MPR-FLOW: EFFICIENT INVERSE PROBLEM SOLVING FOR SEQUENTIAL MEASUREMENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Many real-world imaging techniques rely on combining multiple measurements acquired sequentially. The final image can then be obtained by solving an inverse problem that accounts for the physical model of the acquisition process. In many applications, it is essential to minimize the number of data acquisitions to avoid multiplying time-consuming, costly, or risky procedures. Being able to obtain the best possible reconstruction as measurements are acquired is thus critical to avoid unnecessary acquisitions. In this setting, each reconstruction can be a useful computation state for the next step. With this in mind, we introduce Measurement-Progressive Reconstruction Flow (MPR-Flow), a plug-and-play flow solver that exploits continuity at two scales. First, across acquisition stages, previous reconstructions are used as a prior to speed up the reconstruction. Second, across solver iterations, the re-noising step is constrained so that well-measured regions are preserved. MPR-Flow needs no retraining, nor any changes in the prior. Experiments on simulated and real datasets show that it outperforms state-of-the-art flow-based inverse solvers in terms of reconstruction quality and acquisition time.

1 INTRODUCTION

A wide range of image reconstruction or restoration techniques rely on solving some inverse problem that models the mechanisms at the root of the deterioration or acquisition process. Such problems are typically ill-posed and high-dimensional, making them challenging to solve without suitable priors. Generative models have been successfully used as regularizers (Ongie et al., 2020; Wei et al., 2022; Bora et al., 2017) and have substantially advanced computational imaging. Reconstruction, however, remains largely an offline process that starts only after a predetermined set of measurements has been collected. In most acquisition systems, measurements are gathered sequentially over an extended acquisition process, which may itself be constrained by time, cost, or safety considerations. This is the case, for instance, in positron emission tomography, where every additional projection increases the radiation dose delivered to the patient. Magnetic resonance imaging (MRI) requires long acquisition times where the patient is exposed to very loud noises, making it poorly tolerated by infants, young children, and animals. Seismic exploration and synthetic aperture radar tomography are other examples where individual measurements are costly and require heavy logistics. Another example is atomic force microscopy (AFM), where a probe is scanned across biological surfaces to reconstruct their topography or map their mechanical properties (e.g. elasticity). High-resolution imaging can require several minutes, during which the biological sample may undergo significant changes or lose viability. This relatively long acquisition time therefore limits the use of AFM for studying living biological samples. For these cases, reconstruction could run concurrently and provide useful information long before the final measurement is collected, or make additional measurements superfluous. The resulting challenge is not only to reconstruct an image from incomplete data, but to update that reconstruction efficiently as more data arrives.

Plug-and-play (PnP) methods are well matched to regularize the reconstruction because they decouple the forward model from the learned image prior, alternating an explicit data-consistency step with a learned denoiser (Venkatakrishnan et al., 2013; Kamilov et al., 2023). A prior trained once can therefore be combined with different measurement operators without retraining an end-to-end reconstruction network. While diffusion and score-based priors have demonstrated strong performance on severely ill-posed inverse problems, their reverse-time sampling often entails many

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.

costly neural-network evaluations (Chung et al., 2023; Daras et al., 2024). Flow matching provides an appealing alternative by learning a continuous transport by simulation-free vector-field regression and admitting comparatively short sampling trajectories through numerical ODE integration (Lipman et al., 2023). Methods such as PnP-Flow and FLOWER exploit this structure by coupling pretrained flow priors with explicit measurement consistency for general forward operators (Martin et al., 2025; Pourya et al., 2026), providing the flexibility needed for reconstruction settings in which the available measurements change over time.

The flexibility of these solvers, however, leaves open how to exploit information already computed during previous acquisition stages. In this setting, we consider that one nominally static object is observed through a progressively enlarged measurement set. This progressive setting is formalized in Section 4.1: the target remains unchanged across stages, and each intermediate reconstruction yields an estimate that can be used to guide the subsequent ones. Early estimates may contain plausible structure but might then be unsupported by new measurements, so the solver must balance retaining useful information and letting new data correct it.

We introduce *MPR-Flow*, which exploits the continuity of the reconstruction problem at two distinct scales. Across acquisition stages, the reconstruction obtained at the previous stage is carried over to initialize the next one: as measurements accumulate, the flow is started at a later time, and the inherited estimate serves as the starting point of the data-fidelity updates, which new measurements can then refine. Across solver iterations, instead of sampling independent Gaussian perturbations at each step, we introduce Reflected Noise Coupling (RNC), which correlates the injected noise across iterations to preserve measurement consistency. RNC preserves the standard Gaussian marginal required by the pretrained prior while modifying the joint trajectory, so successive iterations remain coupled rather than independently randomized.

In this paper, we propose to tackle the problem of measurement-progressive reconstruction through the prism of PnP solvers. Our contributions are as follows:

- To leverage the sequential nature of the task, we design an acquisition-dependent warm start that jointly speeds up the reconstruction and increases the average performance by propagating estimates at each reconstruction stage.
- To favor estimations supported by measurements, we propose the Reflected Noise Coupling (RNC) of consecutive re-embedding noises.
- We evaluate the complete method at matched evaluation budgets on controlled inverse problems and in two real-world imaging settings, AFM and MRI.

2 Related Work

Flow-Based Inverse Solvers Learned priors regularize ill-posed reconstruction while separating image statistics from the measurement operator (Ongie et al., 2020). Classical PnP methods realize this separation through learned denoisers (Venkatakrishnan et al., 2013; Kamilov et al., 2023), whereas generative inverse solvers impose measurement information during inference. Diffusion Posterior Sampling, for example, combines a learned score with an approximate likelihood gradient (Chung et al., 2023). Flow-based solvers replace the diffusion trajectory with a learned transport: PnP-Flow and FLOWER couple an endpoint predictor to explicit measurement consistency (Martin et al., 2025; Pourya et al., 2026), while FlowDPS and Flow-ADMM use, respectively, likelihood guidance and an ADMM splitting (Kim et al., 2025; Sommerhoff & Moeller, 2026).

Sequential Acquisition and Informed Initialization Sequential acquisition methods repeatedly reconstruct from the measurements available so far and use intermediate information to guide further acquisition (Zhang et al., 2019; Van Gorp et al., 2021; Yin et al., 2021). Although information from earlier stages influences subsequent acquisition, each updated reconstruction is computed from the enlarged measurement set rather than continuing from the preceding solution. CCDF forward-diffuses an informative initializer to an intermediate noise time and shortens conditional diffusion for a single inverse problem (Chung et al., 2022), SeqDiff accelerates frame-wise posterior sampling by forward-diffusing the preceding reconstructed frame and using it to initialize a shortened diffusion trajectory for the next frame (Stevens et al., 2025), and Accelerated Sequential Flow Matching propagates posterior information for evolving states (Huang et al., 2026). Warm-Start Flow Matching

uses an auxiliary draft generator in ordinary text and image generation, without a forward operator or measurement consistency (Kim, 2026). However, none of these works considers measurement-progressive inverse problems, where a single static target is reconstructed repeatedly from nested, cumulative measurement sets.

3 PRELIMINARIES

Inverse problems and plug-and-play reconstruction. Let $x^\star \in \mathbb{K}^d$ denote an unknown signal and let $A : \mathbb{K}^d \to \mathbb{K}^m$ be a known forward operator. We consider measurements of the form

$$y = Ax^\star + e, \tag{1}$$

where e denotes measurement noise. When A is non-injective or poorly conditioned, multiple signals may be consistent with the observations, making direct inversion unstable or impossible.

Classical reconstruction methods resolve this ambiguity by combining measurement consistency with prior information:

$$\widehat{x} \in \arg\min_x F_y(x) + \lambda \mathcal{R}(x), \qquad F_y(x) = \frac{1}{2\sigma_y^2}\|Ax - y\|_2^2, \tag{2}$$

where F_y measures agreement with the observations and $\mathcal{R}$ encodes prior structure. Under additive Gaussian noise, the displayed F_y is the negative log-likelihood up to an additive constant. If $\mathcal{R}$ is the negative log-density of a prior p_X, the resulting minimizer is a maximum-a-posteriori estimate. A Bayesian formulation instead considers the posterior

$$p(x \mid y) \propto \exp[-F_y(x)]\, p_X(x), \tag{3}$$

whose samples can capture ambiguities that a single point estimate would suppress.

Plug-and-play (PnP) methods replace an explicit regularization step with a learned denoising operator (Venkatakrishnan et al., 2013; Kamilov et al., 2023). An operator-splitting algorithm alternates between a data-consistency update, which depends on (A, y), and a denoising update representing the learned prior. The denoiser need not be the proximal map of an explicit $\mathcal{R}$, allowing a prior trained once to be combined with different forward operators.

Flow matching. Flow matching learns a time-dependent velocity field that transports a tractable source distribution p_0 to the data distribution p_1 (Lipman et al., 2023). Let $X_0 \sim p_0$ and $X_1 \sim p_1$. For generative time $\tau \in [0, 1]$ and its associated straight conditional path, we use:

$$X_\tau = (1-\tau)X_0 + \tau X_1, \qquad D_\tau(x) = x + (1-\tau)v_\theta(x, \tau) = \mathbb{E}[X_1 \mid X_\tau = x]. \tag{4}$$

Thus, D_τ predicts the clean endpoint from an intermediate flow state.

4 METHOD

4.1 MEASUREMENT-PROGRESSIVE INVERSE PROBLEMS

Many inverse problems are solved repeatedly while measurements are still being acquired. Examples include progressive k-space sampling, sequential inpainting masks, or blockwise sensing trajectories. Although each acquisition stage defines a different inverse problem, all stages share the same latent target and nested measurement history. The reconstruction problem is therefore not a collection of unrelated inverse problems, but a continuation problem indexed by acquisition progress.

