

Efficient Flow Matching for Sparse-View CT Reconstruction

Jiayang Shi^{1,2}, Lincen Yang², Zhong Li^{2,3}, Tristan van Leeuwen^{1,4}, Daniël M. Pelt², and K. Joost Batenburg²

¹ Centrum Wiskunde en Informatica

² LIACS, Leiden University

³ Great Bay University

⁴ Mathematical Institute, Utrecht University

`l.yang@liacs.leidenuniv.nl`

Abstract. Generative models, particularly Diffusion Models (DM), have shown strong potential for Computed Tomography (CT) reconstruction serving as expressive priors for solving ill-posed inverse problems. However, diffusion-based reconstruction relies on Stochastic Differential Equations (SDEs) for forward diffusion and reverse denoising, where such stochasticity can interfere with repeated data consistency corrections in CT reconstruction. Since CT reconstruction is often time-critical in clinical and interventional scenarios, improving reconstruction efficiency is essential. In contrast, Flow Matching (FM) models sampling as a deterministic Ordinary Differential Equation (ODE), yielding smooth trajectories without stochastic noise injection. This deterministic formulation is naturally compatible with repeated data consistency operations. Furthermore, we observe that FM-predicted velocity fields exhibit strong correlations across adjacent steps. Motivated by this, we propose an FM-based CT reconstruction framework (FMCT) and an efficient variant (EFMCT) that reuses previously predicted velocity fields over consecutive steps to substantially reduce the number of Neural network Function Evaluations (NFEs), thereby improving inference efficiency. We provide theoretical analysis showing that the error introduced by velocity reuse is bounded when combined with data consistency operations. Extensive experiments demonstrate that FMCT/EFMCT achieve competitive reconstruction quality while significantly improving computational efficiency compared with diffusion-based methods. The codebase is open-sourced at <https://github.com/EFMCT/EFMCT>.

Keywords: CT Reconstruction · Flow Matching · Generative Models.

1 Introduction

Generative models have recently gained popularity for Computed Tomography (CT) reconstruction due to their strong capability to learn expressive image priors [24,14,13,4]. Among them, Diffusion Models (DM) have achieved state-of-the-art (SOTA) performance in sparse-view CT reconstruction [13,4,18] by learning the distribution of full-view reconstructions and leveraging it as a powerful

prior for solving ill-posed inverse problems. Despite their strong reconstruction quality, DM methods are inherently stochastic because they model the forward diffusion and reverse denoising processes as Stochastic Differential Equations (SDEs) [21]. When combined with inverse problem solving, the repeated data consistency correction steps required for CT reconstruction can interfere with the stochastic evolution of the SDE [26,5], leading to a *push-and-pull effect* between prior-driven denoising and physics-based correction. This interaction often results in unstable or inconsistent reconstruction behavior [18,5], and necessitates a large number of iterations and Neural network Function Evaluations (NFEs), typically on the order of thousands. Although deterministic schedulers such as DDIM [19] have been proposed to reduce the number of reverse steps, achieving high-quality reconstructions with substantially fewer iterations remains challenging due to the inherently stochastic formulation of diffusion models.

In clinical practice, CT reconstruction is frequently time-critical. In emergency and interventional settings, rapid image availability directly impacts diagnosis and treatment decisions, making time to first image a key performance metric [8,23]. In addition, modern medical CT systems increasingly target higher spatial resolutions, which lead to larger reconstruction volumes and further exacerbate computational demands [11]. While conventional Filtered Backprojection (FBP) provides fast reconstructions, its image quality degrades substantially under sparse-view or low-dose acquisition, where advanced prior-based methods become necessary. As a result, the excessive computational cost of diffusion-based reconstruction methods poses a significant challenge for their practical deployment in real-world clinical workflows [18].

Flow Matching (FM) offers a fundamentally different modeling perspective that is particularly well suited to this setting. Instead of stochastic sampling, FM models deterministic probability transport through an Ordinary Differential Equation (ODE) [12]. Consequently, FM-based sampling follows a constant trajectory without stochastic noise injection. This deterministic formulation enables more coherent interaction with repeated data consistency corrections in inverse problems [25,17], thereby mitigating the *push-and-pull effect* observed in diffusion-based methods and allowing for more efficient sampling. Furthermore, the velocity fields learned by FM are observed to evolve smoothly along the ODE trajectory, with strong correlations between adjacent time steps, as demonstrated in [2] and empirically confirmed in our experiments (see Fig. 1). This observation suggests that repeatedly re-evaluating the neural network at every integration step may be redundant.

