ADAPTIVE DUAL VELOCITY FIELD LEARNING ALONG RECONSTRUCTION TRAJECTORIES FOR INSUFFICIENT-ANGLE CT

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models and flow-based generative models, most notably flow matching (FM), have recently delivered strong performance in computed tomography (CT) reconstruction. However, for the same scanning system whose acquisition geometry varies in the number of projection angles or in the angular range, existing methods still require retraining a dedicated model for each configuration and thus offer limited cross-geometry adaptability. Moreover, we observe that the quantity learned during intermediate sampling—the noise predictor ϵ_θ or score function in diffusion or the velocity field in flow models—is essentially independent of the current state along the sampling trajectory: in linear FM, for instance, the conditional velocity is a constant determined only by the two endpoints, $\mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0$. Such a position-independent field is fundamentally implausible for the intrinsically nonlinear transport from projection data to images. In this paper, we therefore organize reconstructions obtained under different numbers of projections as an ordered reconstruction trajectory, and learn along this trajectory a dual velocity field consisting of a current (geometry-transition) velocity that describes the local evolution between successive angle settings and a target velocity that points toward the clean image. An adaptive velocity-learning mechanism further conditions the transport on the current state and its position along the trajectory, and the two velocities drive a progressive two-stage reconstruction network. A single trained model thereby generalizes across sparse- and limited-angle settings without retraining, achieving high-precision quantitative imaging on LoDoInd, Mayo, and PCB benchmarks.

Keywords: Computed tomography reconstruction; dual velocity field; adaptive velocity learning; quantitative imaging

1 INTRODUCTION

Computed tomography (CT) reconstructs a cross-sectional image $\mathbf{x} \in \mathbb{R}^n$ from its projection measurements (sinogram) $\mathbf{y} \in \mathbb{R}^m$ acquired along different view angles, under the linear forward model $\mathbf{y} = \mathbf{A}_g\mathbf{x} + \mathbf{n}$, where $\mathbf{A}_g$ denotes the geometry-dependent projection operator and $\mathbf{n}$ the measurement noise (Kak & Slaney, 2001). In practice, radiation dose (Brenner & Hall, 2007) and acquisition time impose strong constraints, giving rise to two canonical ill-posed settings: *sparse-angle* (SA) tomography, which collects only a few projections over the full 360° range, and *limited-angle* (LA) tomography, which restricts projections to a contiguous angular interval. Direct inversion by filtered back-projection (FBP) (Kak & Slaney, 2001; Natterer, 2001) then produces images contaminated by severe streak and missing-wedge artifacts, and classical remedies—algebraic reconstruction (Gordon et al., 1970), sparsity-driven iterative thresholding (Daubechies et al., 2004; Beck & Teboulle, 2009), and total-variation minimization (Rudin et al., 1992; Sidky & Pan, 2008)—are geometry-agnostic but still leave residual artifacts or over-smoothing. Although SA and LA are conventionally studied in isolation, they are two manifestations of one and the same deficiency—*insufficient angular sampling* of the Radon domain: SA lacks angular *density*, whereas LA lacks angular *coverage*. We therefore unify them into a single problem, *insufficient-angle CT* (IA-CT), in which the geometry descriptor g—the number of projection angles, or the angular range—is a *free variable* of

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 very same scanning system, and a reconstruction model is expected to remain accurate for *every* g.

Learning-based reconstruction has evolved through three generations. The first generation is *direct mapping*: a purely feed-forward regressor refines a degraded FBP estimate, from RED-CNN (Chen et al., 2017) and FBPConvNet (Jin et al., 2017) to wavelet-frame U-Nets (Han & Ye, 2018; Kang et al., 2018), adversarial refinement (Goodfellow et al., 2014; Yang et al., 2018; Huang et al., 2022), and implicit adversarial interpolation across the missing wedge (Anirudh et al., 2018). The second generation *unrolls* T iterations of an optimization scheme—proximal gradient (Gregor & LeCun, 2010; Zhang & Ghanem, 2018), ADMM (Yang et al., 2016), or primal–dual updates (Adler & Oktem, 2018; Hammernik et al., 2018; Schlemper et al., 2018; Monga et al., 2021)—so that the regularizer and the step sizes are learned while data-consistency layers built on $\mathbf{A}_g^\top$ are retained, making the reconstruction more stable and interpretable than pure post-processing. The third generation replaces deterministic regression with iterative *sampling* from a learned generative prior: score-based diffusion (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b; Dhariwal & Nichol, 2021) with measurement guidance (Ho & Salimans, 2022; Kawar et al., 2022; Chung et al., 2023; 2022a; Song et al., 2023a; Wang et al., 2024; Mardani et al., 2023; Cardoso et al., 2024; Dou & Song, 2024) and ODE-based flow matching (FM) (Lipman et al., 2023; Liu et al., 2023; Tong et al., 2024; Albergo et al., 2023), implemented as neural ODEs (Chen et al., 2018). This family has rapidly become the dominant paradigm, and it is most convincing in *quantitative low-dose denoising*, where the forward operator is fixed and the degradation is approximately Gaussian (Chung et al., 2022b; Gao et al., 2024; Xie & Li, 2022), with a rapidly growing 2025–2026 literature extending diffusion-based and learned priors to LA/SA CT (Guo & López-Tapia, 2025; Li et al., 2026; Hahn et al., 2026).

For IA-CT, however, two gaps separate this paradigm from quantitative success, as illustrated in Figure 1. **(i) Geometry-bound retraining.** All three generations are designed and trained against one *fixed* operator $\mathbf{A}_g$: the learned mapping or field is frozen to a fixed number of angles or a fixed angular range, so supporting every configuration of a modern scanner requires one separately trained network per configuration—impractical in clinical and industrial settings where scan protocols vary across patients and instruments. **(ii) Position-independent fields versus nonlinear transport.** The quantity learned during intermediate sampling—the noise predictor ϵ_θ or score in diffusion, the velocity field in FM—is essentially *independent of the current state* along the trajectory: in linear FM the conditional target velocity is the constant vector $\mathbf{x}_1 - \mathbf{x}_0$ determined only by the two endpoints. The learned field is therefore nearly *parallel*: applied to an off-path state $\tilde{x}_t$, parallel transport preserves the off-path displacement and terminates at $\hat{x}_1$, leaving a large gap from the true endpoint (Figure 1a). IA-CT makes this failure concrete. Each angular view carries partial image information, and as the angular sampling becomes sufficient, the reconstruction gradually converges to the exact solution along a strongly nonlinear trajectory; when the sampling is insufficient, the trajectory simply stops midway (Figures 1b,c). Yet existing deep methods, given a fixed geometry, learn either a *direct mapping to the endpoint* or a *position-independent velocity field*, thereby ignoring this nonlinear evolution entirely—which yields quantitatively biased reconstructions and hallucinated structures inside the missing wedge.