We consider a sequence of inverse problems involving a fixed unknown $x^\star \in \mathbb{K}^d$, where $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$. At acquisition stage $s \in \{1, \dots, S\}$, the cumulative observation model is

$$y_s = A_s x^\star + e_s, \qquad e_s \sim \mathcal{N}(0, \sigma_s^2 I). \tag{5}$$

Here, $A_s : \mathbb{K}^d \to \mathbb{K}^{m_s}$ contains all measurements available at stage s. The measurement sets are nested: each stage retains the previous observations and appends a new measurement block. Equivalently,

$$A_s = \begin{bmatrix} A_{s-1} \\ \Delta A_s \end{bmatrix}, \qquad y_s = \begin{bmatrix} y_{s-1} \\ \Delta y_s \end{bmatrix}, \qquad m_{s-1} < m_s, \tag{6}$$

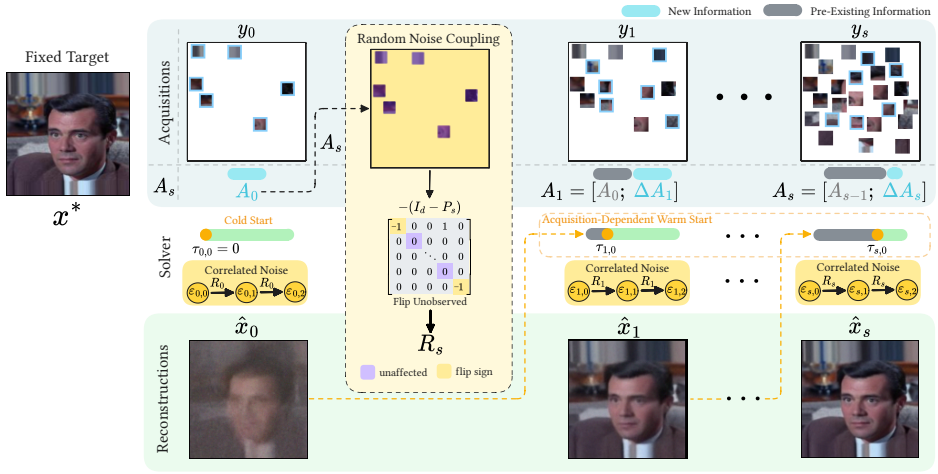

Figure 1: Conceptual illustration of measurement-progressive reconstruction. A fixed target $x^\star$ is observed through a sequence of nested operators $A_s = [A_{s-1}; \Delta A_s]$ (pre-existing measurements in gray, newly acquired block ΔA_s in cyan); reconstructions $\widehat{x}_s$ at each stage are shown below. Within each reconstruction stage, noise is correlated form iteration to iteration.

where Δy_s and ΔA_s are the new measurements and it's associated measurement operator. $m_s \in \mathbb{N}$ is the number of cumulative measurements available at stage s, so that y_s and $e_s \in \mathbb{K}^{m_s}$. This representation covers, for example, cumulative Fourier samples, newly observed image locations, and measurement blocks collected along a prescribed acquisition trajectory.

Let $F_s(x) = \frac{1}{2\sigma_s^2}\|A_s x - y_s\|_2^2$ denote the data-fidelity term at stage s. Although the target $x^\star$ is unchanged, the reconstruction objective changes because F_s incorporates a strictly larger measurement set. We call this setting *measurement-progressive reconstruction*.

At each stage s, we query the generative model N times to produce a reconstruction $\widehat{x}_s$. Within stage s, the current reconstruction is $x_{s,n}$, where n indexes model evaluations. For $s > 1$, the preceding reconstruction $\widehat{x}_{s-1}$ is therefore available as computational state for solving the enlarged inverse problem. A conventional flow-based inverse solver solves each acquisition stage from scratch, traversing the full flow interval $\tau \in [0, 1]$ independently of the preceding reconstruction; we instead reuse $\widehat{x}_{s-1}$ and solve stage s over a reduced interval $[\tau_{0,s}, 1]$.

4.2 MPR-Flow

MPR-Flow performs progressive reconstruction along two nested axes rather than as a collection of isolated solves. For each *acquisition* s, the previous estimate $\widehat{x}_{s-1}$ is carried into the new problem defined by the enlarged pair (A_s, y_s). The generative model is then executed from an acquisition-dependent starting time which accounts for how much information was added. At the *solver scale*, RNC constrains the structure of the noise according to the image subspace of A_s so that directions not related to existing measurements see their noise contribution canceling out during the reconstruction.

Acquisition-Scale Continuation The global MPR-Flow update is built upon the fidelity–renoising–prediction decomposition introduced in (Martin et al., 2025). For $s > 1$, the initial reconstruction state is given by

$$x_{s,0} = \widehat{x}_{s-1}. \tag{7}$$

First, during the data-consistency update, the embedding is moved toward the solution of the inverse problem:

$$z_{s,n} = x_{s,n} - \eta(\tau_{s,n})\nabla F_s(x_{s,n}), \qquad \eta(\tau) = (1 - \tau_{s,n})^\alpha. \tag{8}$$

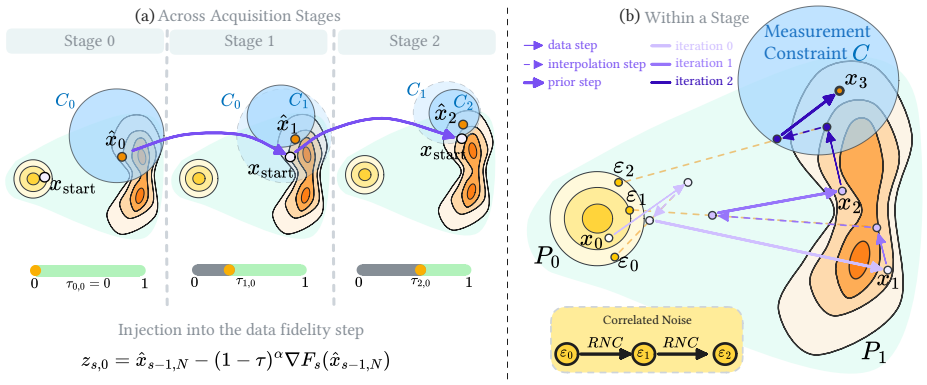

Figure 2: Overview of MPR-Flow's two-scale continuity. **(a)** Across acquisition stages, the preceding reconstruction is injected into the data-fidelity step of next solve, where it's afterwards interpolated back to a latter flow time $\tau_{s,0}$.**(b)** Within each stage, the algorithm takes a gradient step in the direction of the measurement constraint, then it's interpolated back with noise to stay in the probability path, and lastly pushed towards the target distribution by a denoising step. Each noise is correlated iteration to iteration.

Second, the corrected estimate is re-embedded at the current flow time and mapped to a clean endpoint:

$$\widetilde{z}_{s,n} = \tau_{s,n} z_{s,n} + (1-\tau)\epsilon_{s,n}, \qquad \epsilon_{s,n} \sim \mathcal{N}(0,1), \tag{9}$$

$$x_{s,n+1} = D_\tau(\widetilde{z}_{s,n}) = \widetilde{z}_{s,n} + (1-\tau_{s,n}) v_\theta(\widetilde{z}_{s,n}, \tau_{s,n}). \tag{10}$$

The completed stage reconstruction is $\widehat{x}_s = x_{s,N}$. The exponent α controls a separate interaction with data consistency. A larger starting time has two coupled effects in equation 9: it retains more of the data-consistent state and injects less fresh source noise.

Acquisition-Dependent Warm Start For a given stage s, the previous reconstruction $\hat{x}_{s-1}$ is used for initialization as it might provide a good guess for $\hat{x}_s$. However, the impact of previous estimates should be weighted by the amount of information available during their reconstruction. A fixed warm start would assign the same confidence to every preceding reconstruction, even early reconstructions supported by very few measurements. As the reconstructions' reliability improves at each stage, we propose an *Acquisition-Dependent Warm Start* (ADWS) approach that accordingly adjusts the starting time for the flow-matching trajectory.

Let $\tau_{\max} \in [0, 1)$ be the largest permitted starting time. We define:

$$w_s = \min\left\{\frac{s-1}{K}, 1\right\} \qquad \text{and} \qquad \tau_{0,s} = w_s \tau_{\max}, \tag{11}$$

with K being a hyperparameter controlling the slope of w_s. $\tau_{0,1} = 0$ while $\tau_{0,s}$ approaches $\tau_{\max}$ as cumulative measurement support grows. For every $s > 1$, the preceding reconstruction remains the reference initialization, resulting in acquisition-scale continuation from one reconstruction stage to the next. The schedule controls the strength of this continuation by determining how strongly the preceding estimate is re-noised and where the inner flow trajectory begins. The linear ramp in equation 11 encodes increasing trust in the preceding reconstruction without requiring a learned reliability estimator or additional sequence training.

Reflected Noise Coupling Sampling the noise independently from the observation leads to a trade-off between two regimes: with a high enough noise level, the model can regenerate unobserved regions, but it can alter well-measured ones; re-noising too little could limit the prior's ability to revise untrusted ones. Furthermore, noise can also accumulate in unmeasured directions, making the solver drift away from the observations.

To limit the noise accumulation while retaining the prior's ability to revise erroneous reconstructions, we introduce *Reflected Noise Coupling* (RNC). RNC exploits the flow matching temporal

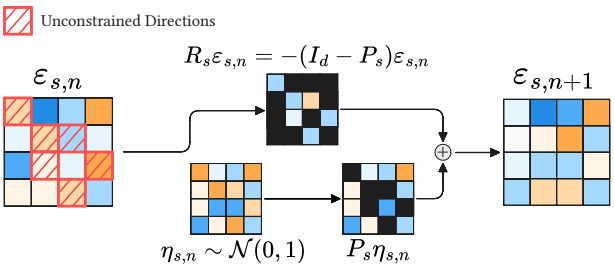

Figure 3: Illustration of Reflected Noise Coupling within an acquisition stage. In constrained dimensions (red), the corresponding pixels are flipped, while all remaining entries are set to zero (black), and populated with independently drawn noise $\eta_{s,n}$

dependence by adding the opposite of the previous perturbation in the measurement nullspace. This trick has two major consequences. First, it keeps the noise level constant in the input, thus being compatible with the flow matching denoising scheme. Second, across iterations, it partially cancels out the carried-over effect of the noise on unobserved image components, thus favoring convergence toward solutions aligned with the measurements.

Let A_s map a signal to its predicted measurements at stage s, and let A_s^* denote its adjoint, which maps measurement-space vectors back into signal space. The subspace $\mathrm{range}(A_s^*)$ contains the signal directions directly constrained by the measurements. Its orthogonal complement is $\ker(A_s)$: adding a component in this nullspace leaves the predicted measurements unchanged. We therefore let P_s be the orthogonal projector onto $\mathrm{range}(A_s^*)$, so that $I - P_s$ projects onto $\ker(A_s)$. We can then define:

$$R_s = -(I - P_s), \tag{12}$$

Thus, the noise is independently sampled in the measurement-supported directions, while in the measurement null-space the noise is reversed. When the null-space perturbations propagate, this coupling can counteract their accumulation over consecutive iterations. Combining the two components, we update the noise according to:

$$\varepsilon_{k+1} = P_s\eta_k + R_s\varepsilon_k, \qquad \eta_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I). \tag{13}$$

where η_k is independent of ε_k. Consecutive noise injections are therefore independent in measurement-supported directions and negatively correlated in the nullspace. This preserves fresh stochastic exploration where the measurements constrain the reconstruction, while introducing exact sign reversal where perturbations cannot be corrected directly by the data. The RNC procedure can be visualized in Figure 3.

Importantly, RNC changes the temporal dependence of the noise sequence without changing its distribution at any individual iteration. If P_s is an orthogonal projector and $\varepsilon_0 \sim \mathcal{N}(0, I)$, then $\varepsilon_k \sim \mathcal{N}(0, I)$ for every k. The two terms in equation 13 are independent Gaussian vectors supported on orthogonal complementary subspaces, with covariance $P_s + (I - P_s) = I$. At the same time, consecutive noise vectors have cross-covariance $\mathrm{Cov}(\varepsilon_{k+1}, \varepsilon_k) = -(I - P_s) = R_s$. We refer the reader to Appendices B.1 to B.2 for the proof and how to estimate P_s.

Complete MPR-Flow Algorithm Algorithm 1 summarizes the complete MPR-Flow procedure. At each acquisition stage, the method selects an acquisition-dependent starting flow time and initializes the reconstruction from the preceding stage, while enforcing data consistency using the complete current measurement pair (A_s, y_s). Each inner iteration then applies a data-fidelity update, re-embeds the corrected estimate at the current flow time, predicts its clean endpoint, and updates the perturbation through RNC.