Motivated by these properties, we propose an efficient FM-based CT reconstruction framework that reuses previously predicted velocity fields to substantially reduce the number of NFEs. The reuse of velocity fields introduces only a controlled integration error, which we show is of the same order as the Euler discretization error and remains bounded through the explicit data consistency corrections. As a result, the proposed strategy achieves significant efficiency gains while maintaining competitive reconstruction quality. Our main contributions are summarized as follows: 1) We propose, to the best of our knowledge, the




Fig. 1: **Left:** Cosine similarity of predicted velocity fields at consecutive iterations in FMCT, with and without data consistency correction, showing strong correlation across adjacent steps. **Right:** Reconstruction performance vs. NFE for DM methods and FMCT/EFMCT, demonstrating that FMCT/EFMCT achieve competitive quality at substantially lower NFEs. DM methods here use the deterministic DDIM sampler for fairness; DPS/MCG results with the *original 1000-step DDPM sampler* are shown as *isolated points*. Results are averaged over the same randomly selected 21 reconstructions from the AAPM dataset across all iteration settings; shaded regions indicate standard deviation.

first flow-matching-based framework for CT reconstruction. 2) We introduce a velocity field reuse strategy that significantly improves the efficiency of FM-based CT reconstruction by reducing the required NFEs. 3) We provably show that the single-step reuse error is of the same order as Euler discretization, and does not alter the overall convergence behavior under bounded consecutive reuse steps. 4) We conduct extensive experiments to demonstrate the reconstruction performance and efficiency of the proposed FM-based CT reconstruction method.

2 Method

Problem Formulation. CT reconstruction aims to recover an unknown object $\mathbf{x} \in \mathbb{R}^m$ from a set of projection measurements $\mathbf{y} \in \mathbb{R}^n$. The measurement process can be mathematically modeled as a linear system

$$\mathbf{y} = \mathbf{A}\mathbf{x}, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{n \times m}$ is the system matrix determined by the acquisition geometry. In sparse-view CT, the number of measurements is insufficient ($n < m$), making the inverse problem underdetermined and ill-posed. In addition, measurement noise further degrades reconstruction quality. In this work, we focus on the sparse-view aspect of CT reconstruction.

Flow Matching. Given a source distribution $p(\mathbf{x}_1)$ and a target distribution $p(\mathbf{x}_0)$, FM defines a time-dependent velocity field $\mathbf{v}_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that transports samples from $p(\mathbf{x}_1)$ to $p(\mathbf{x}_0)$ via the ODE [12]

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_t(\mathbf{x}_t), \quad t \in [0, 1]. \quad (2)$$




Fig. 2: Overview of our proposed FMCT(left) and EFMCT(right) method.

The idea of FM is to train to approximate $\mathbf{v}_t$ by a neural network $\mathbf{v}_t^\theta$, which allows sampling from the target distribution by numerically integrating the ODE. The optimization objective of FM then becomes

$$\mathcal{L}_{FM} = \mathbb{E}_{t \sim \mathcal{U}(0,1), \mathbf{x}_t \sim p(\mathbf{x}_t)} \|\mathbf{v}_t(\mathbf{x}_t) - \mathbf{v}_t^\theta(\mathbf{x}_t)\|^2. \quad (3)$$

However, the marginal velocity field $\mathbf{v}_t(\mathbf{x}_t)$ is generally intractable, as it requires integration over all possible $\mathbf{x}_0$. The common way is to replace the marginal velocity with a conditional velocity field $\mathbf{v}_t(\mathbf{x}_t|\mathbf{x}_0)$, which is proven in [12] to yield gradients equivalent to those of Equation 3.

In this work, we focus on one specific conditional linear flow model, also known as Rectified Flow [15]. It models the transportation along straight-line trajectories through interpolation between $\mathbf{x}_0$ and $\mathbf{x}_1$

$$\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1, \quad \mathbf{v}_t(\mathbf{x}_t|\mathbf{x}_0) = \mathbf{x}_1 - \mathbf{x}_0. \quad (4)$$

This leads to the conditional flow matching loss

$$\mathcal{L}_{CFM} = \mathbb{E}_{t \sim \mathcal{U}(0,1), \mathbf{x}_0 \sim p(\mathbf{x}_0), \mathbf{x}_1 \sim p(\mathbf{x}_1)} \|\mathbf{v}_t^\theta((1-t)\mathbf{x}_0 + t\mathbf{x}_1) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2. \quad (5)$$

In practice, we sample $\mathbf{x}_1$ from a standard Gaussian distribution and use full-view CT images as $\mathbf{x}_0$. Once trained, the model enables deterministic sampling by integrating the learned velocity field backward from $t = 1$ to $t = 0$.

Flow Matching for CT Reconstruction. Given a pretrained FM model, we incorporate physics-based data consistency into the sampling process to enable CT reconstruction. Each integration step consists of three components.

First, a flow transport step is performed using an explicit Euler update

$$\mathbf{x}'_{t-\Delta t} = \mathbf{x}_t - \Delta t \mathbf{v}_t^\theta(\mathbf{x}_t). \quad (6)$$

Second, we estimate the terminal sample at $t = 0$ via linear extrapolation, and apply the data consistency operation $\mathcal{DC}(\cdot)$ to obtain the correction direction $\mathbf{v}_{DC}$ from measurement aware refinement

$$\hat{\mathbf{x}}_0 = \mathbf{x}_t - t \mathbf{v}_t^\theta(\mathbf{x}_t), \quad \mathbf{v}_{DC} = \mathcal{DC}(\hat{\mathbf{x}}_0, \mathbf{y}, \mathbf{A}). \quad (7)$$

In this work, we implement $\mathcal{DC}(\cdot)$ using conjugate gradient [9,6], though the framework is compatible with other data consistency schemes.