This observation leads to our key insight: rather than jumping from a single insufficient-angle reconstruction straight to the endpoint, the model should *keep learning along the stopped trajectory*—preserving the nonlinear direction of the trajectory while injecting a velocity directed at the target. We therefore reformulate IA-CT reconstruction as learning along a *reconstruction trajectory*: reconstructions obtained under different angle configurations are ordered by geometry difficulty into a sequence $x^{(0)}, \ldots, x^{(I-1)}$ of I geometry levels (with $I = 10$ in this paper), and along this trajectory we jointly learn a *dual velocity field*—a *target velocity* $\mathbf{v}_i^{\mathrm{tar}}$ that points from the current state toward the clean image, and a *local (geometry-transition) velocity* $\mathbf{v}_i^{\mathrm{loc}}$ that describes the evolution between successive angle settings. An *adaptive velocity-learning* scheme conditions the transport on the current state and its position along the trajectory, so that the field bends with the state instead of following a straight line, resolving the position-independence criticized above. The idea is realized by a progressive two-stage network (Figure 2), referred to as *ADVF* (**A**daptive **D**ual **V**elocity **F**ield learning): in the first stage, a velocity network F_θ predicts both velocities from the current reconstruction and its geometry index; in the second stage, lightweight transition modules recursively complete the missing geometry levels and a reconstruction network G_ϕ aggregates the completed sequence to

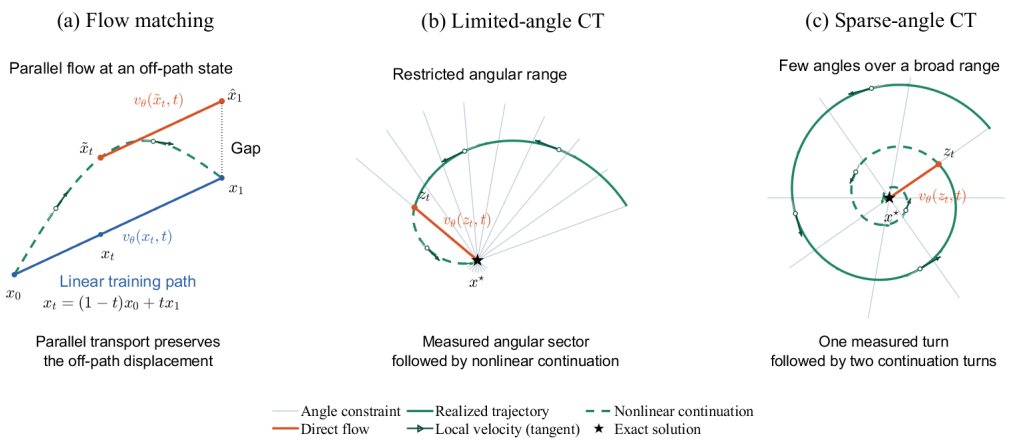


Figure 1: Motivation of the proposed method. (a) Flow matching learns the velocity field $v_\theta(x_t, t)$ along the linear training path $x_t = (1-t)x_0 + tx_1$; due to this linearity, the learned transport is essentially parallel, so an off-path state $\tilde{x}_t$ carried by the same field terminates at $\hat{x}_1$, leaving a large gap from the true endpoint x_1. (b) Limited-angle and (c) sparse-angle CT: each view carries partial image information, and the reconstruction gradually converges to the exact solution $x^\star$ along a nonlinear trajectory as the angular sampling becomes sufficient; when the sampling is insufficient, the trajectory stops midway at z_t. Existing deep models learn, for a fixed geometry, a direct mapping to the endpoint or a position-independent velocity field, ignoring this nonlinear continuation (dashed). We instead keep learning along the dashed trajectory by jointly preserving its nonlinear direction and a target-directed velocity.

estimate the final clean image. A single trained model thereby serves every angle configuration of the same scanning system without retraining.

We instantiate F_θ and G_ϕ with U-Net (Ronneberger et al., 2015) backbones and evaluate on LoDoInd, Mayo (McCollough et al., 2017), and PCB benchmarks under multiple SA and LA settings, comparing against FBP, a direct-mapping baseline (DLPVI (Zhao et al., 2025)), an unrolled baseline (Shi et al., 2026), and flow-based baselines (linear FM and PnP-Flow (Martin et al., 2025)) in terms of PSNR and SSIM (Section 4). Our contributions are summarized as follows:

- We unify sparse-angle and limited-angle reconstruction as insufficient-angle CT (IA-CT), and cast multi-geometry reconstruction as learning along an ordered reconstruction trajectory.
- We reveal the position-independence of the fields learned by diffusion and flow models and its fundamental mismatch with the nonlinear artifact-evolution transport of IA-CT, where insufficient angular sampling stops the reconstruction trajectory midway while existing models learn only direct endpoint mappings or position-independent velocity fields.
- We propose **ADVF**, which keeps learning along the nonlinear trajectory by jointly learning a dual velocity field—a trajectory-direction-preserving local velocity and a target-directed velocity—via adaptive velocity learning, realized as a progressive two-stage network; one model thereby serves all angle configurations without retraining.
- Across LoDoInd, Mayo, and PCB, ADVF achieves state-of-the-art accuracy under both SA and LA settings and generalizes to unseen geometries with high-precision quantitative imaging.

2 Related Work

2.1 CT forward model and reconstruction methods for insufficient-angle CT

Let $\mathbf{x} \in \mathbb{R}^n$ be the (column-vectorized) image and $\mathcal{G}$ the acquisition geometry, characterized by a descriptor $g \in \mathcal{G}$: for sparse-angle CT, g is the number of projection angles; for limited-angle CT, g

is the angular range (or the ending angle). The corresponding forward model is

$$\mathbf{y} = \mathbf{A}_g \, \mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{A}_g \in \mathbb{R}^{m \times n}$ is the geometry-parameterized projection operator and $\mathbf{n}$ is noise. We denote by $\tilde{\mathbf{x}} = \mathcal{B}_g \mathbf{y}$ a degraded estimate produced by a (possibly approximate) back-projection, e.g., an FBP reconstruction. The goal is to learn, from paired data $\{(\mathbf{y}^{(i)}, \mathbf{x}_1^{(i)})\}$ sampled across *multiple* geometries g, a reconstruction operator that is accurate for every g using a single model.