5 RESULTS

5.1 EXPERIMENTAL SETUP

We evaluate MPR-Flow on four measurement-progressive inverse problems. In every episode, the target is fixed and measurements accumulate monotonically. Random inpainting, random-matrix sensing, and AFM start from no measurements; MRI starts from a fixed sampled central region of k-space used for autocalibration. The controlled problems and MRI use $S = 10$ acquisition stages. All images were normalized to the range $[-1, 1]$.

We adapt the inverse-problem benchmarking framework of Martin et al. (2025) to measurement-progressive reconstruction, using a common forward-operator and adjoint interface to supply each method with identical cumulative measurements at every stage. Within each experiment, all methods share the same fixed flow-matching U-Net prior (Tong et al. (2024)).

Random inpainting and matrix sensing Both these inverse problems use a CelebA-trained prior (Liu et al. (2015)), for inpainting A_s is a diagonal pixel mask that reveals equally sized sets of previously unseen pixels at each stage. We consider final sampling rates of 10% and 30% of the pixels in the image. For random sensing, A_s is formed from the first m_s rows of a fixed row-normalized Gaussian matrix, growing from zero to approximately 5% of the number of values in the latent image over the course of 10 equal acquisition stages.

Partial MRI. We train a separate flow prior on the NYU fastMRI brain dataset Zbontar et al. (2018); Knoll et al. (2020); Muckley et al. (2021) and evaluate on a held-out split from the same distribution. The forward operator is $A_s = M_s \odot F$, where F is an orthonormal 2D DFT and M_s is a Cartesian phase-encoding mask. The initial mask contains 31 autocalibration lines, and each of the following 10 stages requests eight additional lines. A single k-space noise realization is drawn per measurement episode and masked at each stage. The encoded images are of dimension 384×384.

AFM AFM measures surface topography of samples sequentially along a scanning trajectory. We train the prior on simulated DNA (DeoxyriboNucleic Acid) AFM images generated using ASAP from the DeepTrack (Midtvedt et al., 2021) framework. Thin DNA filaments occupy only a small fraction of the surface, making early measurements weakly informative on the structures of interest. For AFM, 30 acquisition stages are used, with each carrying an equal quantity of sampling points, up to the final acquisition reaching 30% of the total measurement.

Baselines We compare against cold-started PnP-Flow (Martin et al., 2025), FLOWER (Pourya et al., 2026), OT-ODE (Onken et al. (2020)), and Flow-Priors (Zhang et al., 2024) each run independently at every acquisition stage with no reconstruction state carried across stages and tuned per task and budget. Each candidate configuration is evaluated over a complete acquisition sequence and ranked by mean per-image PSNR at the final stage, with acquisition sequences and evaluation seeds held fixed across candidates. For MPR-Flow and the two most competitive baselines Flower and PnP-Flow, we select hyperparameters via grid search. For all other baselines, we adopt the values reported in their respective papers. The selected configurations are detailed in Appendix F.

Evaluation All methods receive the same N flow-network evaluations per stage, and we report PSNR and SSIM over 128 test images. For AFM, we additionally report fragment-count mean squared error (Frag. MSE), the average squared difference between the numbers of reconstructed and ground-truth fragments. This provides a simple structural measure of whether a reconstruction preserves the expected number of disconnected filament components, complementing pixelwise fidelity metrics that may remain high despite broken or split structures. More detail on the accuracy of DNA strand completion in AFM imaging, and explanations of the metrics used can be found in Appendix D. Unless stated otherwise, results use the same seed and spreads are across images. Results of the test with matching iteration budgets can be seen in Table 1.

5.2 BASELINE COMPARISON

We present the performance achieved by MPR-Flow and baseline models in 1. On inpainting problems, MPR-Flow achieved significantly higher PSNR and SSIM than the baselines, whether the

Method ($N = 10$)	Inpainting 10%		Inpainting 30%		AFM DNA 30%		Random matrix		MRI $\sigma = 0$		MRI $\sigma = 0.03$	
	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	Frag. MSE ↓	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
Flow Priors† (Zhang et al. (2024))	14.6	0.358	14.8	0.379	29.5	7.87	10.6	0.167	21.8	0.544	22.0	0.544
OT-ODE‡ (Onken et al. (2020))	18.4	0.446	24.7	0.744	29.4	6.19	11.1	0.220	29.5	0.603	29.4	0.606
Flower (Pourya et al. (2026))	15.5	0.359	25.0	0.738	33.1	4.23	13.8	0.255	32.2	0.844	32.1	0.850
PnP-Flow (Martin et al. (2025))	16.3	0.392	26.1	0.794	36.7	3.54	14.7	0.303	30.6	0.654	30.6	0.671
MPR-Flow	**24.6**	**0.778**	**32.0**	**0.941**	**42.4**	**1.29**	**25.1**	**0.735**	**36.1**	**0.878**	**35.5**	**0.893**
Method ($N = 50$)	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	Frag. MSE ↓	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
Flow Priors† (Zhang et al. (2024))	23.4	0.727	27.7	0.883	33.2	4.05	9.9	0.225	30.5	0.682	30.4	0.671
OT-ODE‡ (Onken et al. (2020))	17.7	0.362	25.9	0.781	31.1	5.42	11.0	0.196	31.6	0.757	31.7	0.786
Flower (Pourya et al. (2026))	21.9	0.640	30.5	0.914	37.3	2.09	18.7	0.487	34.3	0.869	33.9	0.874
PnP-Flow (Martin et al. (2025))	23.0	0.696	31.0	0.923	40.4	2.20	19.9	0.552	34.3	0.866	34.0	0.875
MPR-Flow	**25.8**	**0.826**	**32.8**	**0.953**	**42.7**	**0.92**	**28.2**	**0.837**	**36.4**	**0.884**	**35.7**	**0.903**

Table 1: Reconstruction quality inverse problems for $N \in \{10, 50\}$ reconstruction steps. Values are means across 128 test images. Higher PSNR and SSIM are better. For AFM DNA, Frag. MSE is the mean squared fragment-count error against ground truth. Within each reconstruction budget, best results are shown in **bold** and second-best results in blue. †Flow Priors uses $2\times$ the stated forward NFEs, plus derivative computations. ‡OT-ODE matches forward NFEs but additionally requires backpropagation at each iteration.

final sub-sampled fraction of the pixels is 30% or 10%. For random inpainting at 10%, MPR-Flow achieves around $+6$dB and $+8$dB PSNR increase over OT-ODE and PnP-Flow, with corresponding SSIM gains of $+0.33$ and $+0.39$. We also evaluate AFM reconstruction of DNA images, where the sequential nature of the scan makes the reconstruction a genuine instance of progressive inpainting. Here, the structures of interest are highly localized, while sequential measurements only uncover up to 30% of the total image. Again, MPR-Flow brings large gains both for PSNR and Frag. MSE shows about $+2.3$dB increase in PSNR and -1.2 Frag. MSE compared to the best baseline.

In matrix sensing, measurements are mixed together and have to be unmixed to solve the inverse problem, making this task sensing more challenging than inpainting. MPR-Flow shows far better PSNR and SSIM compared to methods that do not exploit the forward operator structure across stages, with a $+6$dB PSNR and $+0.32$ SSIM increase compared to PnP-Flow. MRI is a real-world sensing case where the measurements are acquired sequentially under structured Fourier measurements. For this task, MPR-Flow shows strong and consistent gains for both clean and noisy data. Compared to Flower, PSNR is improved by about $+4$dB and SSIM by $+0.34$.

MPR-Flow's main benefit is high reconstruction quality with limited computation. As can be observed in Figure 4, with just 8 model evaluations, our model surpasses the baselines at 100 model evaluations in both PSNR and SSIM. Across problems, Table 3 in Appendix C shows that MPR-Flow matches the performance of FLOWER and PnP-Flow in inverse problems with NFE reductions from $10\times$ to $33\times$.

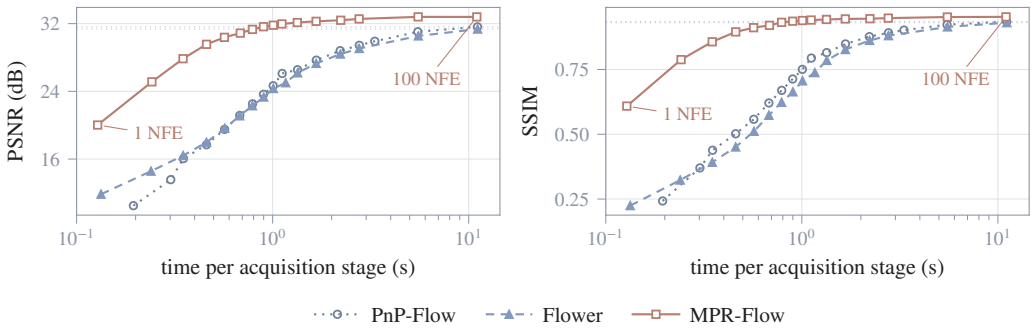

Figure 4: Quality against measured reconstruction time on the 30% inpainting problem, with $N \in \{1, \ldots, 10, 12, 15, 20, 25, 50, 100\}$. Values are means across 128 test images. PnP-Flow is cold-started at each budget; "continuation only" omits RNC. Time is the reported latency per acquisition stage for a batch of 32 images. The time axis is logarithmic.

Table 2: Ablation of RNC and ADWS at $N = 10$ reconstruction steps. PnP-Flow uses neither component, while MPR-Flow combines both. The intermediate rows enable each of these components independently. Values are means across 128 test images. Higher PSNR and SSIM are better. For each metric within each inverse problem, best results are shown in **bold** and second-bests in blue.

			Inpainting 10%		Inpainting 30%		Random matrix		MRI $\sigma = 0$		MRI $\sigma = 0.03$	
Method	RNC	ADWS	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
PnP-Flow	–	–	16.3	0.392	26.1	0.794	14.7	0.303	30.6	0.654	30.6	0.671
+ RNC	✓	–	16.6	0.415	27.0	0.826	14.7	0.317	30.8	0.632	30.8	0.646
+ ADWS	–	✓	23.6	0.738	30.8	0.924	24.2	0.703	35.4	**0.892**	34.7	0.888
MPR-Flow	✓	✓	**24.6**	**0.778**	**32.0**	**0.941**	**25.1**	**0.735**	**36.1**	0.878	**35.5**	**0.893**

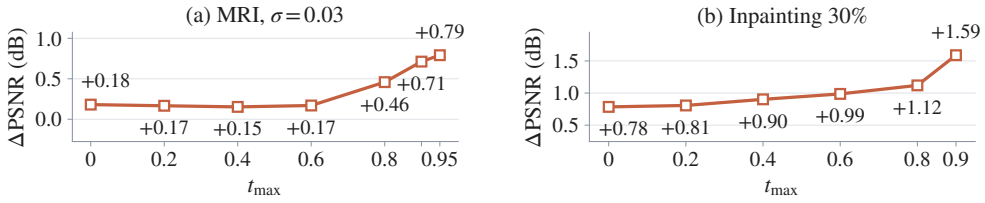

Figure 5: PSNR gain obtained by combining RNC with ADWS, relative to ADWS alone, as a function of the start time $t_{\max}$.