Finally, the data-consistent update is applied as

$$\mathbf{x}_{t-\Delta t} = \mathbf{x}'_{t-\Delta t} + \mathbf{v}_{\mathcal{DC}}. \quad (8)$$

These steps are repeated iteratively until $t = 0$ to obtain the final reconstruction. An overview of the proposed framework is illustrated in Fig. 2.

Velocity Reuse for Efficient Sampling. Empirically, we observe that the predicted velocities $\mathbf{v}_t^\theta$ are highly correlated in consecutive steps (Fig. 1), resulting in near-straight trajectories toward the reconstructed image. This behavior is consistent with the smooth evolution of velocity fields along the underlying ODE in flow matching. These observations suggest that repeatedly re-evaluating the neural network at every integration step may be redundant. Moreover, when combined with explicit data consistency operations, we argue that the small deviations introduced by reusing a previously predicted velocity can be effectively compensated by physics-based corrections, as conceptually illustrated in Fig. 2.

Specifically, after computing $\mathbf{v}_t = \mathbf{v}_t^\theta(\mathbf{x}_t)$ at time t , we reuse this velocity for up to M consecutive steps:

$$\mathbf{x}'_{t-(j+1)\Delta t} = \mathbf{x}_{t-j\Delta t} - \Delta t \mathbf{v}_t, \quad j \in [1, M]. \quad (9)$$

Since $\mathbf{v}_t$ is reused without additional network evaluations, this strategy directly reduces the number of NFEs. To ensure stability and data fidelity, we perform an adaptive refinement check at each reuse step

$$\|\mathbf{A}\mathbf{x}'_{t-(j+1)\Delta t} - \mathbf{y}\|^2 \leq \eta \|\mathbf{A}\mathbf{x}'_{t-j\Delta t} - \mathbf{y}\|^2, \quad (10)$$

where $\eta > 1$ is a relaxation factor. If the condition is violated, velocity reuse is terminated and the velocity is recomputed using the neural network. This adaptive mechanism balances efficiency and reconstruction accuracy. η is empirically set at 1.05. The complete algorithm is summarized in the appendix.

Proposition 1. *Assume that $\mathbf{v}$ is locally Lipschitz in both $\mathbf{x}$ and t , i.e., $\|\mathbf{v}_t(\mathbf{x}) - \mathbf{v}_t(\mathbf{x}')\| \leq L_x \|\mathbf{x} - \mathbf{x}'\|$ and $\|\mathbf{v}_t(\mathbf{x}) - \mathbf{v}_{t'}(\mathbf{x})\| \leq L_t |t - t'|$ for all relevant $\mathbf{x}, \mathbf{x}', t, t'$. Then, for a single reuse step, the local deviation between the reuse update $\tilde{\mathbf{x}}_{k+1}$ and the standard Euler update $\mathbf{x}_{k+1}$ satisfies $\|\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_{k+1}\| = O(\Delta t^2)$, i.e., velocity reuse introduces a local error of the same order as the Euler discretization itself. Moreover, if the same velocity is reused for at most M consecutive steps (independent of Δt), and the data-consistency correction is non-expansive with respect to its image argument, then the deviation remains controlled. The accumulated deviation after one reuse block is $O(M^2 \Delta t)$.*

Implication. The proposition implies that velocity reuse introduces a local error of the same order as the Euler discretization itself. When the reuse length is bounded, the accumulated error remains controlled over consecutive reuse steps.

Proofs are in Appendix.

Table 1: Quantitative comparison of reconstruction performance across different methods. Mean and standard deviation are reported for PSNR, SSIM, and data fidelity. The best average result for each metric is shown in **bold**, and the second best is underlined. NFE* counts only neural network forward evaluations. ↓ indicates efficiency improvement in NFE and computation time (on RTX4090) of FMCT/EFMCT compared with the most efficient diffusion-based method.