Classical reconstruction relies on analytical algorithms such as FBP (Kak & Slaney, 2001; Natterer, 2001) and on regularized iterative methods—algebraic reconstruction (Gordon et al., 1970), sparsity-driven iterative thresholding (Daubechies et al., 2004; Beck & Teboulle, 2009), and total-variation minimization (Rudin et al., 1992; Sidky & Pan, 2008)—which are geometry-agnostic but leave significant residual artifacts or over-smoothing under SA/LA acquisition. Learning-based methods fall into three families: (i) *direct-mapping* networks that refine an FBP estimate (Chen et al., 2017; Jin et al., 2017; Han & Ye, 2018; Kang et al., 2018; Yang et al., 2018; Huang et al., 2022; Anirudh et al., 2018); (ii) *unrolled* networks that unfold a variational iteration and learn its parameters (Gregor & LeCun, 2010; Yang et al., 2016; Adler & Oktem, 2018; Zhang & Ghanem, 2018; Hammernik et al., 2018; Schlemper et al., 2018; Monga et al., 2021); and (iii) *generative* methods that learn the image prior (Bora et al., 2018; Jalal et al., 2021; Ho et al., 2020; Song et al., 2021b; Lipman et al., 2023); see (Ongie et al., 2020) for a survey. Very recent work continues to refine SA/LA reconstruction—tri-domain transformers (Li et al., 2025; Bai et al., 2026), unsupervised deep Radon priors (Xu et al., 2025; Fu et al., 2025), deep image priors for limited-angle tomography (Hori et al., 2025), analytic regularization theory (Hahn et al., 2026), and diffusion-based solvers (Guo & López-Tapia, 2025; Li et al., 2026). Across all families and these recent advances, however, one property is shared: the network is trained against *one fixed* geometry—a fixed number of views or a fixed angular range—so a new scan protocol demands retraining, and none of them exploits the ordered family of reconstructions that a single scanner naturally produces across geometries. This is the precise gap targeted by our progressive multi-geometry framework.

2.2 Flow matching and inverse-problem solvers

Flow matching (FM) (Lipman et al., 2023), rectified flow (Liu et al., 2023), optimal-transport variants (Tong et al., 2024), and the unifying view of stochastic interpolants (Albergo et al., 2023) provide a simulation-free, deterministic ODE alternative to diffusion modeling (Chen et al., 2018): a network regresses the velocity field of a probability path that transports a source distribution p_0 (typically $\mathcal{N}(\mathbf{0}, \mathbf{I})$) to the data distribution p_1. Given a data sample $\mathbf{x}_1 \sim p_1$ and a noise sample $\mathbf{x}_0 \sim p_0$, the standard conditional flow defines the linear interpolation

$$\mathbf{x}_t = (1 - t)\, \mathbf{x}_0 + t\, \mathbf{x}_1, \quad t \in [0, 1], \tag{2}$$

whose exact conditional velocity (the *target velocity*) is the *constant* vector

$$\mathbf{u}(\mathbf{x}_t \mid \mathbf{x}_1) \;=\; \frac{\mathrm{d}\mathbf{x}_t}{\mathrm{d}t} \;=\; \mathbf{x}_1 - \mathbf{x}_0. \tag{3}$$

A network $v_\theta(\mathbf{x}_t, t)$ is trained by the simulation-free objective

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_1}\Big[\big\| v_\theta(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0) \big\|^2 \Big], \tag{4}$$

and new samples are drawn by integrating the ODE $\mathrm{d}\mathbf{x}/\mathrm{d}t = v_\theta(\mathbf{x}_t, t)$ from $t = 0$ to $t = 1$. Crucially, the path in Eq. equation 2 is a straight line and the target velocity in Eq. equation 3 is constant, so the transport is *linear* and does not adapt its curvature to the geometry g.

Owing to its straight conditional paths and simple regression objective, FM has been rapidly adopted for inverse problems, where the velocity is additionally conditioned on a degraded observation or the measurement model, giving rise to measurement-guided flow solvers, plug-and-play designs that embed pretrained denoisers into the FM sampler, and conditional FM for medical reconstruction. Fast samplers, distillation, and consistency training (Song et al., 2021a; Lu et al., 2022; Salimans & Ho, 2022; Karras et al., 2022; Song et al., 2023b) further reduce the integration cost. Nevertheless, existing FM formulations operate on a *single* input state: their straight-line, constant target velocity is independent of the current state (see equation 3), a property whose consequences for IA-CT we scrutinize in Section 2.3.

2.3 THE NONLINEARITY GAP: ILL-POSEDNESS OF POSITION-INDEPENDENT VELOCITY FIELDS

We now analyze the obstacle that arises when position-independent fields meet IA-CT. In linear FM, the target velocity equation 3 is constant along the entire path, and from the regression objective equation 4 a deterministic network can identify, at best, a conditional expectation of the endpoint displacement given the current state and time. Two premises make this benign in denoising: the underlying transport is (nearly) linear, and test states lie on the training paths. IA-CT violates both. First, the transport from an insufficient-angle reconstruction to the clean image is a strongly *nonlinear artifact-evolution* process: each angular view carries partial image information, and the reconstruction converges to the exact solution along a curved, view-dependent trajectory that stops midway when the sampling is insufficient (Figures 1b,c). Second, realistic inference states lie *off* the linear training paths. For such an off-path state $\tilde{x}_t$, the learned field outputs a velocity nearly parallel to the endpoint displacement, so integration degenerates into parallel transport that preserves the off-path displacement: the terminal gap $\|\hat{x}_1 - x_1\|$ is of the same order as the off-path distance $\|\tilde{x}_t - x_t\|$ (Figure 1a). Worse, trajectories induced by different geometries pass through visually similar intermediate reconstructions while requiring different continuations, so a single position-independent field faces a one-to-many regression target—an ill-posed learning problem—whereas a network trained for one fixed geometry simply learns the direct endpoint map $x_g \mapsto x_1$ and discards the nonlinear evolution altogether. Both routes ignore the trajectory structure that IA-CT naturally provides. This analysis motivates a field that is *state- and position-dependent*: we keep learning along the nonlinear continuation by jointly regressing the local transition between successive geometries and the target-directed velocity, which is precisely the dual velocity field of ADVF.