5.3 Ablation Study

Table 2 isolates the contributions of MPR-Flow by applying its components to PnP-Flow one at a time at the low-budget setting of $N = 10$ flow evaluations per stage. ADWS provides the largest individual improvement, substantially increasing both PSNR and SSIM across all tasks. RNC produces consistent gains on top of PnP-Flow. Combining the two components yields the highest PSNR and SSIM in nearly all settings, which shows the complementarity of the two contributions. Indeed, RNC gives the denoiser more freedom to correct later-time updates by removing perturbations accumulating in the null-space. This complementarity is highlighted in Figure 5, where the latter is the warm start; the higher values are gains brought by combining RNC and ADWS.

6 Discussion and Limitations

The gain in reconstruction quality depends on how much useful structure can be reused from the preceding estimate.As inpainting stages expose non-overlapping pixel subsets and random-matrix stages add i.i.d. rows, $\widehat{x}_{s-1}$ provides precisely the components the current reconstruction needs, yielding the potentially large improvements. By contrast, MRI stages have greater overlap, so a separate trajectory can reconstruct more on its own and the resulting gain might be lower.

7 Conclusion

We introduced MPR-Flow, a plug-and-play solver for measurement-progressive inverse problems that combines acquisition-dependent warm starting with reflected noise coupling. Our experiments demonstrated higher reconstruction quality than the evaluated baselines, even at lower computational budgets. These results support treating reconstruction of a static target from an expanding set of measurements as a continuation problem, rather than a series of disconnected inverse problems, and it warrants a dedicated type of solver. Passing $\widehat{x}_{s-1}$ forward to the next stage and initiating the flow later, once sufficient support has accumulated, captures most of the attainable gains, and correlating the re-embedding noise through an orthogonal reflection provides an additional, reliable boost. These improvements in the sequential setting carry over across different forward operators, and they are substantial enough to outperform stagewise independent solvers even when those use larger reconstruction budgets.

AI USE STATEMENT

AI tools were used to assist in the formulation of mathematical claims, writing and debugging code, provide feedback on experiment results, to create, edit, and debug code, to identify relevant literature, and to improve the clarity and readability of the manuscript. The authors assume full responsibility for the final manuscript, including mathematical claims, code, and reported results.

ETHICS STATEMENT

The MRI experiments use the NYU fastMRI dataset Zbontar et al. (2018); Knoll et al. (2020); Muckley et al. (2021), whose metadata are described as de-identified in the dataset sharing agreement. Reconstructed images are evaluated for research purposes, including qualitative comparisons, and thus are not intended for diagnosis or patient care. Learned reconstruction methods may introduce plausible structures unsupported by the measurements or suppress relevant features; therefore, the reported reconstruction performance does not establish clinical safety.

REPRODUCIBILITY STATEMENT

The paper describes the measurement-progressive formulation, reconstruction procedure, and evaluation protocol. The algorithm is given in Appendix A.1. Section 5 specifies the datasets, acquisition schedules, baselines, evaluation metrics, and descriptions of the inverse problems. Hyperparameters used are given in F. In addition, the source code is available in anonymous.4open.science/r/MPR-Flow-7BE6/README.md.

REFERENCES

ASAP. AFM Simulation and Analysis Notebooks. `https://github.com/DeepTrackAI/ASAP`. GitHub repository.

Haim Avron and Sivan Toledo. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM*, 58(2):8:1–8:34, 2011. doi: 10.1145/1944345.1944349.

Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 537–546. PMLR, 06–11 Aug 2017. URL `https://proceedings.mlr.press/v70/bora17a.html`.

Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12403–12412, New Orleans, LA, USA, June 2022. doi: 10.1109/CVPR52688.2022.01209. URL `https://ieeexplore.ieee.org/document/9879691/`.

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL `https://openreview.net/forum?id=OnD9zGAGT0k`.

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.

Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53 (2):217–288, 2011. doi: 10.1137/090771806.

Yinan Huang, Hans Hao-Hsun Hsu, Junran Wang, Bo Dai, and Pan Li. Accelerated Sequential Flow Matching: A Bayesian Filtering Perspective, May 2026. URL `http://arxiv.org/abs/2602.05319`. arXiv:2602.05319 [cs.LG].

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.

Jeongsol Kim, Bryan Sangwoo Park, and Jong Chul Ye. FlowDPS: Flow-driven posterior sampling for inverse problems. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.

Minyoung Kim. Warm-Start Flow Matching for Guaranteed Fast Text/Image Generation, March 2026. URL http://arxiv.org/abs/2603.19360. arXiv:2603.19360 [cs.LG].

Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1):e190007, 2020. doi: 10.1148/ryai.2020190007.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 3730–3738, 2015. doi: 10.1109/ICCV.2015.425.

Ségolène Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. PnP-Flow: Plug-and-play image restoration with flow matching. In *International Conference on Learning Representations (ICLR)*, 2025. URL https://openreview.net/forum?id=5AtHrq3B5R.

Benjamin Midtvedt, Saga Helgadottir, Aykut Argun, Jesús Pineda, Daniel Midtvedt, and Giovanni Volpe. Quantitative digital microscopy with deep learning. *Applied Physics Reviews*, 8(1):011310, 2021. doi: 10.1063/5.0034891.

Matthew J. Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, Simon Arberet, Dominik Nickel, Zaccharie Ramzi, Philippe Ciuciu, Jean-Luc Starck, Jonas Teuwen, Dimitrios Karkalousos, Chaoping Zhang, Anuroop Sriram, Zhengnan Huang, Nafissa Yakubova, Yvonne W. Lui, and Florian Knoll. Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Transactions on Medical Imaging*, 40(9):2306–2317, 2021. doi: 10.1109/TMI.2021.3075856.

Gregory Ongie, Ajil Jalal, Christopher A. Metzler, Richard G. Baraniuk, Alexandros G. Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020. doi: 10.1109/JSAIT.2020.2991563.

Derek Onken, Samy Wu Fung, Xingjian Li, and Lars Ruthotto. Ot-flow: Fast and accurate continuous normalizing flows via optimal transport. In *AAAI Conference on Artificial Intelligence*, 2020. URL https://api.semanticscholar.org/CorpusID:219176997.

Mehrsa Pourya, Bassam El Rawas, and Michael Unser. Flower: A flow-matching solver for inverse problems. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://openreview.net/forum?id=QGd34p02mI.

Hendrik Sommerhoff and Michael Moeller. FlowADMM: Plug-and-play ADMM with flow-based renoise–denoise priors. *arXiv preprint arXiv:2605.08640*, 2026.

Tristan S. W. Stevens, Oisín Nolan, Jean-Luc Robert, and Ruud J. G. van Sloun. Sequential Posterior Sampling with Diffusion Models. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, April 2025. doi: 10.1109/ICASSP49660.2025.10889752. URL http://arxiv.org/abs/2409.05399. arXiv:2409.05399 [cs.CV].

Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=CD9Snc73AW. Expert Certification.

Hans Van Gorp, Iris Huijben, Bastiaan S Veeling, Nicola Pezzotti, and Ruud J. G. Van Sloun. Active deep probabilistic subsampling. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 10509–10518. PMLR, 2021.

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

Xinyi Wei, Hans van Gorp, Lizeth Gonzalez-Carabarin, Daniel Freedman, Yonina C. Eldar, and Ruud J. G. van Sloun. Deep unfolding with normalizing flow priors for inverse problems. *IEEE Transactions on Signal Processing*, 70:2962–2971, 2022. doi: 10.1109/TSP.2022.3179807.

Tianwei Yin, Zihui Wu, He Sun, Adrian V. Dalca, Yisong Yue, and Katherine L. Bouman. End-to-end sequential sampling and reconstruction for MRI. In *Proceedings of Machine Learning for Health (ML4H)*, volume 158 of *Proceedings of Machine Learning Research*, pp. 261–281, 2021.

Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael G. Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated mri. *ArXiv*, abs/1811.08839, 2018. URL https://api.semanticscholar.org/CorpusID:53759905.

Yasi Zhang, Peiyu Yu, Yaxuan Zhu, Yingshan Chang, Feng Gao, Ying Nian Wu, and Oscar Leong. Flow priors for linear inverse problems via iterative corrupted trajectory matching. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 57389–57417. Curran Associates, Inc., 2024. doi: 10.52202/079017-1829. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/698570ae5ec88e07e9f4547f831b4593-Paper-Conference.pdf.

Zizhao Zhang, Adriana Romero, Matthew J. Muckley, Pascal Vincent, Lin Yang, and Michal Drozdzal. Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2049–2058, June 2019.

A Detailed Algorithms and Solver Updates

A.1 MPR-Flow Algorithm

The resulting algorithm shown in Algorithm 1 makes use of both stage-wise continuation and RNC for iteration-level continuation.

Algorithm 1: MPR-Flow for nested measurements

Input: $\{(A_s, y_s)\}_{s=1}^{S}$, flow predictor D_τ, budgets $\{N_s\}_{s=1}^{S}$, and $\tau_{\max}, K_{\text{meas}}, \alpha$
Output: stage reconstructions $\{\widehat{x}_s\}_{s=1}^{S}$
for $s \leftarrow 1$ **to** S **do**
  Compute $\tau_{0,s}$ by equation 11 and set $\tau_{s,0:N_s} \leftarrow \textsc{Grid}(\tau_{0,s}, 1, N_s)$;
  if $s = 1$ **then**
    Draw $x_{s,0} \sim \mathcal{N}(0, I)$;
  else
    $x_{s,0} \leftarrow \widehat{x}_{s-1}$;
  Construct the orthogonal projector P_s onto $\mathrm{range}(A_s^*)$ and set $R_s \leftarrow -(I - P_s)$;
  Draw $\varepsilon_{s,0} \sim \mathcal{N}(0, I)$ independently;
  for $n \leftarrow 0$ **to** $N_s - 1$ **do**
    $\tau \leftarrow \tau_{s,n}$;
    $z_{s,n} \leftarrow x_{s,n} - (1-\tau)^\alpha \nabla F_s(x_{s,n})$;
    $\widetilde{z}_{s,n} \leftarrow \tau z_{s,n} + (1-\tau)\varepsilon_{s,n}$;
    $x_{s,n+1} \leftarrow D_\tau(\widetilde{z}_{s,n})$;
    if $n < N_s - 1$ **then**
      Draw $\xi_{s,n+1} \sim \mathcal{N}(0, I)$ independently;
      $\varepsilon_{s,n+1} \leftarrow P_s \xi_{s,n+1} + R_s \varepsilon_{s,n}$;
  $\widehat{x}_s \leftarrow x_{s,N_s}$;
return $\{\widehat{x}_s\}_{s=1}^{S}$;

A.2 Details of Acquisition-Dependent Warm Start on Flower

Acquisition-Dependent Warm Start is a versatile approach that can easily be combined with other Flow-based inverse problem solvers such as Flower Pourya et al. (2026). Let $\tau_{s,0} = \tau_{\text{start}}(s)$ denote the scheduled start time. The destination consumed at inner iteration n is

$$d_{s,n} = \begin{cases} \widehat{x}_{s-1}, & n = 0,\ s > 1, \\ D_\tau(x_{s,n}) = x_{s,n} + (1-\tau) v_\theta(x_{s,n}, \tau), & \text{otherwise}, \end{cases} \qquad \tau = \tau_{s,n}. \tag{14}$$

The first stage is unaffected: $\tau_{1,0} = 0$ and $x_{1,0} \sim p_0$, so the standard destination prediction is used throughout.