Dataset	Method	40 views			20 views			NFE*	Time/s
		PSNR	SSIM	Data Fit	PSNR	SSIM	Data Fit		
AAPM	FBP	26.98 ± 0.45	0.691 ± 0.027	2783.64 ± 233.30	24.01 ± 0.49	0.547 ± 0.029	1798.77 ± 171.00	-	0.02
	ADMM-TV	31.94 ± 0.42	0.836 ± 0.027	<u>22.11</u> ± 2.29	29.45 ± 0.46	0.798 ± 0.029	<u>17.53</u> ± 2.00	-	9.85
	DPS	30.05 ± 0.69	0.795 ± 0.033	261.35 ± 20.58	28.78 ± 0.81	0.764 ± 0.032	142.93 ± 35.99	1000	142.16
	MCG	30.00 ± 0.79	0.791 ± 0.032	198.07 ± 14.20	27.57 ± 0.42	0.736 ± 0.031	180.83 ± 11.87	1000	143.44
	PGDM	30.26 ± 0.61	0.803 ± 0.030	208.98 ± 13.72	29.43 ± 0.60	0.775 ± 0.031	106.51 ± 5.30	100	29.79
	DDS	31.06 ± 0.40	<u>0.834</u> ± 0.030	19.09 ± 1.14	29.21 ± 0.36	0.785 ± 0.031	9.87 ± 0.61	50	3.83
	FMCT	<u>31.63</u> ± 0.39	0.810 ± 0.031	<u>60.57</u> ± 18.52	29.98 ± 0.23	0.769 ± 0.030	46.46 ± 14.86	<u>25</u> (↓ 50%)	<u>1.92</u> (↓ 50%)
	EFMCT	31.52 ± 0.28	0.832 ± 0.027	182.80 ± 104.39	<u>29.52</u> ± 0.23	<u>0.795</u> ± 0.027	111.41 ± 79.80	7 (↓ 75%)	0.83 (↓ 78%)
Decathlon	FBP	28.01 ± 0.76	0.763 ± 0.037	2772.79 ± 434.36	24.85 ± 0.96	0.618 ± 0.039	1731.97 ± 305.77	-	0.02
	ADMM-TV	34.40 ± 1.23	0.900 ± 0.021	<u>15.73</u> ± 0.55	30.26 ± 1.03	0.851 ± 0.028	<u>14.76</u> ± 0.50	-	9.85
	DPS	33.85 ± 1.25	0.870 ± 0.018	250.11 ± 18.61	31.69 ± 1.44	0.842 ± 0.025	121.25 ± 18.57	1000	142.16
	MCG	30.03 ± 1.20	0.824 ± 0.030	301.17 ± 35.25	28.67 ± 1.26	0.798 ± 0.038	198.22 ± 22.68	1000	143.44
	PGDM	34.78 ± 1.23	0.899 ± 0.014	147.10 ± 15.20	31.82 ± 1.19	0.860 ± 0.018	132.96 ± 13.69	100	29.79
	DDS	37.66 ± 1.20	<u>0.923</u> ± 0.009	15.10 ± 9.29	33.00 ± 1.29	<u>0.873</u> ± 0.018	5.01 ± 0.36	100	7.65
	FMCT	<u>37.65</u> ± 1.53	0.912 ± 0.020	334.44 ± 74.21	33.41 ± 1.57	0.841 ± 0.036	294.97 ± 58.87	<u>50</u> (↓ 50%)	<u>4.09</u> (↓ 47%)
	EFMCT	37.56 ± 1.37	0.940 ± 0.014	621.99 ± 48.99	<u>33.21</u> ± 1.29	0.888 ± 0.024	470.58 ± 33.62	11 (↓ 89%)	1.72 (↓ 78%)

3 Results

Comparison Methods. We compare the proposed method with several representative diffusion-based CT reconstruction approaches, including Diffusion Posterior Sampling (DPS) [5], Manifold Constrained Gradient (MCG) [7], Pseudoinverse-Guided Diffusion Models (PGDM) [20], and Decomposed Diffusion Sampler (DDS) [6]. These methods represent a diverse set of strategies for incorporating measurement information into diffusion-based reconstruction and serve as strong SOTA baselines. In addition to learning-based methods, we include two classical reconstruction approaches, FBP and Model-Based Iterative Reconstruction (MBIR) method, Alternating Direction Method of Multipliers with Split-Bregman Total Variation (ADMM-TV) [3,10].

Datasets. We conduct experiments on two publicly available CT datasets to evaluate the reconstruction performance of the proposed method. The first dataset is the AAPM 2016 Low Dose CT Grand Challenge dataset [16]. From this dataset, CT images (512×512 pixels) from nine patients are used to train the diffusion and flow matching models, while images from patient *L506* are reserved for testing. The second dataset is the Medical Segmentation Decathlon CT dataset [1]. From the *Task_06 Lung* subset (512×512 pixels), we randomly select ten patients for training and evaluate reconstruction performance on images from patient *017*. To obtain sparse-view measurements, we simulate 20-view and 40-view projection data using the ASTRA Toolbox [22]. All experiments are conducted using a parallel-beam geometry. This choice is made deliberately to provide a controlled and widely used experimental setting that minimizes confounding effects introduced by geometric complexity and allows us to isolate the impact of the reconstruction methodology itself. The proposed FMCT and




Fig. 3: Visual comparison of reconstructions across different methods, datasets, and views. PSNR, SSIM, and reconstruction time are shown in the lower-left, lower-right, and upper-right corners of each image, respectively. We visualize only the best three diffusion-based methods for better visibility. Zoom in for more details.

EFMCT with velocity reuse strategy do not rely on geometry-specific assumptions and can be readily extended to other acquisition geometries by modifying the forward operator.

Reconstruction Performance. Table 1 reports the quantitative reconstruction results. FMCT/EFMCT are often among the highest PSNR and SSIM across most datasets and view configurations, while requiring substantially fewer NFEs and lower computation time than diffusion-based baselines. ADMM-TV and DDS achieve the best and second-best data fidelity, reflecting their strong emphasis on measurement consistency. Compared with FMCT, EFMCT exhibits a slight reduction in PSNR and data fidelity but attains comparable or improved SSIM, accompanied by a significant efficiency improvement. This demonstrates that the proposed velocity reuse strategy effectively reduces computational cost while preserving reconstruction quality.