3 METHOD

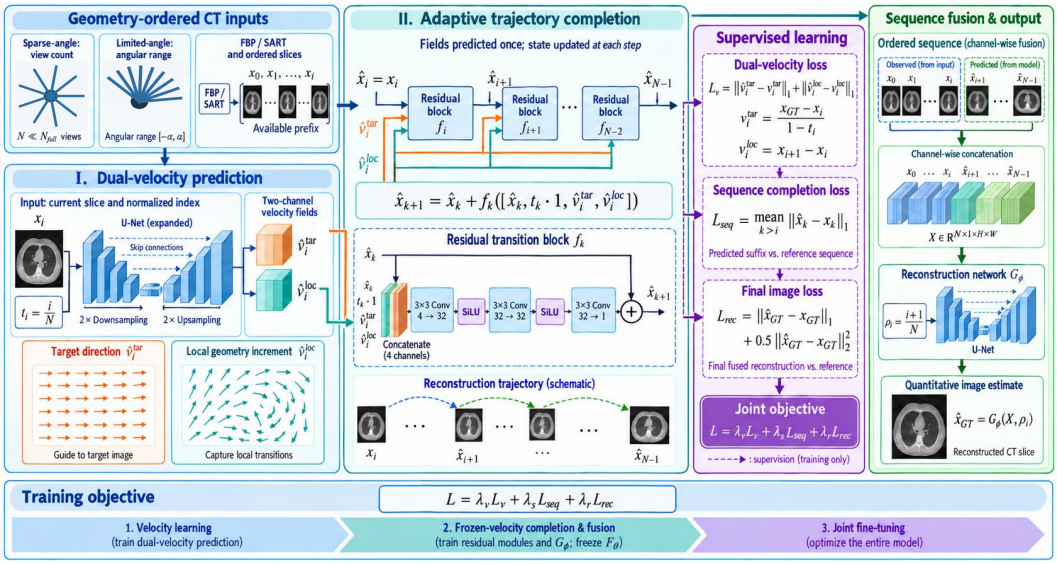


Figure 2: Overview of ADVF. **Stage I** predicts, from the current slice and its normalized geometry index, a target velocity toward the clean image and a local geometry-transition velocity toward the next geometry level. **Stage II** recursively completes the missing geometry levels with lightweight transition modules and aggregates the completed sequence with the reconstruction network G_ϕ to produce the final quantitative reconstruction.

3.1 OVERVIEW

Let x_1 denote the ground-truth image. We discretize the acquisition geometry into I steps, with normalized geometry times $t_i = i/I$ spanning $[0, 1]$ for $i = 0, \ldots, I$; we set $I = 10$ in this paper. Let $x^{(i)}$ denote the analytical reconstruction at the i-th geometry level, ordered by geometry difficulty,

with $x^{(I)} = x_1$ the clean image. For LA CT, $x^{(0)}$ corresponds to the smallest angular range and $x^{(I-1)}$ to the largest range; for SA CT, $x^{(0)}$ corresponds to the smallest number of views and $x^{(I-1)}$ to the largest. Each $x^{(i)}$ is produced by reconstructing from one specific acquisition geometry. The goal is to learn a single model that maps any partially observed geometry sequence $\{x^{(0)}, \ldots, x^{(i)}\}$ to the clean image x_1.

Unlike diffusion and flow models, whose position-independent field cannot follow the nonlinear artifact evolution of insufficient-angle CT, ADVF learns *along* the reconstruction trajectory. It is a progressive two-stage network driven by an *adaptive dual velocity field* conditioned on the current state and on its position along the trajectory, so that one trained model serves every angle configuration of the same scanning system without retraining. As illustrated in Figure 2, the pipeline (i) constructs the ordered trajectory from the geometry levels, (ii) predicts a target velocity toward the clean image and a local geometry-transition velocity between neighbouring levels with the velocity network F_θ (Stage I), (iii) rolls out lightweight transition modules $\{f_k\}$ with these two velocities to complete the missing levels and aggregates the completed sequence with the reconstruction network G_ϕ (Stage II), and (iv) trains the three components with a three-phase curriculum, detailed in Section 3.4 and Appendix A.1.

3.2 Stage I: geometry-aware velocity prediction

The first stage models how the reconstruction evolves as the geometry improves. For a randomly sampled index $i \in \{0, \ldots, I-1\}$, we feed the current reconstruction $x^{(i)}$ and the normalized geometry time $t_i = i/I$ into a velocity network F_θ. The network predicts two velocity fields:

$$v_i^{(1)} = \frac{x_1 - x^{(i)}}{1 - t_i}, \tag{5}$$

$$v_i^{(2)} = x^{(i+1)} - x^{(i)}. \tag{6}$$

The first term is a target velocity that points from the current geometry level toward the clean image. The second term is a local transition velocity that captures the change between two neighboring geometry levels. We regress both targets with a shared backbone and two output channels:

$$[\hat{v}_i^{(1)}, \hat{v}_i^{(2)}] = F_\theta(x^{(i)}, t_i), \tag{7}$$

and optimize

$$\mathcal{L}_{\text{vel}} = \left\| \hat{v}_i^{(1)} - v_i^{(1)} \right\|_1 + \left\| \hat{v}_i^{(2)} - v_i^{(2)} \right\|_1. \tag{8}$$

3.3 Stage II: progressive sequence completion and final reconstruction

At geometry index i, only the prefix $\{x^{(0)}, \ldots, x^{(i)}\}$ is directly available. The suffix $\{x^{(i+1)}, \ldots, x^{(I-1)}\}$ is missing and must be synthesized. We therefore introduce a set of lightweight transition modules $\{f_0, \ldots, f_{I-1}\}$, each implemented as a shallow convolutional block. Starting from the current state $x^{(i)}$, we recursively generate the missing geometry levels:

$$x'^{(k+1)} = f_k(x'^{(k)}, t_k, \hat{v}_i^{(1)}, \hat{v}_i^{(2)}), \quad k = i, \ldots, I-2, \tag{9}$$

where $x'^{(i)} = x^{(i)}$. The completed sequence is

$$\tilde{X}_i = \{x^{(0)}, \ldots, x^{(i)}, x'^{(i+1)}, \ldots, x'^{(I-1)}\}. \tag{10}$$

We then concatenate the I geometry levels along the channel dimension and feed them into a reconstruction network G_ϕ:

$$\hat{x}_1 = G_\phi(\tilde{X}_i, t_i). \tag{11}$$

This second stage explicitly aggregates observed low-level geometry evidence and predicted high-level geometry evidence, instead of relying on a single reconstruction as input.