With residual variance

$$\lambda_\tau = \left(\frac{1-\tau}{\sqrt{\tau^2 + (1-\tau)^2}} \right)^2, \tag{15}$$

the refined destination $\mu_{s,n}$ solves

$$\left(\frac{A_s^* A_s}{\sigma_{\text{F}}^2} + \frac{I}{\lambda_\tau} \right) \mu_{s,n} = \frac{A_s^* y_s}{\sigma_{\text{F}}^2} + \frac{d_{s,n}}{\lambda_\tau}, \tag{16}$$

approximately, by conjugate gradients initialized at $d_{s,n}$. Here σ_{F} is the precision parameter of the Flower refinement, distinguished from the injected measurement noise. The destination is then embedded at the next flow time,

$$x_{s,n+1} = \tau_{s,n+1} \mu_{s,n} + \big(1 - \tau_{s,n+1}\big) \epsilon_{s,n}, \qquad \epsilon_{s,n} \sim p_0. \tag{17}$$

Since $\tau_{s,N_s} = 1$, the stage terminates on a refinement and $\widehat{x}_s = \mu_{s,N_s-1}$ contains no injected noise.

In the null space of A_s, the measurement term in equation 16 vanishes and the refinement retains the seed. Along measured singular directions, current data and the seed are balanced by σ_{F} and $\lambda_{\tau_{s,0}}$.

New measurements can then overwrite inherited components. Substituting $\widehat{x}_{s-1}$ for one predicted destination saves one network evaluation but adds a seeded linear solve. A continued stage uses N evaluations of v_θ and $N+1$ refinement solves in the matched-NFE convention.

B Reflective Noise Coupling: Construction and Guarantees

This appendix describes the operator construction and proves preservation of the Gaussian noise marginal. Sections B.1 and B.2 give the projector construction and the exact Gaussian guarantees. Section B.3 addresses the numerical projector error.

Within each acquisition stage, the reflection is fixed and constructed independently of the reconstruction noise. Proofs are given in real coordinates, the complex case uses Hermitian adjoints and proper complex Gaussian noise.

B.1 Operator construction

Let A_s map a signal to its measurements, and let A_s^* denote its adjoint. The subspace $(A_s^*) = \ker(A_s)^\perp$ contains the directions supported by the measurements, while perturbations in $\ker(A_s)$ are unsupported by the measurements. Let P_s be orthogonal projector onto (A_s^*), RNC uses

$$R_s = -(I_d - P_s), \qquad R_s\varepsilon_k = -(I_d - P_s)\varepsilon_k. \tag{18}$$

This operator reverses the inherited noise in the measurement nullspace and removes its measurement-supported component. The latter is replaced by independently sampled noise in the update below.

Proposition 1 (Measurement-informed reflection)**.** *Assume that the normal operator $H_s = A_s^* A_s$ satisfies*

$$H_s^2 = \lambda_s H_s, \qquad \lambda_s > 0. \tag{19}$$

Then $P_s = \frac{H_s}{\lambda_s}$ is the orthogonal projector onto $\mathrm{range}(A_s^*)$*. Consequently, for $R_s = -(I_d - P_s)$, we have:*

$$R_s^* = R_s, \qquad P_s R_s = R_s P_s = 0 \qquad R_s^2 = R_s^* R_s = I_d - P. \tag{20}$$

Proof. Since H_s is self-adjoint,

$$P_s^* = P_s, \qquad P_s^2 = \frac{H_s^2}{\lambda_s^2} = \frac{H_s}{\lambda_s} = P_s. \tag{21}$$

Thus P_s is an orthogonal projector. The identity $\langle v, H_s v\rangle = \|A_s v\|^2$ implies $\ker(H_s) = \ker(A_s)$, and hence

$$\mathrm{range}(P_s) = \ker(H_s)^\perp = \ker(A_s)^\perp = \mathrm{range}(A_s^*). \tag{22}$$

Finally, $R_s = P_s - I_d$ is self-adjoint and satisfies

$$P_s R_s = R_s P_s = P_s^2 - P_s = 0, \qquad R_s^2 = (P_s - I_d)^2 = I_d - P_s. \tag{23}$$

□

Consequently the nullspace noise in inherited with a flipped sign, while the preceding noise is discarded for the measurement-supported components and then resampled independently.

For inpainting, $H_s = M_s^{\mathrm{pix}}$ is a binary pixel mask. For ideal Cartesian MRI, $H_s = F^* M_s^{\mathrm{F}} F$, where F is unitary and M_s^{F} is a binary Fourier mask. Both satisfy $H_s^2 = H_s$, and therefore yield

$$\begin{aligned} &P_s = M_s^{\mathrm{pix}}, && R_s = -(I_d - M_s^{\mathrm{pix}}) && \text{(inpainting)},\\ &P_s = F^* M_s^{\mathrm{F}} F, && R_s = -F^*(I_d - M_s^{\mathrm{F}})F && \text{(ideal Cartesian MRI)}. \end{aligned} \tag{24}$$

The split is thus made in pixel or Fourier coordinates, as appropriate. When the scaled-projector construction is rejected, including in the random-matrix configuration, we use $R_s = -I_d$. This is also an exact reflection and it reverses the entire inherited perturbation.

B.2 GAUSSIAN MARGINALS AND TEMPORAL DEPENDENCE

Starting from $\varepsilon_0 \sim \mathcal{N}(0, I_d)$, the update is

$$\varepsilon_{k+1} = R\varepsilon_k + P\eta_{k+1}, \qquad \eta_{k+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d), \tag{25}$$

with innovations independent of ε_0. Here, the stage index is suppressed, $P = P_s$ is a fixed orthogonal projector, and $R = -(I_d - P)$. The global-alternation fallback corresponds to $P = 0$ and $R = -I_d$.

Proposition 2 (Marginal preservation and lag covariance)**.** *For a fixed orthogonal projector P and $R = -(I_d - P)$, the noise vectors in equation 25 are jointly Gaussian and satisfy*

$$\varepsilon_k \sim \mathcal{N}(0, I_d), \qquad \mathrm{Cov}(\varepsilon_{k+\ell}, \varepsilon_k) = R^\ell, \quad k, \ell \geq 0. \tag{26}$$

Proof. Every finite collection of noise vectors is a linear transformation of ε_0 and the independent Gaussian innovations, and is therefore jointly Gaussian with zero mean. Suppose $\mathrm{Cov}(\varepsilon_k) = I_d$. Independence of the next innovation yields

$$\begin{aligned}\mathrm{Cov}(\varepsilon_{k+1}) &= R\mathrm{Cov}(\varepsilon_k)R^\top + PP^\top \\ &= RR^\top + PP^\top = (I_d - P) + P = I_d.\end{aligned}$$

Since the initialization has covariance I_d, induction proves the marginal claim. For $\ell \geq 1$, expanding the recursion gives

$$\varepsilon_{k+\ell} = R^\ell \varepsilon_k + \sum_{j=1}^{\ell} R^{\ell-j} P \eta_{k+j}.$$

The innovations in the sum are independent of ε_k, so the sum has zero cross-covariance with ε_k. The remaining term gives $R^\ell \mathrm{Cov}(\varepsilon_k) = R^\ell$. The case $\ell = 0$ is the marginal covariance. □

For $R = -(I_d - P)$, we have $R^\ell = (-1)^\ell (I_d - P)$ for $\ell \geq 1$. Consequently, lag-one correlations are zero on $\mathrm{range}(P)$ and -1 on its complement. At each iteration the noise vector still has independent $\mathcal{N}(0, 1)$ entries, and independent injections are recovered when $P = I_d$. This guarantee concerns the noise marginal and does not imply independence from the reconstruction state.

B.3 NUMERICAL CONSTRUCTION

In order to keep the compute requirements of the algorithm low, we apply the normal operator through $H_s v = A_s^*(A_s v)$ without forming it explicitly. To estimate the scale and assess equation 19, we take a probe block $\Omega = [\omega_1, \ldots, \omega_m]$ with independent $\omega_j \sim \mathcal{N}(0, I_d)$ and compute

$$U = H_s\Omega, \qquad V = H_s U, \qquad \widehat{\lambda}_s = \frac{\mathrm{Re}\langle U, V\rangle_F}{\|U\|_F^2}. \tag{27}$$

For $U \neq 0$, this is the unique real minimizer of $\|V - \lambda U\|_F^2$: differentiating the quadratic gives $-2\,\mathrm{Re}\langle U, V\rangle_F + 2\lambda\|U\|_F^2 = 0$. If $H_s^2 = \lambda_s H_s$ exactly, then $V = \lambda_s U$, so the fitted scale is exact and the residual vanishes.

For real symmetric $H = H_s$, the same estimate can be written as

$$\widehat{\lambda}_s = \frac{\sum_{j=1}^m \omega_j^\top H^3 \omega_j}{\sum_{j=1}^m \omega_j^\top H^2 \omega_j}. \tag{28}$$

Equation equation 28 is a ratio of Gaussian trace estimators and converges to $(H^3)/(H^2)$ as the number of probes increases Avron & Toledo (2011). Under $H^2 = \lambda_s H$, this ratio equals λ_s exactly.

We assess the fitted relation using

$$\delta_s = \frac{\|V - \widehat{\lambda}_s U\|_F}{\|U\|_F}. \tag{29}$$

For a positive, numerically usable scale, a residual below the prescribed tolerance accepts $\widehat{P}_s = H_s/\widehat{\lambda}_s$ and $\widehat{R}_s = \widehat{P}_s - I_d$. Otherwise, we use $P_s = 0$ and $R_s = -I_d$. The residual provides a

randomized diagnostic of the scaled-projector relation (Halko et al., 2011, Section 4.3). The fitted residual measures agreement on the probe block. Scale estimation and screening are performed once per stage, using probes independent of the reconstruction noise. Each subsequent affine reflection requires one normal-operator product, while the fallback requires none.

Numerical accuracy. For an approximate accepted operator, given $X = H/\widehat{\lambda}$ and $\widehat{R} = X - I_d$. The update becomes $\varepsilon_{k+1} = X\eta_{k+1} + \widehat{R}\varepsilon_k$. Self-adjointness of X gives

$$XX^\top + \widehat{R}\widehat{R}^\top - I_d = 2(X^2 - X). \tag{30}$$

Writing $E_k = \mathrm{Cov}(\varepsilon_k) - I_d$, independence of the innovation yields

$$\begin{aligned}\mathrm{Cov}(\varepsilon_{k+1}) &= X^2 + (I_d - X)(I_d + E_k)(I_d - X) \\ &= I_d + 2(X^2 - X) + (I_d - X)E_k(I_d - X).\end{aligned} \tag{31}$$

Therefore

$$E_{k+1} = (I_d - X)E_k(I_d - X) + 2(X^2 - X), \qquad E_0 = 0. \tag{32}$$

In particular, $E_1 = 0$ if and only if $X^2 = X$. The exact operator models above therefore preserve unit covariance. For a fixed approximate operator, the noise remains Gaussian, with covariance error governed by equation 32.

