Physics-Informed Cold Diffusion Framework with Time-Conditioned U-Net for Multi-Coil Brain MRI Reconstruction

Nives Grčić, Marija Habijan, Krešimir Matić, Irena Galić
Faculty of Electrical Engineering, Computer Science and Information Technology Osijek, Croatia
E-mail: marija.habijan@ferit.hr

Keywords: Accelerated MRI, brain imaging, cold diffusion models, k-space reconstruction, medical image processing

Received: October 15, 2025

Magnetic Resonance Imaging (MRI) reconstruction from undersampled multi-coil k-space data remains a central challenge for accelerating clinical scans. This paper proposes a cold diffusion based reconstruction framework that integrates a deterministic k-space forward operator with a time-conditioned U-Net denoiser. Unlike conventional diffusion models that rely on stochastic Gaussian noise, the proposed approach explicitly models the physical undersampling process, yielding interpretable and data-consistent reconstructions. The model was trained and evaluated on the fastMRI multi-coil brain dataset using 2D slices center-cropped to 320×320 *and normalized in complex form (real + imaginary channels). Ablation studies compare the proposed method against a supervised U-Net baseline and a Gaussian DDPM trained under identical conditions. Quantitatively, the cold diffusion model with 500 diffusion steps achieved 37.8 dB PSNR, 0.42 SSIM, and 0.46 HFEN, outperforming both the DDPM (32.8 dB, 0.33 SSIM, 0.60 HFEN) and supervised U-Net (28.9 dB, 0.29 SSIM, 0.71 HFEN) while reducing inference time by 70% relative to conventional diffusion sampling. Qualitative results indicate improved anatomical sharpness, stable convergence and consistent reconstruction across slices, suggesting that the physics-informed diffusion framework is a reliable and interpretable approach for accelerated MRI reconstruction.*

Povzetek: Raziskava predstavlja okvir za rekonstrukcijo večzavojnih možganskih MRI-posnetkov, zasnovan na hladni difuziji, ki modelira nezadostno vzorčenje v k-prostoru kot deterministični napredni proces ter uporablja časovno pogojeno mrežo U-Net za postopno obnavljanje visokokakovostnih, artefaktov prostih slik z izboljšano strukturno in zaznavno zvestobo.

1 Introduction

Medical imaging has a significant role in modern healthcare. It enables early diagnosis, treatment planning, and personalized interventions [39]. Among the available imaging modalities, Magnetic Resonance Imaging (MRI) is particularly valuable because of its ability to non-invasively capture high-resolution images of soft tissues with superior contrast [11]. This makes MRI crucial for diagnosing neurological disorders, tumors, and other pathologies where fine anatomical details are critical [18]. Despite these advantages, a significant limitation of MRI is its long acquisition times, which contribute to patient discomfort, motion-induced artifacts, and reduced clinical throughput.

Accelerated MRI reconstruction has emerged as a key strategy for shortening scan times by acquiring only a fraction of k-space data and then reconstructing a diagnostic-quality image from undersampled measurements [14, 27]. Traditional methods such as parallel imaging (PI) and compressed sensing (CS) have demonstrated that acceleration is possible, but they typically rely on handcrafted priors and iterative optimization. This often results in residual artifacts or loss of detail under aggressive undersampling. The rise of machine and deep learning models, such as the U-Net architecture [29], has significantly advanced reconstruction quality by learning complex mappings between undersampled inputs and fully sampled references. Other approaches, such as deep unrolling, integrate the interpretability of optimization algorithms with the flexibility of neural networks. While effective, these methods face challenges. U-Nets may oversmooth critical details, unrolling networks can be computationally intensive, while generative adversarial networks (GANs) may introduce hallucinated features that limit clinical trust.

Recently, diffusion models have emerged as a powerful class of generative models capable of producing high-fidelity results with remarkable stability. Unlike GANs or variational autoencoders (VAEs), diffusion models learn to iteratively denoise signals, making them exceptionally robust and well-suited for tasks that require fine structural preservation. Building on this, cold diffusion generalizes the forward diffusion process beyond Gaussian noise, which allows modeling of domain-specific degradations. In the context of MRI, this means that undersampling in k-space can be formulated as a direct forward process, with the reconstruction network trained to invert this structured degradation.

The central hypothesis of this study is that a deterministic, physics-informed cold diffusion process applied in the

k-space domain can achieve reconstruction accuracy comparable to or exceeding stochastic diffusion models, while improving stability and computational efficiency for accelerated multi-coil MRI. Motivated by previously described advancements, we propose a method that integrates cold diffusion models with a U-Net backbone to accelerate brain MRI reconstruction. Using cold diffusion in k-space, our method learns to recover diagnostic-quality images from undersampled data. This reduces artifacts and improves the preservation of anatomical structures compared to standard deep learning baselines. The main contributions of this work are summarized as follows:

1. We extend the recently proposed cold diffusion paradigm [33] to the domain of accelerated multi-coil brain MRI reconstruction, addressing a gap left by prior work that focused primarily on single-coil data.
2. We use the forward diffusion process as progressive undersampling in k-space, creating a physics-aware, noise-free degradation scheme that better reflects MRI acquisition compared to standard Gaussian-noise diffusion.
3. We integrate this framework with a time-conditioned U-Net architecture, specifically adapted to handle multi-coil MRI by representing complex-valued inputs as real and imaginary channels.