3.4 TRAINING OBJECTIVE AND CURRICULUM

Besides the dual-velocity loss $\mathcal{L}_{\text{vel}}$ defined in Stage I, the second stage is trained with two losses. The first enforces consistency between the synthesized suffix and the corresponding ground-truth geometry sequence,

$$\mathcal{L}_{\text{roll}} = \sum_{k=i+1}^{I-1} \left\| x'^{(k)} - x^{(k)} \right\|_1 . \tag{12}$$

The second supervises the final clean reconstruction,

$$\mathcal{L}_{\text{rec}} = \|\hat{x}_1 - x_1\|_1 + \frac{1}{2} \|\hat{x}_1 - x_1\|_2^2 . \tag{13}$$

The full objective of the second stage is

$$\mathcal{L} = \lambda_{\text{roll}} \mathcal{L}_{\text{roll}} + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}}, \tag{14}$$

which does not involve the velocity loss $\mathcal{L}_{\text{vel}}$, used only to train the velocity network F_θ in Stage I. The three-phase training curriculum that schedules the optimization of F_θ, $\{f_k\}$, and G_ϕ, as well as the complete training and inference procedures (Algorithms 1 and 2), are given in the Appendix.

4 EXPERIMENTS

4.1 DATASETS

We evaluate on three IA-CT benchmarks. **LoDoInd** (Shi et al., 2024) is a public benchmark that provides experimental (not simulated) low-dose industrial CT images of a test object at five noise levels together with a reference scan[1]; we take the reference images as ground truth, and the images used in our experiments are of size 640×640. The benchmark contains 4000 slices, of which the first 3000 are used for training and the remaining 1000 are held out. **Mayo** is derived from the AAPM Low-Dose CT Grand Challenge (McCollough et al., 2017) and provides 512×512 slices; following the same protocol, the first 3000 slices form the training set and all remaining slices are held out. **PCB** is built from real printed circuit boards (PCBs) scanned in our laboratory and contains 2450 slices of size 256×1024, split into 1960 training and 490 testing slices. Its construction—the SART reconstruction of the acquired volumes and the re-simulation required by the mismatch between the scanner geometry (SOD/ODD) and the fan-beam geometry assumed by our simulation—is detailed in Appendix A.4.

For every benchmark, each ground-truth slice is converted into a reconstruction trajectory by simulating fan-beam projections with the native detector and source–object–detector geometry and reconstructing subsets of the views analytically; these multi-geometry reconstructions are the network inputs and the corresponding clean slice is the prediction target. All quantitative results are reported in terms of PSNR and SSIM. The full acquisition and simulation parameters—detector configuration, SOD/SDD, pixel size, and the angular ranges and view counts of all geometry levels—are given in Appendix Table 2.

4.2 COMPARISON METHODS

We compare ADVF against three representative families of methods. **(i) Direct-mapping methods:** DLPVI (Zhao et al., 2025), which integrates projection, view-by-view backprojection, and image domains into a single deep framework for high- and ultra-sparse-view CBCT reconstruction. **(ii) Unrolled methods:** a prompting Lipschitz-constrained network (Shi et al., 2026), which learns a single multiple-in-one model for sparse-view CT reconstruction. **(iii) Flow-based methods:** PnP-Flow (Martin et al., 2025), which couples pretrained denoisers with flow matching for plug-and-play image restoration.

Our method (ADVF) keeps the same two-stage architecture and instantiates both F_θ and G_ϕ with U-Net (Ronneberger et al., 2015) backbones. To evaluate our method at different model capacities, we report two variants that differ only in the channel width of the backbone: a *tiny* variant with $\{32, 64, 128\}$ channels and 1.76M parameters, and a *base* variant with $\{64, 128, 256\}$ channels and 6.70M parameters, i.e., about $3.8\times$ more parameters than the tiny variant.

[1] https://github.com/jiayangshi/LoDoInd

4.3 RESULTS

Figure 3 compares the reconstructions of all competing methods with ADVF at a representative limited-angle setting of 160° and a sparse-angle setting of 120 views. The average PSNR / SSIM over all geometry levels is reported in Table 1, and the corresponding per-level results are given in Appendix Tables 3–8.