C ITERATIONS AND RUNTIME TO MATCH THE BASELINES

For each problem, we report the smallest evaluated MPR-Flow iteration budget that reaches or exceeds each baseline's mean PSNR or SSIM at 100 iterations, this can be seen in Table 3

Table 3: Minimum MPR-Flow budgets for matching 100-iteration baselines. Times are wall-clock seconds per acquisition stage (t_{stage}) and per full evaluation (T_{run}). Reduction factors compare against the corresponding baseline. A dash denotes a target not reached by any evaluated budget up to 100 iterations.

Problem	Baseline	Metric	$N_{\min}$	t_{stage} (s)	T_{run} (s)	NFE reduction	Time reduction
Inpainting 10	PnP-Flow	PSNR	9	1.054	42.17	11.11×	10.54×
		SSIM	8	0.900	36.02	12.50×	12.34×
	Flower	PSNR	8	0.900	36.02	12.50×	12.41×
		SSIM	7	0.791	31.65	14.29×	14.12×
Inpainting 30	PnP-Flow	PSNR	8	0.898	35.92	12.50×	12.35×
		SSIM	8	0.898	35.92	12.50×	12.35×
	Flower	PSNR	8	0.898	35.92	12.50×	12.36×
		SSIM	7	0.789	31.55	14.29×	14.08×
MRI ($\sigma = 0$)	PnP-Flow	PSNR	4	0.413	66.07	25.00×	23.86×
		SSIM	–	–	–	–	–
	Flower	PSNR	4	0.413	66.07	25.00×	26.49×
		SSIM	15	1.491	238.55	6.67×	7.34×
MRI ($\sigma = 0.03$)	PnP-Flow	PSNR	3	0.316	50.52	33.33×	31.18×
		SSIM	4	0.414	66.18	25.00×	23.80×
	Flower	PSNR	3	0.316	50.52	33.33×	31.57×
		SSIM	4	0.414	66.18	25.00×	24.10×
Random matrix	PnP-Flow	PSNR	10	1.302	52.08	10.00×	9.39×
		SSIM	10	1.302	52.08	10.00×	9.39×
	Flower	PSNR	10	1.302	52.08	10.00×	13.94×
		SSIM	10	1.302	52.08	10.00×	13.94×

D AFM: EXPANDED RESULTS AND METRIC EXPLANATION

Fragment extraction We measure fragment preservation by counting connected components in the thresholded AFM/topology channel using 4-neighbor connectivity. Reconstruction and ground truth are restricted to the same DNA-region mask generated alongside the dataset, to select the regions of interest before component counting. These counts describe image components rather than physical DNA chains.

$$\text{Frag. MSE} = \frac{1}{M}\sum_{i=1}^{M} d_i^2. \tag{33}$$

Count-error metrics Since fragment-count MSE summarizes only the magnitude of the count error, we also report over-count, under-count, and exact-count rates on images containing at least one ground-truth fragment. These metrics distinguish reconstructions that split structures into too many components from those that merge or miss components, while exact-count rate measures how often the recovered topology has the correct number of fragments.

The over-count, under-count, and exact-count rates are computed per image and sum to one before rounding, so errors cannot cancel across images. Exact-count agreement alone does not guarantee topological correctness, since different structural errors can preserve the total component count.

Masked PSNR. To assess reconstruction fidelity specifically within DNA regions, we additionally report masked PSNR, computed by restricting the mean squared error to pixels within the DNA-region mask and using the same intensity range as for global PSNR. This complements global PSNR by focusing on filament regions, whose reconstruction errors may be diluted given that most of the measurement is composed of the sample's background.

N	Method	PSNR ↑	SSIM ↑	Frag. MSE ↓	Over ↓	Under ↓	Exact ↑	Masked PSNR ↑
10	Flower	33.13	0.798	4.23	**19.5%**	25.8%	54.7%	19.72 ± 1.49
	PnP-Flow	36.69	0.900	3.54	38.3%	**18.0%**	43.8%	23.59 ± 2.01
	MPR-Flow	**42.41**	**0.983**	**1.29**	21.1%	20.3%	**58.6%**	$\mathbf{28.76} \pm 3.22$
50	Flower	37.34	0.937	2.09	**14.1%**	31.3%	54.7%	23.60 ± 1.91
	PnP-Flow	40.37	0.975	2.20	29.7%	**18.8%**	51.6%	26.71 ± 2.79
	MPR-Flow	**42.75**	**0.990**	**0.92**	**14.1%**	**18.8%**	**67.2%**	$\mathbf{28.77} \pm 2.98$

Table 4: AFM reconstruction fidelity and fragment-count accuracy at 30% coverage after 10 acquisition stages. N denotes the per-stage reconstruction budget. Over, Under, and Exact are the percentages of images whose reconstructed component counts exceed, fall below, or match ground truth, respectively. Higher PSNR, SSIM, masked PSNR, and Exact are better; lower Frag. MSE, Over, and Under are better. Best results within each budget are bold; second-best distinct results are blue, with ties receiving the same formatting.

Table 4 reports masked PSNR alongside the per-image over-count, under-count, and exact-count rates. MPR-Flow achieves the highest masked PSNR at both budgets, exceeding PnP-Flow, the strongest baseline on this metric, by 5.17 dB at $N = 10$ and 2.06 dB at $N = 50$. Increasing the budget improves masked PSNR from 19.72 to 23.60 dB for Flower and from 23.59 to 26.71 dB for PnP-Flow, whereas MPR-Flow remains nearly unchanged (28.76 versus 28.77 dB). Thus, MPR-Flow already achieves strong fidelity within DNA regions at the smaller budget.

The directional count rates reveal different error tendencies. At $N = 10$, Flower under-counts more frequently than it over-counts (25.8% versus 19.5%), whereas PnP-Flow predominantly over-counts (38.3% versus 18.0%). MPR-Flow exhibits similar frequencies of over- and under-counting (21.1% and 20.3%) and the highest exact-count accuracy (58.6%). Although it does not minimize either directional error rate individually, it produces fewer images with an incorrect count overall.

E QUALITATIVE RECONSTRUCTIONS

The aggregate tables showed earlier report one number per episode, however this appendix shows one full acquisition episode per problem, stage by stage, with every row receiving identical measurements and an identical per-stage budget. Figure 6 covers random inpainting, Figure 7 partial MRI, Figure 8 AFM, and Figure 9 sequential random-matrix sensing. Per-panel PSNR is printed under each frame, so the rows can be read either as pictures or as stagewise curves.

Three things are visible in all three episodes and are not visible in the tables. The cold-started rows do not only contain a lower PSNR, but in early iterations they appear to be inconsistent, containing successive frames disagreeing about content the measurements do not yet determine, because each stage draws a fresh trajectory from the prior. The continued rows change smoothly, so an operator watching them sees refinement rather than resampling. The advantage is largest in the middle of an episode rather than at either end, early on there is little to inherit, and late on the measurements determine the answer by themselves.



Figure 6: **Random inpainting, 30% terminal sampling.** One episode of $S = 8$ acquisitions at $N = 20$ steps per stage, with $\alpha = -0.1$, $\tau_{\max} = 0.9$, $K_{\text{meas}} = 3$, and $\rho = 0.97$. The reconstruction rows compare PnP-Flow, ADWS only, and MPR-Flow (RNC + ADWS). All methods receive identical measurements and use the same per-stage reconstruction budget. PnP-Flow reconstructs independently at each stage, while ADWS only and MPR-Flow carry forward structure that subsequent measurements can refine. PSNR is reported beneath each reconstruction. MPR-Flow's PSNR advantage over PnP-Flow is largest at $s = 2$ ($+5.53$ dB). At the final stage, MPR-Flow reaches 32.02 dB, compared with 28.65 dB for PnP-Flow and 30.27 dB for ADWS only.

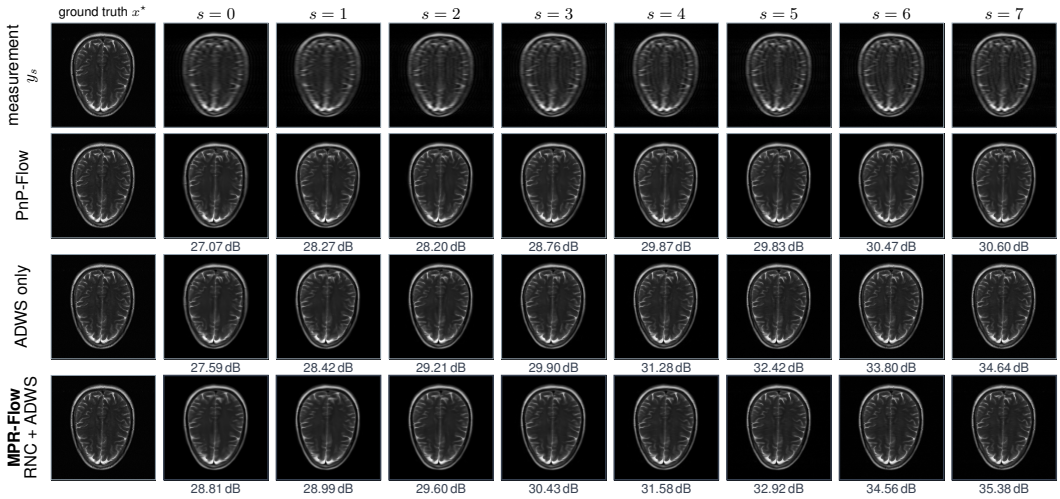

Figure 7: **Partial MRI.** One episode of $S = 8$ acquisitions on a held-out fastMRI brain slice, with eight phase-encoding lines added per stage to an 8% auto-calibration band, increasing sampling coverage from 10.2% to 24.7%. The reconstruction rows compare PnP-Flow, ADWS only, and MPR-Flow (RNC + ADWS). All methods receive identical measurements and use $N = 10$ reconstruction steps per stage, with $\alpha = -0.1$, $\tau_{\max} = 0.95$, $K_{\text{meas}} = 6$, and $\rho = 0.97$. PSNR is reported beneath each reconstruction. At the final stage, MPR-Flow reaches 35.38 dB, compared with 30.60 dB for PnP-Flow and 34.64 dB for ADWS only.