Figure 3 provides a visual comparison. Although ADMM-TV achieves high PSNR/SSIM, its reconstructions are overly smooth and lack fine details. Diffusion-based methods and FMCT/EFMCT produce visually sharper and more detailed reconstructions, though they do not always exactly match the reference. This behavior is expected in sparse-view CT, where the problem is severely ill-posed and multiple reconstructions can satisfy the measurements. Generative methods bal-




Fig. 4: Ablation of the velocity reuse strategy. Left: Reuse enabled after different iteration indices. Right: Varying the maximum number of consecutive reuse steps. Shaded regions indicate standard deviation.

ance data consistency with adherence to the learned prior, which may lead to small deviations in fine-scale details.

Impact of Velocity Reuse Strategy. We analyze the effect of the proposed velocity reuse strategy through an ablation study shown in Fig. 4. The experiments are conducted on 21 randomly selected test CT images, with the total number of sampling iterations fixed to 50. In the first ablation, we vary the iteration at which velocity reuse is enabled while fixing the maximum number of consecutive reuse steps to 10. Enabling reuse too early leads to an initial drop in reconstruction performance, whereas delaying reuse to later iterations gradually improves PSNR and SSIM. This behavior suggests that early-stage sampling is more sensitive to approximation errors, while reuse becomes safer once the reconstruction has approached a more stable trajectory. Consequently, velocity reuse is better suited to mid- or late-stage sampling. In the second ablation, velocity reuse is enabled from the first iteration, and the maximum number of allowed consecutive reuse steps is varied. Increasing the reuse limit initially improves PSNR and SSIM, indicating that moderate reuse does not adversely affect reconstruction quality. However, allowing excessive consecutive reuse results in a general performance degradation, revealing a trade-off between reconstruction quality and computational efficiency. Aggressive reuse substantially reduces computational cost but may compromise reconstruction accuracy. For all EFMCT results, we adopt a reuse strategy that enables velocity reuse after the first iteration and allows up to 10 consecutive reuse steps in all experiments.

4 Conclusion

In this work, we proposed FMCT/EFMCT, representing the first application of flow matching to CT reconstruction. By exploiting the deterministic transport structure of flow matching, we introduced a velocity reuse strategy that substantially reduces NFEs. We provided theoretical analysis showing that the error introduced by velocity reuse is of the same order as Euler discretization and remains bounded when combined with explicit data consistency operations.

Extensive experiments on multiple datasets demonstrate that FMCT/EFMCT achieve competitive reconstruction quality while significantly improving computational efficiency compared with diffusion-based methods. We believe that the proposed framework offers a practical and principled approach toward efficient generative CT reconstruction, with the potential to facilitate the deployment of advanced reconstruction methodologies in real-world clinical settings.

Acknowledgements. The authors are supported by European Union H2020-MSCA-ITN-2020 under grant agreement no. 956172 and Dutch Research Council under grant no. ENWSS.2018.003, NWA.1160.18.316.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature communications* **13**(1), 4128 (2022)
2. Bajpai, D.J., Agarwal, S., Saxena, A., Kulkarni, K., Mitra, S., Hanawal, M.K.: Flowcast: Trajectory forecasting for scalable zero-cost speculative flow matching. In: *The Fourteenth International Conference on Learning Representations* (2026), <https://openreview.net/forum?id=gv8R88PVEG>
3. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning* **3**(1), 1–122 (2011)
4. Chen, J., Lin, Y., Qin, Y., Wang, H., Li, X.: Cross-view generalized diffusion model for sparse-view ct reconstruction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 140–150. Springer (2025)
5. Chung, H., Kim, J., McCann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=0nD9zGAGT0k>
6. Chung, H., Lee, S., Ye, J.C.: Decomposed diffusion sampler for accelerating large-scale inverse problems. In: *The Twelfth International Conference on Learning Representations* (2024), <https://openreview.net/forum?id=DsEhqQtFAG>
7. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) *Advances in Neural Information Processing Systems* (2022), <https://openreview.net/forum?id=nJJjv0JDJju>
8. Fanucci, E., Fiaschetti, V., Rotili, A., Floris, R., Simonetti, G.: Whole body 16-row multislice ct in emergency room: effects of different protocols on scanning time, image quality and radiation exposure. *Emergency radiology* **13**(5), 251–257 (2007)
9. Fletcher, R., Reeves, C.M.: Function minimization by conjugate gradients. *The computer journal* **7**(2), 149–154 (1964)
10. Goldstein, T., Osher, S.: The split bregman method for l1-regularized problems. *SIAM journal on imaging sciences* **2**(2), 323–343 (2009)