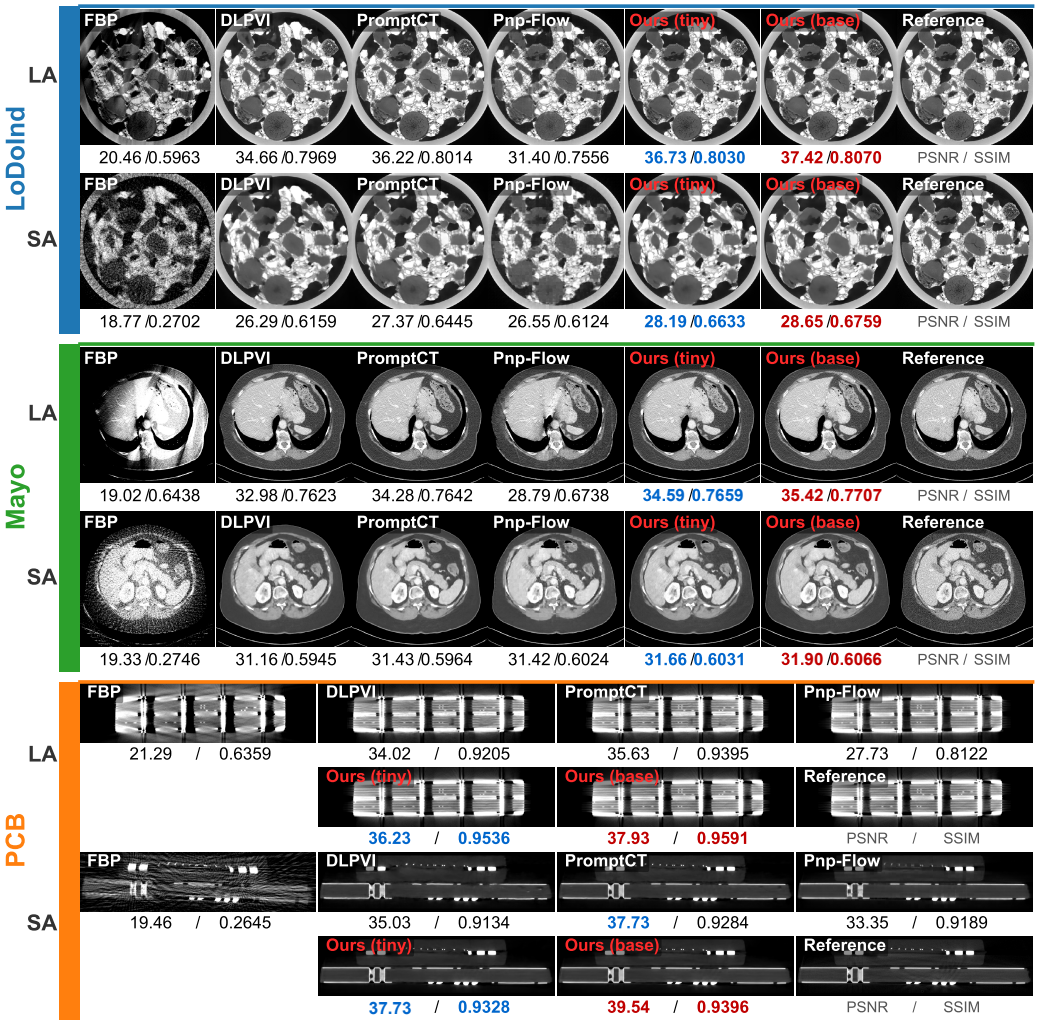


Figure 3: Qualitative comparison of ADVF against different types of reconstruction methods, shown for a limited-angle (LA) scan of 160° and a sparse-angle (SA) scan of 120 views. From left to right, the columns show FBP, DLPVI (direct mapping), PromptCT (unrolled), Pnp-Flow (flow-based), Ours (tiny), Ours (base), and the reference. From top to bottom, the rows show LoDoInd under LA, LoDoInd under SA, Mayo under LA, Mayo under SA, PCB under LA, and PCB under SA. The PSNR / SSIM of each reconstruction is reported below the corresponding panel.

Three observations can be made. First, the direct-mapping baseline DLPVI trails the iterative PromptCT and our method on all six settings. Since the missing angular information differs from one geometry level to the next, a single feed-forward mapping has to absorb all of this variation into one set of weights: the network is forced to learn at the image level how to restore the information that is missing at every setting, which increases its learning burden considerably and leaves it with limited robustness across settings. Second, the iterative baseline PromptCT keeps a certain robustness and is the strongest of the compared methods, but its number of refinement iterations is fixed: every step is subject to the same optimization constraint, whereas the amount of information that

still has to be recovered varies substantially along the trajectory. This mismatch bounds its accuracy, so it remains below our method in every setting, and it nevertheless stays clearly ahead of the flow-based Pnp-Flow. Third, although PnP-Flow also takes steps of different sizes and regularizes them with a pretrained denoiser, the step sizes are not matched to the geometry sequence: on PCB this regularization does lift the reconstruction above the plain direct mapping, but on the two medical benchmarks PnP-Flow remains the weakest of the learned methods.

Our method removes the fixed-step restriction and at the same time introduces the adaptive dual velocity field, which explicitly models the nonlinear transport of insufficient-angle reconstruction—precisely the nature of both the sparse-angle and the limited-angle processes. It therefore attains state-of-the-art accuracy on all six settings. Quantitatively, ADVF (base) improves over the best competing method by 0.92–2.88 dB in PSNR (1.81 dB on average) and by up to 0.026 in SSIM, and even the much smaller tiny variant, which contains only about $1/3.8$ of the parameters of the base model, still improves over the best competitor in every setting (by 0.61–1.65 dB). Table 1 further reports the averaged PSNR and SSIM of every method over the different scan angles and view counts, where the same ordering holds.

Table 1: Average PSNR / SSIM over scan levels t_0–t_8 for each method and dataset. Best results are marked in red and second-best results in blue.

Dataset	Setting	FBP	DLPVI	PromptCT	Pnp-Flow	Ours (tiny)	Ours (base)
LoDoInd	LA	23.12 / 0.5967	36.88 / 0.8003	37.72 / 0.8028	32.74 / 0.7563	38.66 / 0.8039	39.22 / 0.8102
LoDoInd	SA	22.95 / 0.4465	33.19 / 0.7326	33.70 / 0.7338	32.33 / 0.7160	35.11 / 0.7510	35.70 / 0.7593
Mayo	LA	23.19 / 0.6954	36.78 / 0.7967	37.60 / 0.7990	33.83 / 0.7463	38.32 / 0.8058	38.82 / 0.8059
Mayo	SA	23.92 / 0.4792	33.81 / 0.6980	34.05 / 0.7008	33.47 / 0.7028	34.66 / 0.7139	34.97 / 0.7146
PCB	LA	20.63 / 0.6103	34.44 / 0.9113	36.71 / 0.9268	34.82 / 0.9056	38.36 / 0.9412	39.59 / 0.9471
PCB	SA	20.48 / 0.5606	36.63 / 0.9224	38.13 / 0.9310	35.06 / 0.9217	39.10 / 0.9369	40.45 / 0.9478

5 Conclusion

We presented ADVF, a framework that recasts insufficient-angle CT reconstruction as learning along an ordered reconstruction trajectory and, along this trajectory, jointly learns an adaptive dual velocity field—a target-directed velocity and a local geometry-transition velocity. This departs from the position-independent fields learned by diffusion and flow models, whose straight-line transport cannot follow the nonlinear artifact evolution of IA-CT. Conditioning the transport on the current state and on its position along the trajectory, and realizing it as a progressive two-stage network in which lightweight transition modules complete the missing geometry levels and a reconstruction network aggregates them, allows a single trained model to serve every angle configuration of the same scanning system without retraining. Across three benchmarks (LoDoInd, Mayo, and PCB) and under both limited-angle and sparse-angle acquisition, ADVF improves over the strongest competing method by 0.92–2.88 dB in PSNR, while its much smaller variant remains competitive at a fraction of the parameters.

Two limitations remain. First, a single network covers different view numbers and angular ranges, but the remaining acquisition parameters must stay fixed: the model is tied to one scanning system, with a given source–object distance, source–detector distance, detector configuration, and pixel size, and transferring it to another system or to a different magnification still requires retraining. A reconstruction network that is general across acquisition geometries, and not only across the angular sampling, therefore remains an open challenge. Second, the current implementation reconstructs each slice independently in 2D and thus does not exploit the volumetric consistency between neighbouring slices; lifting the trajectory formulation from single slices to three-dimensional volumes is a natural but non-trivial next step.

Nevertheless, we believe this work is meaningful. It offers a new perspective on insufficient-angle reconstruction by making the geometry trajectory itself the object of learning, shows that a state- and position-dependent velocity field is what distinguishes this problem from conventional denoising or restoration, and delivers a single model that covers both sparse-angle and limited-angle settings with state-of-the-art accuracy. We hope these ideas will encourage further work on trajectory-aware reconstruction, on extending the framework to three-dimensional and volumetric settings, and on making such models general across acquisition geometries.

AI USE STATEMENT

(This section is **required** and does not count toward the page limit.)

The main idea, the method design and the core implementation of this work are our own. We used generative AI tools for the auxiliary tasks described below, and all AI-assisted artifacts were verified by the authors.

