, classical FWI recovers major boundaries but exhibits strong artifacts and noise, particularly at depth. TV-regularized FWI produces sharper, more block-like interfaces and suppresses some noise, but can oversimplify smoothly varying structures and introduce staircase artifacts. RED-DiffEq gives a cleaner reconstruction with improved structural continuity, while the DiffPIR-based method recovers major structures but remains comparatively smooth and less accurate at depth. Naive SDE-based DPS misplaces several central and deep interfaces. Our PF-ODE method achieves the highest PSNR in this comparison (32.54 dB).

This OpenFWI model experiments illustrate the advantage of the proposed two-stage guidance strategy for highly nonlinear inverse problems. Conventional posterior sampling frequently converges to local minima due to cycle skipping when full waveform residuals are used from the beginning of sampling. By first matching phase information and subsequently switching to full-waveform proximal optimization, the proposed SS-PFODE progressively aligns large-scale geological structures before recovering fine velocity contrasts. The reconstructed salt boundaries, sediment interfaces, and deep structures are significantly closer to the ground truth while maintaining stable optimization throughout the PF-ODE trajectory.

Overall, the proposed framework consistently achieves better reconstruction quality while reducing computational cost across linear and nonlinear inverse problems. These results demonstrate that state-aligned likelihood guidance and SS-PFODE integration complement each other, enabling efficient posterior sampling with large integration steps.



Figure 8: Qualitative comparison of the two-stage guidance ablation. From left to right: **(a) I-Full**—DPS-ODE with the full likelihood gradient as inner guidance, no outer proximal solution; **(b) I-LF**—DPS-ODE with only the low-frequency component of the gradient as inner guidance, no outer proximal solution; **(c) P-only**—outer proximal solution alone; **(d) Ours-PhaseImg**—proposed two-stage sampler with phase information; **(e) Ours-LF**—proposed two-stage sampler using the low-frequency-component; Inner guidance alone (a, b) is over-smoothed and drifts from the measurements, whereas adding the outer proximal solution (d, e) restores data-consistent high-frequency detail; Ours-LF (e) is closest to the ground truth, consistent with Table 3.

5.3 Ablation Study

We perform ablation experiments to quantify the effect of two-stage guidance.

To isolate the contributions of the two components of the proposed sampler—the *inner gradient guidance* (a DPS-style likelihood-gradient step along the PF-ODE trajectory) and the *outer proximal solution* (a data-consistency correction applied after each denoising step)—we ablate each stage independently and in combination (Table 3). Figure 8 shows the comparison of inverted image results.Two things worth flagging for the writeup: PSNR/SSIM and LPIPS disagree slightly at the top (Ours-LF wins the former, Ours-Img the latter), both proposed phase retrieval component and low frequency component gives the optimal results.

Also from CT reconstruction, the proximal refinement stage increases edge sharpness and decreases streak artifacts. For seismic inversion, the proposed method improves convergence robustness and final velocity accuracy.

Table 3: Ablation of the two-stage guidance. Inner guidance alone is insufficient; the outer proximal solution is the stronger single component; and their combination is best. Note that low-frequency guidance is worst alone (I-LF) but best combined (Ours-LF), showing the two stages are complementary.

Configuration	Inner	Prox.	PSNR↑	SSIM↑	LPIPS↓
I-Full	full grad.	✗	22.69	0.9634	0.2288
I-LF	low-freq.	✗	17.11	0.8825	0.3448
P-only	—	✓	24.24	0.9715	0.1544
Ours-PhaseImg	image grad.	✓	25.71	0.9817	**0.1496**
Ours-LF	low-freq.	✓	**25.93**	**0.9828**	0.1525

6 Conclusion

In this work, we propose the high-order probability flow ODE posterior sampling with two stage measurement alignment for large scale inverse problems. To resolve the posterior mean limitation, the PFODE framework incorporates likelihood guidance at the current diffusion state to the Probability Flow ODE sampler. An efficient exponential ODE solver designed with Strang-splitting is developed that improves stability and efficiency, allowing large time steps and therefore significantly reducing needed number of steps for denoising. More importantly, we borrowed the two stage inversion method which has been very successfully strategy for seismic waveform inverson for field observation data. In the proposed two-stage measurement guidance strategy, we start at stage one to exploit measurement phase information or low frequency band information, and at stage two we use all the full measurement information to recover finer details. Our experiments especially the metric and the inverted images quantitatively and also visually demonstrate a general improvement in reconstruction quality, and significant reduced evaluation requirements.

AI USE STATEMENT

LLMs are used to assist with grammar checking, improving writing fluency, managing computational resources and providing suggestions after the initial draft. All observations, ideas, methodologies in this paper are developed entirely by the authors. Any content generated with the help of LLMs is created under detailed author instructions and thoroughly verified by the authors before inclusion.