11. Hussain, S., Mubeen, I., Ullah, N., Shah, S.S.U.D., Khan, B.A., Zahoor, M., Ullah, R., Khan, F.A., Sultan, M.A.: Modern diagnostic imaging technique applications and risk factors in the medical field: a review. *BioMed research international* **2022**(1), 5164970 (2022)
12. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=PqvMRDCJT9t>
13. Liu, J., Anirudh, R., Thiagarajan, J.J., He, S., Mohan, K.A., Kamilov, U.S., Kim, H.: Dolce: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10498–10508 (2023)
14. Liu, W., Ding, H.: Solving low-dose ct reconstruction via gan with local coherence. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 524–534. Springer (2023)
15. Liu, X., Gong, C., qiang liu: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=XVjTT1nw5z>
16. McCollough, C.H., Bartley, A.C., Carter, R.E., Chen, B., Drees, T.A., Edwards, P., Holmes III, D.R., Huang, A.E., Khan, F., Leng, S., et al.: 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), e339–e352 (2017)
17. Pourya, M., Rawas, B.E., Unser, M.: FLOWER: A flow-matching solver for inverse problems. In: The Fourteenth International Conference on Learning Representations (2026), <https://openreview.net/forum?id=QGd34p02mI>
18. Shi, J., Pelt, D.M., Batenburg, K.J.: DM4CT: Benchmarking diffusion models for computed tomography reconstruction. In: The Fourteenth International Conference on Learning Representations (2026), <https://openreview.net/forum?id=YE5scJekg5>
19. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=St1giarCHLP>
20. Song, J., Vahdat, A., Mardani, M., Kautz, J.: Pseudoinverse-guided diffusion models for inverse problems. In: International Conference on Learning Representations (2023), https://openreview.net/forum?id=9_gsMA8MRKQ
21. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=PxtIG12RRHS>
22. Van Aarle, W., Palenstijn, W.J., Cant, J., Janssens, E., Bleichrodt, F., Dabrvolski, A., De Beenhouwer, J., Joost Batenburg, K., Sijbers, J.: Fast and flexible x-ray tomography using the astra toolbox. *Optics express* **24**(22), 25129–25147 (2016)
23. Willeminck, M.J., Schilham, A.M., Leiner, T., Mali, W.P.T.M., de Jong, P.A., Budde, R.P.: Iterative reconstruction does not substantially delay ct imaging in an emergency setting. *Insights into imaging* **4**(3), 391–397 (2013)
24. Ying, X., Guo, H., Ma, K., Wu, J., Weng, Z., Zheng, Y.: X2ct-gan: reconstructing ct from biplanar x-rays with generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10619–10628 (2019)
25. Zhang, Y., Yu, P., Zhu, Y., Chang, Y., Gao, F., Wu, Y.N., Leong, O.: Flow priors for linear inverse problems via iterative corrupted trajectory matching. *Advances in Neural Information Processing Systems* **37**, 57389–57417 (2024)

26. Zhang, Y., Liu, Z., Li, Z., Li, Z., Clark, J.J., Si, X.: Decoupling training-free guided diffusion by admm. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 23292–23302 (2025)

5 Appendix: Proof of Proposition 1

Proof. Let $\mathbf{v} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a velocity field. Define the explicit noise-to-image Euler update as

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \Delta t \mathbf{v}_{t_k}(\mathbf{x}_k), \quad t_{k+1} = t_k - \Delta t. \quad (11)$$

Define the velocity-reuse update for M consecutive steps using a previously evaluated velocity computed at $(t_{k_0}, \mathbf{x}_{k_0})$:

$$\tilde{\mathbf{x}}_{k+1} = \mathbf{x}_k - \Delta t \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0}), \quad k = k_0, \dots, k_0 + M - 1. \quad (12)$$

Assume Lipschitz continuity $L_x, L_t \geq 0$ such that for all relevant $\mathbf{x}, \mathbf{x}', t, t'$,

$$\|\mathbf{v}_t(\mathbf{x}) - \mathbf{v}_t(\mathbf{x}')\| \leq L_x \|\mathbf{x} - \mathbf{x}'\|, \quad \|\mathbf{v}_t(\mathbf{x}) - \mathbf{v}_{t'}(\mathbf{x})\| \leq L_t |t - t'|. \quad (13)$$

Single Step Error. Assume reuse the velocity at the first step $k_0 = k - 1$ gives error e_k ,

$$\begin{aligned} e_k &= \|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k-1}}(\mathbf{x}_{k-1})\| \\ &\leq \|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k-1}}(\mathbf{x}_k)\| + \|\mathbf{v}_{t_{k-1}}(\mathbf{x}_k) - \mathbf{v}_{t_{k-1}}(\mathbf{x}_{k-1})\| \\ &\leq L_t |t_k - t_{k-1}| + L_x \|\mathbf{x}_k - \mathbf{x}_{k-1}\| \\ &\leq L_t \Delta t + L_x \Delta t \|\mathbf{v}_{t_{k-1}}(\mathbf{x}_{k-1})\|. \end{aligned} \quad (14)$$

Then the error in $\tilde{\mathbf{x}}_{k+1}$ and $\mathbf{x}_{k+1}$ becomes

$$\begin{aligned} \|\mathbf{x}_{k+1} - \tilde{\mathbf{x}}_{k+1}\| &= \Delta t \|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k-1}}(\mathbf{x}_{k-1})\| \\ &\leq \Delta t^2 (L_t + L_x \|\mathbf{v}_{t_{k-1}}(\mathbf{x}_{k-1})\|). \end{aligned} \quad (15)$$

It shows that the reusing the velocity gives single step error of $O(\Delta t^2)$, which matches the error of standard Euler $x(t - \Delta t) = x(t) - \Delta t \dot{x} + O(\Delta t^2)$.