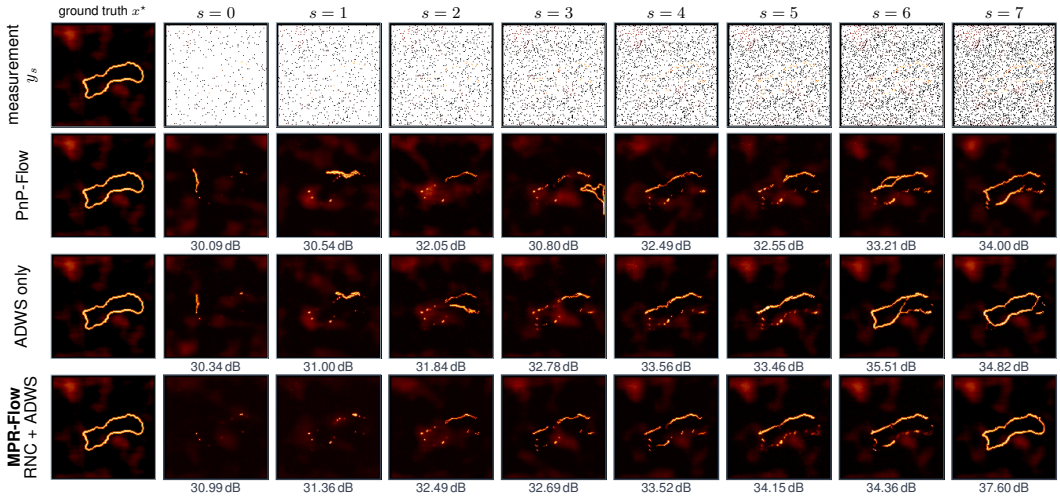

Figure 8: **Progressive reconstruction of synthetic DNA AFM images.** One episode of $S = 8$ acquisitions on a 128×128 image, with nested random pixel sampling increasing coverage from 2% to 15%. Unobserved pixels are shown in white in the measurement row. The reconstruction rows compare PnP-Flow, ADWS only, and MPR-Flow (RNC + ADWS). All methods receive identical measurements and use $N = 20$ reconstruction steps per stage, with $\alpha = -0.3$ and $\rho = 0.97$ for RNC. Global PSNR is reported beneath each reconstruction. At the final stage, MPR-Flow reaches 37.60 dB, compared with 34.00 dB for PnP-Flow and 34.82 dB for ADWS only.

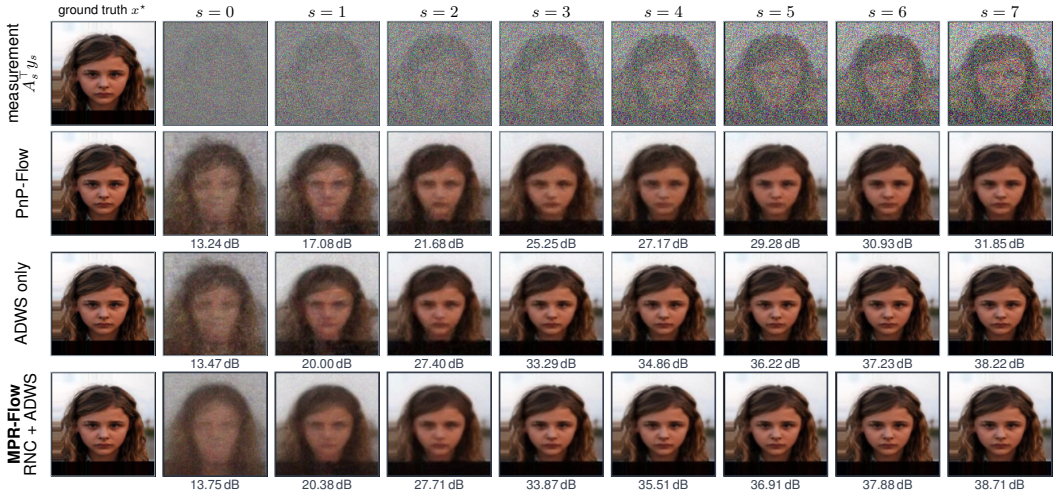

Figure 9: **Sequential random-matrix sensing.** One episode of $S = 8$ acquisitions at $N = 20$ steps per stage, $\alpha = -0.1$, $\tau_{\max} = 0.85$, $K_{\text{meas}} = 3$, and $\rho = 0.97$. The reconstruction rows compare PnP-Flow, ADWS only, and MPR-Flow (RNC + ADWS). The measurement row shows the adjoint of the accumulated observations, $A_s^\top y_s$, since the dense sensing operator does not yield a partial image. At $s = 3$, MPR-Flow improves PSNR over PnP-Flow by 8.62 dB. The improvement is consistent with successive Gaussian measurement rows constraining previously undetermined directions, while ADWS carries forward information recovered at earlier stages.

F HYPERPARAMETER SETTINGS AND SENSITIVITY

F.1 HYPERPARAMETER SETTINGS

In this section we report the settings used for $N = 10$ and $N = 50$ flow-network evaluations per acquisition stage (See Table 5). Here α controls the data-consistency step, $\tau_{\max}$ is the maximum starting flow time, and K_{meas} controls the acquisition-stage ramp.

	Inpainting 10	Inpainting 30	MRI $\sigma = 0$	MRI $\sigma = 0.03$	Random-Matrix
			$N = 10$		
FLOWER					
σ_n (assumed noise level)	10^{-3}	10^{-3}	10^{-3}	10^{-3}	10^{-3}
OT-ODE					
t (initial time)	0.2	0.2	0.2	0.2	0.2
γ	constant	constant	constant	constant	constant
Flow-Priors					
λ (regularization)	10000	10000	1000	1000	100
PnP-Flow and PnP-Flow + RNC					
α (step exponent)	-0.30	-0.20	-0.10	-0.10	-0.35
Continuation and MPR-Flow					
α (step exponent)	-0.30	-0.20	-0.10	-0.10	-0.35
$\tau_{\max}$ (maximum start time)	0.80	0.80	0.95	0.95	0.80
K_{meas} (ramp length)	9	3	6	6	3
			$N = 50$		
FLOWER					
σ_n (assumed noise level)	10^{-3}	10^{-3}	10^{-3}	10^{-3}	10^{-3}
OT-ODE					
t (initial time)	0.2	0.2	0.2	0.2	0.2
γ	constant	constant	constant	constant	constant
Flow-Priors					
λ (regularization)	10000	10000	1000	1000	100
PnP-Flow					
α (step exponent)	-0.20	-0.15	-0.10	-0.10	-0.15
MPR-Flow					
α (step exponent)	-0.20	-0.15	-0.10	-0.10	-0.15
$\tau_{\max}$ (maximum start time)	0.70	0.75	0.90	0.90	0.90
K_{meas} (ramp length)	3	3	1	3	3

Table 5: Hyperparameter settings for $N = 10$ and $N = 50$ flow-network evaluations per acquisition stage.

F.2 HYPERPARAMETER SENSITIVITY

Figure 10 examines the joint effect of the maximum starting flow time $\tau_{\max}$ and the data-consistency exponent α at a fixed budget of $N = 10$ iterations per acquisition stage. values of α closer to zero allow larger starting times while maintaining stable reconstruction. More negative α can also perform well, but become unstable at large starting times. Divergent runs are omitted from the figure, accounting for the truncated curves.

Across the evaluated settings, several configurations yield similar reconstruction quality. Inpainting and random-matrix sensing generally favor $\tau_{\max} = 0.80$–0.85, while MRI benefits from larger starting times when α is closer to zero.

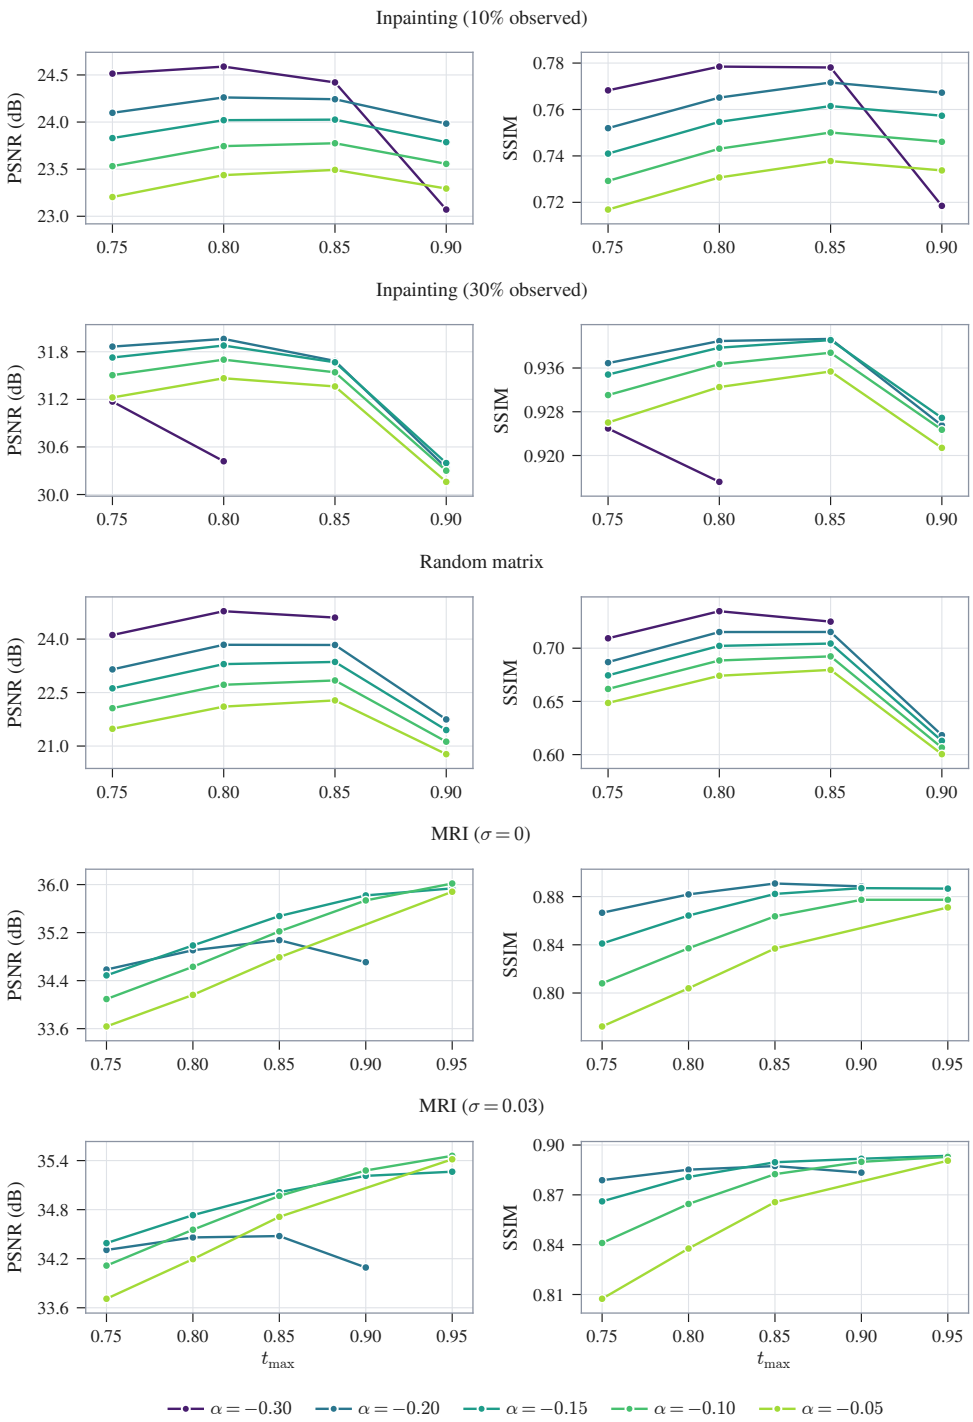

Figure 10: Parameter sensitivity at $N = 10$ iterations per acquisition stage. PSNR (left) and SSIM (right) are plotted against the maximum starting flow time $\tau_{\max}$, with separate curves for each data-consistency exponent α.

G Reconstruction Quality per Acquisition Stage

Reconstruction quality at the final acquisition stage captures only part of the performance of a measurement-progressive method. During the imaging process, the intermediate reconstructions' quality can also be helpful in assesing whether more measurements are needed.