REFERENCES

Hossein Askari, Fred Roosta, and Hongfu Sun. Bi-level guided diffusion models for zero-shot medical imaging inverse problems. *IEEE Transactions on Medical Imaging (under review)*, 2024.

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, 2022.

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.

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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22542–22551, 2023b.

Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems. In *International conference on learning representations*, volume 2024, pp. 38922–38949, 2024.

Hyungjin Chung, Jeongsol Kim, and Jong Chul Ye. Diffusion models for inverse problems. *arXiv preprint arXiv:2508.01975*, 2025.

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.

Chengyuan Deng, Shihang Feng, Hanchen Wang, Xitong Zhang, Peng Jin, Yinan Feng, Qili Zeng, Yinpeng Chen, and Youzuo Lin. OpenFWI: Large-scale multi-structural benchmark datasets for seismic full waveform inversion. *arXiv preprint arXiv:2111.02926*, 2021.

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

A Gayon-Lombardo, L Mosser, NP Brandon, and SJ Cooper. Pores for thought: The use of generative adversarial networks for the stochastic reconstruction of 3d multi-phase electrode microstructures with periodic boundaries (2020). *arXiv preprint arXiv:2003.11632*, 2003.

Dave Hale. Dynamic warping of seismic images. *Geophysics*, 78(2):S105–S115, 2013.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.

Yudai Inada, Shingo Takemoto, and Shunsuke Ono. Efficient and accurate full-waveform inversion with total variation constraint. *arXiv preprint arXiv:2501.08210*, 2025.

Vivek Jayaram, Ira Kemelmacher-Shlizerman, Steven M. Seitz, and John Thickstun. Constrained diffusion implicit models. *International Conference on Learning Representations (under review)*, 2025.

Enze Jiang, Jishen Peng, Zheng Ma, and Xiong-Bin Yan. Ode-dps: Ode-based diffusion posterior sampling for inverse problems in partial differential equation. *arXiv preprint arXiv:2404.13496*, 2024.

Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9): 4509–4522, 2017. doi: 10.1109/TIP.2017.2713099.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *ArXiv*, abs/2201.11793, 2022. URL https://api.semanticscholar.org/CorpusID:246411364.

Xiang Li, Soo Min Kwon, Shijun Liang, Ismail Alkhouri, Saiprasad Ravishankar, and Qing Qu. Dcdp: Decoupled data consistency via diffusion purification for solving general inverse problems. *arXiv preprint arXiv:2403.06054*, 2025.

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.

Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11451–11461, 2022. URL https://api.semanticscholar.org/CorpusID:246240274.

Yong Ma and Dave Hale. Wave-equation reflection traveltime inversion with dynamic warping and full-waveform inversion. *Geophysics*, 78(6):R223–R233, 2013.

Siming Shan, Min Zhu, Youzuo Lin, and Lu Lu. RED-DiffEq: Regularization by denoising diffusion models for solving inverse PDE problems with application to full waveform inversion. *arXiv preprint arXiv:2509.21659*, 2025.

Ye Shen, Ailong Cai, Xinrui Zhang, Junru Ren, Yizhong Wang, Bing Chen, Yarui Xi, Zhizhong Zheng, Ningning Liang, Lei Li, and Bin Yan. Ode-driven deterministic diffusion posterior sampling for sparse-view ct image reconstruction. *arXiv preprint*, 2025.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a.

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

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

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*, 2021b.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations (ICLR)*, 2021c.

Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2022.

Mohammad H. Taufik and Tariq Alkhalifah. Diffusion model-based posterior sampling in full waveform inversion. *arXiv preprint*, 2025.

Darshan Thaker, Abhishek Goyal, and René Vidal. Frequency-guided posterior sampling for diffusion-based image restoration. *arXiv preprint arXiv:2411.15295*, 2025.

Wim Van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K Joost Batenburg, and Jan Sijbers. The astra toolbox: A platform for advanced algorithm development in electron tomography. *Ultramicroscopy*, 157:35–47, 2015.

Jean Virieux, Stéphane Operto, Hafedh Ben-Hadj-Ali, Romain Brossier, Vincent Etienne, Florent Sourbier, Luc Giraud, and Azzam Haidar. Seismic wave modeling for seismic imaging. *The Leading Edge*, 28(5):538–544, 2009.

Jing Wang, Tianfang Li, Hongbing Lu, and Zhengrong Liang. Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose x-ray computed tomography. *IEEE transactions on medical imaging*, 25(10):1272–1283, 2006.