Accumulated Error. According to our proposed algorithm, we define standard single-step state update as

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \Delta t \mathbf{v}_{t_k}(\mathbf{x}_k) + \mathcal{DC}(\mathbf{x}_k - t_k \mathbf{v}_{t_k}(\mathbf{x}_k), \mathbf{y}, \mathbf{A}), \quad (16)$$

and the corresponding state update with velocity-reuse steps as

$$\tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_k - \Delta t \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0}) + \mathcal{DC}(\tilde{\mathbf{x}}_k - t_k \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0}), \mathbf{y}, \mathbf{A}), \quad k = k_0, \dots, k_0 + M - 1, \quad (17)$$

where $\mathcal{DC}(\mathbf{x}, \mathbf{y}, \mathbf{A})$ is the data consistency operation given measurement $\mathbf{y}$ and forward operator $\mathbf{A}$.

Here the same velocity $\mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0})$ is reused for M consecutive steps. Assume $\mathcal{DC}(\mathbf{x}, \mathbf{y}, \mathbf{A})$ is κ -non-expansive with respect to its image argument,

$$\|\mathcal{DC}(\mathbf{x}, \mathbf{y}, \mathbf{A}) - \mathcal{DC}(\mathbf{x}', \mathbf{y}, \mathbf{A})\| \leq \kappa \|\mathbf{x} - \mathbf{x}'\|, \quad 0 \leq \kappa \leq 1. \quad (18)$$

Let $\delta_k = \tilde{\mathbf{x}}_k - \mathbf{x}_k$. Then

$$\begin{aligned} \|\delta_{k+1}\| &\leq \|\delta_k\| + \Delta t \|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0})\| \\ &\quad + \kappa \|\delta_k - t_k (\mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0}) - \mathbf{v}_{t_k}(\mathbf{x}_k))\| \\ &\leq (1 + \kappa) \|\delta_k\| + (\Delta t + \kappa t_k) \underbrace{\|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0})\|}_{\text{velocity mismatch}}. \end{aligned} \quad (19)$$

Using Lipschitz continuity, we have

$$\underbrace{\|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0})\|}_{\text{velocity mismatch}} \leq L_x \underbrace{\|\mathbf{x}_k - \mathbf{x}_{k_0}\|}_{\text{state drift}} + L_t |t_k - t_{k_0}|. \quad (20)$$

Assume the practical updates have bounded increments,

$$\|\mathbf{x}_{j+1} - \mathbf{x}_j\| \leq \Delta t U_{\max}, \quad (21)$$

which holds when both the velocity field and the DC correction remain bounded along the trajectory. Then

$$\begin{aligned} \underbrace{\|\mathbf{x}_k - \mathbf{x}_{k_0}\|}_{\text{state drift}} &\leq \sum_{j=k_0}^{k-1} \|\mathbf{x}_{j+1} - \mathbf{x}_j\| \\ &\leq (k - k_0) \Delta t U_{\max}. \end{aligned} \quad (22)$$

Since $|t_k - t_{k_0}| = (k - k_0) \Delta t$, we obtain

$$\underbrace{\|\mathbf{v}_{t_k}(\mathbf{x}_k) - \mathbf{v}_{t_{k_0}}(\mathbf{x}_{k_0})\|}_{\text{velocity mismatch}} \leq (k - k_0) \Delta t (L_t + L_x U_{\max}). \quad (23)$$

Let $C = L_t + L_x V_{\max}$ and $t_{\max} = \max_k t_k$. Then

$$\|\delta_{k+1}\| \leq (1 + \kappa) \|\delta_k\| + (\Delta t + \kappa t_{\max})(k - k_0) \Delta t C. \quad (24)$$

Unrolling over M consecutive reuse steps with $\delta_{k_0} = 0$ gives

$$\begin{aligned} \underbrace{\|\delta_{k_0+M}\|}_{\text{accumulated deviation}} &\leq (\Delta t + \kappa t_{\max}) C \Delta t \sum_{j=0}^{M-1} (1 + \kappa)^{M-1-j} j \\ &\leq (\Delta t + \kappa t_{\max}) C \Delta t (1 + \kappa)^{M-1} \frac{M(M-1)}{2}. \end{aligned} \quad (25)$$

Therefore for bounded M ,

$$\|\delta_{k_0+M}\| = O(M^2 \Delta t), \quad (26)$$

i.e. the accumulated error remains controlled for bounded consecutive reuse steps.

6 Appendix: Algorithm

We summarize the complete EFMCT algorithm in 1.