Tasks with required disclosure. We used ChatGPT 5.6 to (i) produce and then refine the method architecture diagram (Figure 2) from a sketch drawn by the authors, in which the architecture, the data flow and the layout were already fixed: the tool was used to render the diagram from this sketch with BioRender and to optimize its visual presentation, and every element of the final figure was checked by the authors, and (ii) debug and adapt open-source implementations of the comparison methods for training and evaluation, and to generate the initial U-Net backbone scaffold of our model. The core code that realizes the proposed method—the reconstruction-trajectory construction, the adaptive dual velocity field, the transition modules and the training curriculum—was written manually by the authors, who also verified that the adapted baseline implementations reproduce the intended comparison methods. We did not use generative AI tools to generate research ideas, to derive the proposed method, or to produce and analyse the experimental results.

Tasks with recommended disclosure. We used DeepSeek to retrieve and download related references; the relevance and the accuracy of every cited reference were checked manually by the authors before citation. We used DeepSeek V4.1 Flash to polish the language of the manuscript: all sentences were written by the authors first, the tool was used only to improve grammar, clarity and style, and it did not generate new claims, results or citations. Every suggested change was reviewed and manually accepted or rejected by the authors.

We have reviewed all AI-assisted work, and we verified the diagram, the text and the adapted baseline code against our own implementation and the experimental results. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

(This section is **recommended** and does not count toward the page limit.)

This work does not involve the collection of new data from human subjects. Two of the three benchmarks are public datasets used under their original licences: LoDoInd is a public benchmark of experimental low-dose industrial CT images of non-human test objects, and the Mayo benchmark is derived from the AAPM Low-Dose CT Grand Challenge, whose images were acquired and de-identified by the original providers under their own institutional review board approval and patient consent; we only use the released, de-identified slices and do not attempt to re-identify any individual. The third benchmark (PCB) was acquired in our laboratory on printed circuit boards, i.e., on non-human industrial objects that carry no personal or sensitive information, so no privacy or security concern arises from its use. We follow the respective terms of use of the public datasets and release no data that would violate them.

The proposed method is a general reconstruction technique for insufficient-angle CT. Its intended use is to improve image quality when the angular sampling is limited, which may in turn help reduce the radiation dose in medical scanning and the inspection time in industrial non-destructive testing; we are not aware of any direct harmful application. As with any learned reconstruction method, the accuracy depends on the training distribution: the models are trained and evaluated per scanning system, and the benchmark slices are limited in number and in anatomical or material diversity. We therefore make no claim of unbiased performance across populations, scanners or object types, and we report the acquisition settings of every benchmark so that the scope of our evaluation is explicit.

The authors declare no competing interests and no sponsorship that influenced the design, the results or the conclusions of this work. All experiments were conducted on public or laboratory-owned data in compliance with the corresponding licences. The use of generative AI tools in this work is disclosed in the AI use statement above.

Reproducibility statement

(This section is **recommended** and does not count toward the page limit.)

The proposed method is specified in Section 3.1 and Section 3.4, and its end-to-end training and inference procedures are stated step by step as Algorithm 1 and Algorithm 2 in Appendix A.1. The instantiation of the two networks and of the transition modules is described in Appendix A.2, and all training hyperparameters—optimizer, learning rate, weight decay, batch size, number of epochs, trajectory length and loss weights—together with the epoch boundaries of the three-phase curriculum are reported in Appendix A.5 and Appendix A.6, so that both the tiny and the base variant can be retrained from scratch.

All data used in this work are either public or documented exhaustively. The LoDoInd and Mayo benchmarks are public and their splitting protocol is stated in Section 4, while the in-house PCB benchmark is documented in Appendix A.4, which specifies the two acquired volumes, the SART reconstruction, the slice selection, the resizing and zero padding, the resulting number of slices and the training/test split. The simulation and reconstruction pipeline that turns every ground-truth slice into its multi-geometry trajectory—detector configuration, source–object and source–detector distances, pixel size, and the angular ranges and view counts of all geometry levels—is fully specified in Appendix Table 2, so that the inputs of every experiment can be regenerated from the ground-truth images alone.

The comparison methods are trained and evaluated with their publicly released implementations, adapted to our acquisition pipeline and run with the same data splits, the same trajectory inputs and the same evaluation code as our method, so that all methods are compared under identical conditions. Every quantitative result is a PSNR / SSIM average over the test split of a benchmark; the per-level values underlying each reported average are listed in Appendix Tables 3–8, and additional qualitative comparisons and intermediate rollouts are provided in Appendix Figures 4–9 (scan-angle sweeps) and 10–13 (intermediate rollouts). To support full reproduction, the source code, the configuration files and the scripts that build the trajectories are provided as anonymous supplementary material.

Author Contributions

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors.

Acknowledgments

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper.

References

Jonas Adler and Ozan Oktem. Learned primal-dual reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1322–1332, 2018.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

Anders H. Andersen and Avinash C. Kak. Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm. *Ultrasonic Imaging*, 6(1):81–94, 1984.

Rushil Anirudh, Hyojin Kim, Jayaraman J. Thiagarajan, K. Aditya Mohan, Kyle Champley, and Timo Bremer. Lose the views: Limited angle CT reconstruction via implicit sinogram completion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6343–6352, 2018.

Rui-Jun Bai, Yihong Li, Zhong Li, Luyang Li, Hanming Wei, Chenkai Zhao, Jia Guo, and Ping Chen. TD-SRN: Degradation-modulated tri-domain synergistic reconstruction network for sparse-view CT. *Pattern Recognition*, 2026. doi: 10.1016/j.patcog.2026.114441.

Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

Ashish Bora, Eric Price, and Alexandros G. Dimakis. Compressed sensing with generative models. In *International Conference on Machine Learning (ICML)*, 2018.

David J. Brenner and Eric J. Hall. Computed tomography—an increasing source of radiation exposure. *New England Journal of Medicine*, 357(18):1777–1784, 2007.

Gabriel Cardoso, Yazid Janati, Sylvain Le Corff, and Eric Moulines. Monte carlo guided diffusion for bayesian linear inverse problems. In *International Conference on Learning Representations (ICLR)*, 2024.

Hu Chen, Yi Zhang, Mannudeep K. Kalra, Feng Lin, Yang Chen, Peixi Liao, Jiliu Zhou, and Ge Wang. Low-dose ct with a residual encoder-decoder convolutional neural network. *IEEE Transactions on Medical Imaging*, 36(12):2524–2535, 2017.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.

Hyungjin Chung, Byeongsu Sim, Dohun Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 2022a.

Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022b.

Hyungjin Chung, Jeongsol Kim, Michael T. Mccann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *International Conference on Learning Representations (ICLR)*, 2023.

Ingrid Daubechies, Michel Defrise, and Christine De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021.

Zehao Dou and Yang Song. Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *International Conference on Learning Representations (ICLR)*, 2024.

Jintao Fu, Peng Cong, Shuo Xu, Jiahao Chang, Ximing Liu, and Yuewen Sun. Neural architecture search with deep radon prior for sparse-view CT image reconstruction. *Medical Physics*, 2025. doi: 10.1002/mp.17685.

Qi Gao, Zilong Li, Junping Zhang, Yi Zhang, and Hongming Shan. CoreDiff: Contextual error-modulated generalized diffusion model for low-dose CT denoising and generalization. *IEEE Transactions on Medical Imaging*, 43(2):745–759, 2024.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27, 2014.

Richard Gordon, Robert Bender, and Gabor T. Herman. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology*, 29(3):471–481, 1970.

Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *International Conference on Machine Learning (ICML)*, pp. 399–406, 2010.

Jiaqi Guo and Santiago López-Tapia. Diffusion-based limited-angle CT reconstruction under noisy conditions. In *IEEE International Conference on Image Processing Workshops (ICIPW)*, 2025. doi: 10.1109/icipw68931.2025.11386431.

Bernadette N. Hahn, Gaël Rigaud, and Richard Schmähl. The method of the approximate inverse for limited-angle CT. *Inverse Problems*, 2026. doi: 10.1088/1361-6420/ae5844.

Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018.

Yoseob Han and Jong Chul Ye. Framing U-Net via deep convolutional framelets: Application to sparse-view CT. *IEEE Transactions on Medical Imaging*, 37(6):1418–1429, 2018.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020.

Kensuke Hori, Fumio Hashimoto, Kazuya Koyama, and Takeyuki Hashimoto. Limited-angle SPECT image reconstruction using deep image prior. *Physics in Medicine and Biology*, 2025. doi: 10.1088/1361-6560/adea09.

Zhizhong Huang, Junping Zhang, Yi Zhang, and Hongming Shan. DU-GAN: Generative adversarial networks with dual-domain U-Net-based discriminators for low-dose CT denoising. *IEEE Transactions on Instrumentation and Measurement*, 71:1–12, 2022.

Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G. Dimakis, and Jonathan I. Tamir. Robust compressed sensing MRI with deep generative priors. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021.

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.

Avinash C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging*. SIAM, 2001.

Eunhee Kang, Won Chang, Jaejun Yoo, and Jong Chul Ye. Deep convolutional framelet denoising for low-dose CT via wavelet residual network. *IEEE Transactions on Medical Imaging*, 37(6): 1358–1369, 2018.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 2022.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Siyi Li, Chuandong Tan, Qi Li, Ruqian Song, Kuan Shen, and Yufang Cai. Physics-informed cold diffusion for sparse-view CT reconstruction under rapid battery scanning. *Displays*, 2026. doi: 10.1016/j.displa.2026.103563.

Yu Li, Xueqin Sun, Sukai Wang, Lina Guo, Yingwei Qin, Jinxiao Pan, and Ping Chen. TD-STrans: Tri-domain sparse-view CT reconstruction based on sparse transformer. *Computer Methods and Programs in Biomedicine*, 2025. doi: 10.1016/j.cmpb.2024.108575.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *International Conference on Learning Representations (ICLR)*, 2023.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *International Conference on Learning Representations (ICLR)*, 2023.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 2022.

Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*, 2023.

Samuel T. 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.

Cynthia H. McCollough, Baiyu Chen, David R. Holmes, III, Xinhui Duan, Zhicong Yu, Lifeng Yu, and Shuai Leng. Low-dose CT for the detection and classification of metastatic liver lesions: Results of the 2016 low dose CT grand challenge. *Medical Physics*, 44(10):e162–e176, 2017.

Vishal Monga, Yuelong Li, and Yonina C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.

Frank Natterer. *The Mathematics of Computerized Tomography*. SIAM, 2001.

Gregory Ongie, Ajil Jacob, and Jeffrey A. Fessler. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Topics in Signal Processing*, 14(4):704–726, 2020.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 234–241, 2015.

Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4):259–268, 1992.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations (ICLR)*, 2022.

Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Transactions on Medical Imaging*, 37(2):491–503, 2018.

Baoshun Shi, Ke Jiang, Qiusheng Lian, Xinran Yu, and Huazhu Fu. Prompting lipschitz-constrained network for multiple-in-one sparse-view CT reconstruction. *IEEE Transactions on Medical Imaging*, 45(4):1446–1460, 2026.

Jiayang Shi, Omar Elkilany, and Andreas Fischer. LoDoInd: Introducing a benchmark low-dose industrial CT dataset and enhancing denoising with 2.5d deep learning techniques. *e-Journal of Nondestructive Testing*, 2024. doi: 10.58286/29228.

Emil Y. Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine and Biology*, 53(17): 4777–4807, 2008.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, pp. 2256–2265, 2015.

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

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

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)*, 2021b.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning (ICML)*, 2023b.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Aaron Courville, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research (TMLR)*, 2024.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations (ICLR)*, 2024.

Yutong Xie and Quanzheng Li. Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 655–664, 2022.

Shuo Xu, Jintao Fu, Yuewen Sun, Peng Cong, and Xincheng Xiang. Deep radon prior: A fully unsupervised framework for sparse-view CT reconstruction. *Computers in Biology and Medicine*, 2025. doi: 10.1016/j.compbiomed.2025.109853.

Qingsong Yang, Pingkun Yan, Yanbo Zhang, Hengyong Yu, Yongyi Shi, Xuanqin Mou, Mannudeep K. Kalra, Yi Zhang, Ling Sun, and Ge Wang. Low-dose CT image denoising using a generative adversarial network with wasserstein distance and perceptual loss. *IEEE Transactions on Medical Imaging*, 37(6):1348–1357, 2018.

Yan Yang, Jian Sun, Hongbiao Li, and Zongben Xu. Deep ADMM-Net for compressive sensing MRI. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.

Jian Zhang and Bernard Ghanem. ISTA-Net: Iterative shrinkage-thresholding algorithm learned deep neural network for compressive sensing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1828–1837, 2018.

Xuzhi Zhao, Yi Du, and Yahui Peng. DLPVI: Deep learning framework integrating projection, view-by-view backprojection, and image domains for high- and ultra-sparse-view CBCT reconstruction. *Computerized Medical Imaging and Graphics*, 121:102508, 2025.