Min Wang, Yi Xie, Weiquan Xu, and Stephen Kefeng Xin. Device and method for full waveform inversion. U.S. Patent No. 10,324,205, June 2019. URL https://patents.google.com/patent/US10324205B2/en. Granted June 18, 2019.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *ArXiv*, abs/2212.00490, 2022. URL https://api.semanticscholar.org/CorpusID:254125609.

Michael Warner and Lluís Guasch. Adaptive waveform inversion: Theory. *Geophysics*, 81(6): R429–R445, 2016.

Suttisak Wizadwongsa and Supasorn Suwajanakorn. Accelerating guided diffusion sampling with splitting numerical methods. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2301.11558.

Hongjie Wu, Linchao He, Mingqin Zhang, Dongdong Chen, Kunming Luo, Mengting Luo, Ji-Zhe Zhou, Hu Chen, and Jiancheng Lv. Diffusion posterior proximal sampling for image restoration. In *ACM Multimedia*, 2024.

Yuke Xie, Hervé Chauris, and Nicolas Desassis. Diffusion prior as a direct regularization term for full waveform inversion. *arXiv preprint arXiv:2506.10141*, 2025.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM computing surveys*, 56(4):1–39, 2023.

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*, pp. 20895–20905, 2025a.

Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022.

Shaorong Zhang, Rob Brekelmans, Yunshu Wu, and Greg Ver Steeg. Measurement-aligned sampling for inverse problems. *arXiv preprint arXiv:2506.11893*, 2025b.

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. *arXiv preprint arXiv:2305.08995*, 2023.

A ALGORITHM

Here shows the workflow of the proposed algorithm.

Algorithm 1 Two-Stage Phase Retrieval–Proximal Posterior Sampling with a PF-ODE Solver

Require: score net s_θ, operator A, measurement $\mathbf{y}$, noise level σ_y, schedule $\{\bar{\alpha}_t\}$, step size Δt, number of steps N, proximal weight λ, guidance scale ζ, switch time $t^\star$, scaling $s^\star = s^\star(N)$

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$; precompute $\rho_t = \lambda\sigma_y^2/\bar{\sigma}_t^2$
for $n = 1$ **to** N $\quad$ ($t: T \to 0$ along the PF-ODE, step Δt) **do**
  $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}}\big(\mathbf{x}_t + (1-\bar{\alpha}_t)\, s_\theta(\mathbf{x}_t, t)\big)$ {Tweedie mean, Eq. 9}
  if $t > t^\star$ **then**
    // Stage 1: phase-matching MSE gradient (posterior score)
    $\mathbf{u}^\star \leftarrow \mathbf{f}_{\text{phase}}(A(\hat{\mathbf{x}}_0), \mathbf{y})$; $\tilde{\mathbf{y}} \leftarrow \mathcal{W}_{\mathbf{u}^\star}\mathbf{y}$; $\zeta_t \leftarrow \zeta/\|\mathbf{u}^\star\|_2$
    $\hat{s}(\mathbf{x}, \tau) \leftarrow s_\theta(\mathbf{x}, \tau) - \zeta_t \nabla_{\mathbf{x}} \frac{1}{2}\|\mathbf{u}^\star\|_2^2$
  else
    // Stage 2: full-data proximal solution
    $\tilde{\mathbf{x}}_0 \leftarrow \arg\min_{\mathbf{x}} \|\mathbf{y} - A(\mathbf{x})\|_2^2 + \rho_t\|\mathbf{x} - \hat{\mathbf{x}}_0\|_2^2$ {CG, Eq. 15}
    $\hat{s}(\mathbf{x}, \tau) \leftarrow (\sqrt{\bar{\alpha}_t}\,\tilde{\mathbf{x}}_0 - \mathbf{x}_t)/(1-\bar{\alpha}_t)$ {data-consistent score}
  end if
  // PF-ODE Strang step (Eq. 24), scaling $s^\star$, integrating $\hat{s}$
  $\mathbf{x} \leftarrow \exp\big(\frac{1}{4}\beta_t\Delta t\big)\,\mathbf{x}_t$ {$\Phi^L_{\Delta t/2}$, Eq. 25}
  $\mathbf{x}_{\text{mid}} \leftarrow \mathbf{x} + \frac{1}{4}\beta_t\Delta t\; \hat{s}(\mathbf{x}, t)$ {midpoint predictor}
  $\mathbf{x} \leftarrow \mathbf{x} + s^\star \cdot \frac{1}{2}\beta_t\Delta t\; \hat{s}(\mathbf{x}_{\text{mid}}, t_{\text{mid}})$ {scaled midpoint update}
  $\mathbf{x}_{t-1} \leftarrow \exp\big(\frac{1}{4}\beta_t\Delta t\big)\,\mathbf{x}$ {$\Phi^L_{\Delta t/2}$}
end for
return $\mathbf{x}_0$