Algorithm 1 EFMCT: Flow Matching with Velocity Reuse

Require: Forward operator $\mathbf{A}$, measurements $\mathbf{y}$, FM model $\mathbf{v}_t^\theta(\cdot)$, number of steps N , max reuse steps M , relaxation factor η

- 1: Initialize $t \leftarrow 1$, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\Delta t \leftarrow 1/N$ ▷ Start from Gaussian
- 2: $\mathbf{v}_{\text{prev}} \leftarrow \emptyset$, $r \leftarrow \|\mathbf{A}\mathbf{x}_t - \mathbf{y}\|^2$, $m \leftarrow 0$ ▷ Residual r and reuse counter m
- 3: **for** $i = 1$ to N **do**
- // Attempt velocity reuse if it improves data consistency
- 4: **if** $\mathbf{v}_{\text{prev}} \neq \emptyset$ **and** $m < M$ **then**
- 5: $\mathbf{x}'_{t-\Delta t} \leftarrow \mathbf{x}_t - \Delta t \mathbf{v}_{\text{prev}}$, $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}'_{t-\Delta t} - t \mathbf{v}_{\text{prev}}$, $\tilde{r} \leftarrow \|\mathbf{A}\mathbf{x}'_{t-\Delta t} - \mathbf{y}\|^2$ ▷
- Reuse velocity
- 6: **if** $\tilde{r} \leq \eta r$ **then**
- 7: $\mathbf{x}_{t-\Delta t} \leftarrow \mathbf{x}'_{t-\Delta t} + \mathbf{v}_{\text{DC}}(\hat{\mathbf{x}}_0, \mathbf{y}, \mathbf{A})$, $r \leftarrow \tilde{r}$, $m \leftarrow m + 1$ ▷ DC update
- 8: **else**
- 9: $\mathbf{v}_{\text{prev}} \leftarrow \emptyset$, $m \leftarrow 0$
- 10: **continue**
- 11: **end if**
- 12: **else**
- // Evaluate FM network at current time step
- 13: $\mathbf{x}'_{t-\Delta t} \leftarrow \mathbf{x}_t - \Delta t \mathbf{v}_t^\theta(\mathbf{x}_t)$, $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}'_{t-\Delta t} - t \mathbf{v}_t^\theta(\mathbf{x}_t)$
- 14: $\mathbf{x}_{t-\Delta t} \leftarrow \mathbf{x}'_{t-\Delta t} + \mathbf{v}_{\text{DC}}(\hat{\mathbf{x}}_0, \mathbf{y}, \mathbf{A})$, $r \leftarrow \|\mathbf{A}\mathbf{x}_{t-\Delta t} - \mathbf{y}\|^2$
- 15: $\mathbf{v}_{\text{prev}} \leftarrow \mathbf{v}_t^\theta(\mathbf{x}_t)$, $m \leftarrow 0$
- 16: **end if**
- 17: $t \leftarrow t - \Delta t$, $\mathbf{x}_t \leftarrow \mathbf{x}_{t-\Delta t}$
- 18: **end for**
- 19: **return** $\mathbf{x}_0$

7 Appendix: Extra Results

Table 2: Quantitative comparison of reconstruction perceptual quality across different methods. Mean and standard deviation are reported for LPIPS, where lower values indicate better perceptual similarity. The best average result for each setting is shown in **bold**, and the second best is underlined. NFE* counts only neural network forward evaluations. ↓ indicates efficiency improvement in NFE and computation time (on RTX4090) of FMCT/EFMCT compared with the most efficient diffusion-based method.

Dataset	Method	40 views	20 views	NFE*	Time/s
		LPIPS ↓	LPIPS ↓		
AAPM	FBP	0.391 ± 0.014	0.554 ± 0.021	-	0.02
	ADMM-TV	0.272 ± 0.027	0.380 ± 0.028	-	9.85
	DPS	0.128 ± 0.018	<u>0.166 ± 0.025</u>	1000	142.16
	MCG	0.145 ± 0.017	0.208 ± 0.028	1000	143.44
	PGDM	0.159 ± 0.022	0.187 ± 0.025	100	29.79
	DDS	<u>0.114 ± 0.013</u>	0.181 ± 0.021	50	3.83
	FMCT	0.101 ± 0.010	0.148 ± 0.015	<u>25</u> (↓ 50%)	<u>1.92</u> (↓ 50%)
	EFMCT	0.129 ± 0.014	0.194 ± 0.021	7 (↓ 75%)	0.83 (↓ 78%)
Decathlon	FBP	0.348 ± 0.018	0.512 ± 0.026	-	0.02
	ADMM-TV	0.195 ± 0.025	0.319 ± 0.021	-	9.85
	DPS	0.128 ± 0.037	0.140 ± 0.038	1000	142.16
	MCG	0.164 ± 0.034	0.186 ± 0.026	1000	143.44
	PGDM	0.100 ± 0.129	0.124 ± 0.032	100	29.79
	DDS	<u>0.072 ± 0.015</u>	0.110 ± 0.017	100	7.65
	FMCT	0.109 ± 0.021	0.180 ± 0.026	<u>50</u> (↓ 50%)	<u>4.09</u> (↓ 47%)
	EFMCT	0.070 ± 0.011	<u>0.120 ± 0.024</u>	11 (↓ 89%)	1.72 (↓ 78%)