This appendix explores the reconstruction quality for each stage, in addition we report the first stage at which MPR-Flow matches the best baseline's final PSNR, showing potential to achieve comparable quality with fewer acquisitions, shown in 6.

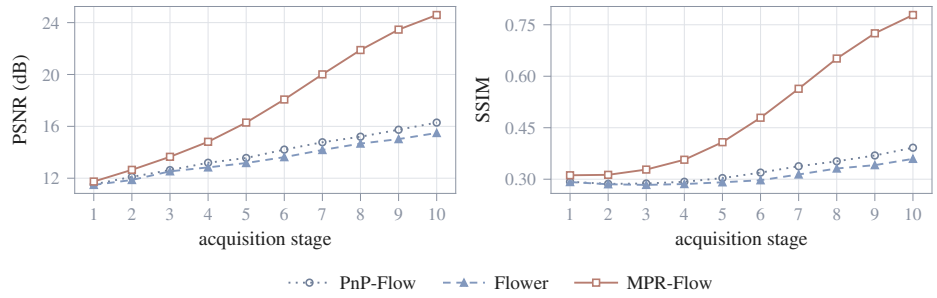

Figure 11: Reconstruction quality per acquisition stage for 10% inpainting at $N = 10$: PSNR (left) and SSIM (right), averaged over 128 test images. MPR-Flow's advantage increases throughout acquisition, and its PSNR reaches the best baseline's final PSNR at stage 5 of 10.

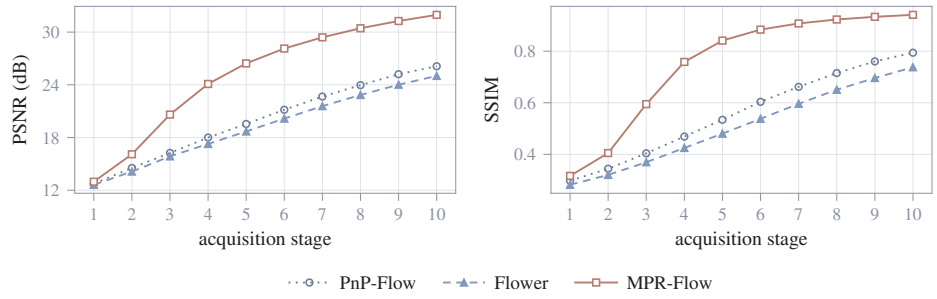

Figure 12: Reconstruction quality per acquisition stage for 30% inpainting at $N = 10$: PSNR (left) and SSIM (right), averaged over 128 test images. The quality gap widens as measurements accumulate, with MPR-Flow reaching the best baseline's final PSNR at stage 5 of 10.

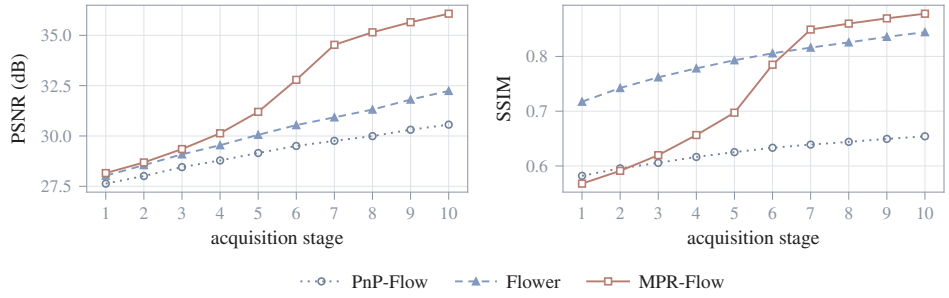

Figure 13: Reconstruction quality per acquisition stage for noiseless partial Cartesian MRI ($\sigma = 0$) at $N = 10$: PSNR (left) and SSIM (right), averaged over 128 test images. MPR-Flow reaches the best baseline's final PSNR at stage 6 of 10. FLOWER achieves higher SSIM during the early stages, before MPR-Flow overtakes it at stage 7.

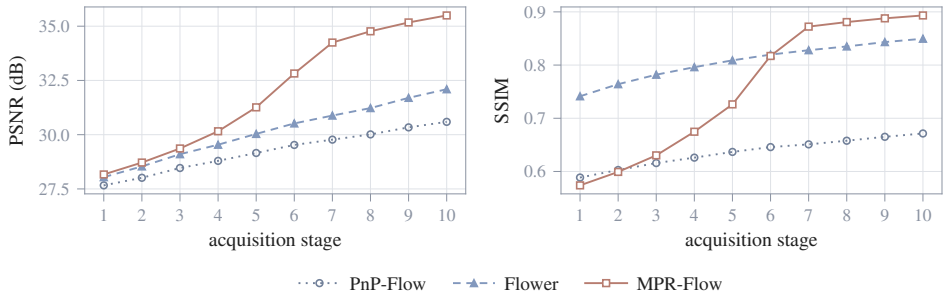

Figure 14: Reconstruction quality per acquisition stage for noisy partial Cartesian MRI ($\sigma = 0.03$) at $N = 10$: PSNR (left) and SSIM (right), averaged over 128 test images. MPR-Flow reaches the best baseline's final PSNR at stage 6 of 10 and achieves the highest final-stage PSNR and SSIM.

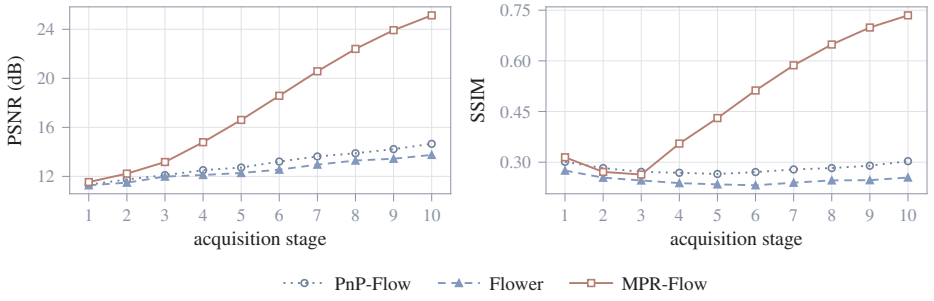

Figure 15: Reconstruction quality per acquisition stage for random-matrix sensing at $N = 10$: PSNR (left) and SSIM (right), averaged over 128 test images. MPR-Flow reaches the best baseline's final PSNR at stage 4 of 10, and its final PSNR advantage is approximately 10.5 dB. Its SSIM briefly falls below PnP-Flow at stages 2–3 before increasing sharply.

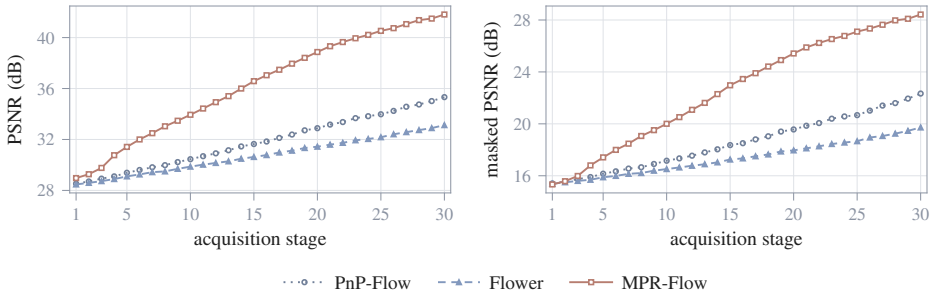

Figure 16: Reconstruction quality per acquisition stage for synthetic DNA AFM with 30% sampling at $N = 10$: global PSNR (left) and PSNR restricted to the DNA mask (right), averaged over 128 test images. MPR-Flow reaches the best baseline's final global PSNR at stage 13 of 30 and achieves a final global PSNR advantage of approximately 6.5 dB.

Table 6: First acquisition stage at which MPR-Flow reaches or exceeds the best baseline's final PSNR, with $N = 10$ flow-network evaluations per stage.

Problem	Stage to match	Total stages
Inpainting 10%	5	10
Inpainting 30%	5	10
MRI, $\sigma = 0$	6	10
MRI, $\sigma = 0.03$	6	10
Random-matrix sensing	4	10
DNA AFM 30%	13	30

H INTERACTION BETWEEN ADWS AND RNC

To better understand the interaction between ADWS and RNC, we report reconstruction quality across maximum warm-start times for both the complete method and its continuation-only variant. Although later starts preserve more of each preceding reconstruction, they also restrict the scope of the remaining updates, creating a trade-off between retaining useful structure and correcting inherited errors caused by the lack of information at the previous measurement stage. The increasing benefit of RNC seen in Figure 17 and Figure 18 might show that as this corrective capacity decreases, RNC may help the solver retain useful information while refining the inherited estimate.

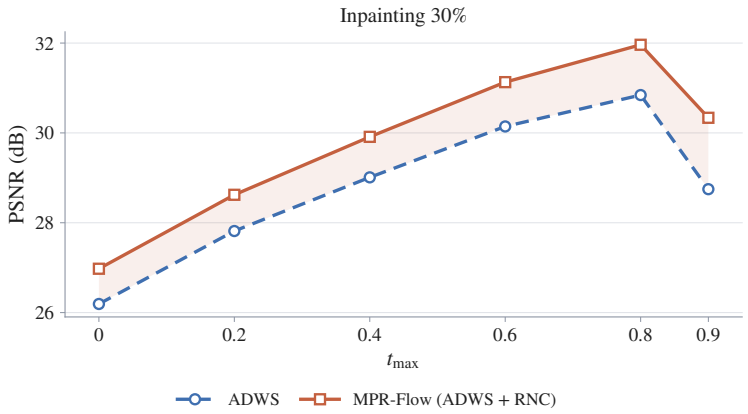

Figure 17: PSNR vs t_{max} for Inpainting 30% with N=10. MPR-Flow shows to be better at every t_{max}, and increases along with t_{max} even in the region where reconstruction quality begins to drop.

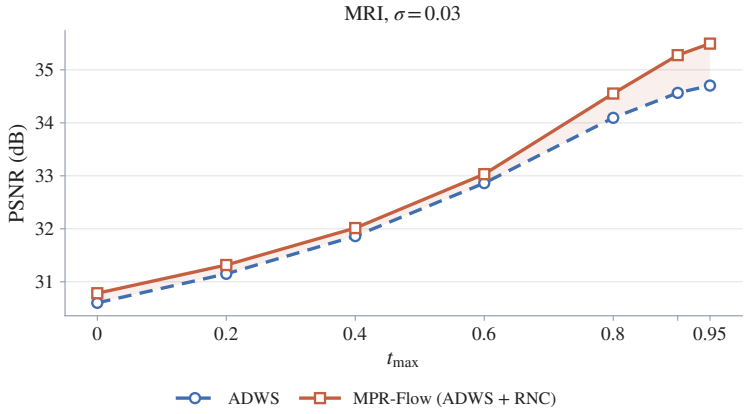

Figure 18: t_{max} for MRI data ($\sigma = 0.03$) with N=10. RNC appears to provide a modest improvement at lower values of t_{max} and show a steeper increase in PSNR at higher t_{max} values.