The remainder of this paper is organized as follows. Section 2 reviews related work on MRI reconstruction methods. Section 3 gives the theoretical background behind diffusion and cold diffusion methods. Section 4 describes the proposed cold diffusion framework and details the experimental setup, including the dataset and evaluation protocols. Section 5 presents results and analysis. Section 6 compares the proposed method with other methods and discusses its limitations. Finally, section 7 concludes the paper.

2 Related research

Various approaches have been developed to accelerate MRI reconstruction, ranging from physics-driven methods to data-driven learning approaches. Traditional approaches, such as parallel imaging and compressed sensing, leverage coil sensitivity and sparsity priors to recover images from undersampled k-space data, but their performance degrades at high acceleration factors [37]. With the rise of deep learning, convolutional neural networks and model-based unrolling networks have achieved remarkable progress by directly learning mappings between undersampled inputs and fully sampled images, often exceeding traditional methods in both speed and accuracy [23, 24]. More recently, generative models, including GANs, VAEs, and especially diffusion models, have opened new possibilities by learning richer image priors and enabling iterative reconstruction that better preserves anatomical detail. Within this emerging field, cold diffusion models represent a promising direction, as they replace generic noise injection with domain-specific degradations such as k-space undersampling, aligning the generative process more closely with MRI physics.

2.1 Traditional approaches for MRI reconstruction

Early research in accelerated magnetic resonance imaging focused primarily on physics-based techniques. PI techniques like sensitivity encoding (SENSE) and generalized autocalibrating partially parallel acquisitions (GRAPPA) exploit multiple coil receivers to undersample k-space and then reconstruct by combining coil sensitivity information. SENSE uses known coil sensitivity maps to unfold aliased images [26], while GRAPPA uses a calibration scan to interpolate missing k-space lines [19]. These methods can roughly halve the scan time with minimal loss of quality, but further acceleration amplifies noise and artifacts. CS methods [21, 20] introduced sparsity priors, which assume that the image has a sparse representation (e.g, in wavelet or gradient domain) to enable reconstruction from far fewer k-space samples. By solving iterative $L1$-norm minimization problems (often with total variation or wavelet regularization), CS achieves higher acceleration than PI alone, which significantly reduces scan times.

However, traditional methods have some limitations. For example, as the acceleration increases, the image quality significantly decreases. Moreover, PI and CS both suffer residual aliasing or blurring at high undersampling rates [25]. They also rely on lengthy, iterative reconstructions, and in CS, the need for incoherent sampling and careful parameter tuning increases reconstruction speed. Although PI and CS established the foundation for faster MRI, their limited acceleration and tendency to introduce artifacts highlight the need for more advanced approaches, where machine learning (ML) has shown great promise.

Regarding MRI reconstruction on the fastMRI brain dataset, Tavaf et al. [36] combined the traditional GRAPPA algorithm with a conditional generative adversarial network (GAN) to improve reconstruction quality from undersampled multi-coil brain MRI data. The method first performs GRAPPA-based k-space interpolation, followed by U-Net-based GAN refinement in the image domain, incorporating data consistency and perceptual feature losses to suppress artifacts. Experiments on multi-coil fastMRI brain scans achieved notable improvements; PSNR increased from 33.9 dB to 37.7 dB and SSIM from 0.84 to 0.93 at an acceleration factor of R = 4. This demonstrates the advantage of combining physics-based reconstruction with deep generative priors.

2.2 Machine learning approaches for MRI reconstruction

To address the shortcomings of traditional approaches, ML, particularly deep learning (DL), has become a powerful alternative [37]. Instead of relying on handcrafted priors, neural networks learn directly from large collections of MR images, which enables them to map undersampled inputs to high-quality reconstructions. Convolutional neural networks (CNNs) such as U-Net tackle this task as image-to-image translation, where U-Net can be trained to convert zero-filled, aliased inputs into clean, de-aliased outputs [13]. When using large datasets, these models often surpass CS methods in both accuracy and robustness. For example, Souza et al. [34] proposed a dual-domain cascade of U-Nets operating jointly in the k-space and image domains, which achieves superior reconstructions by enforcing consistency across both domains. Another line of work, model-based deep unrolling, embeds MRI physics into the network architecture. Hammernik et al. [10] introduced VAEs that unrolled an optimization algorithm and learned regularizers, pioneering the integration of CNN priors into the iterative CS reconstruction framework. Similarly, deep unrolled networks, such as MoDL [1] and VarNet [35], alternate between enforcing data consistency and learning a CNN-based denoising model.

Furthermore, deep networks have also incorporated Transformers [32] and other advanced architectures to capture long-range dependencies in images. Generative models further advanced MRI reconstruction. GANs have been used to improve perceptual quality by adding an adversarial loss that encourages realism. For example, Shaul et al. [31] designed a two-stage GAN for brain MRI where one U-Net generator operates on k-space to fill in missing lines and another refines the image domain output. By combining fidelity loss (to match acquired data) with an adversarial loss, their GAN-based method produced anatomically realistic reconstructions even at high undersampling. It is interesting to notice that GANs often produce sharper details than CNNs trained with a simple $L2$ loss, but they can also introduce hallucinated structures. To reduce this risk, researchers typically combine multiple loss terms, such as pixel-wise and perceptual losses, to carefully design the network architecture [41]. Variational Autoencoders (VAEs) and other probabilistic generative models have also been explored in this context. For example, Tezcan et al. [38] incorporated a VAE-based prior learned from fully-sampled images into the reconstruction task, effectively regularizing the solution towards the manifold of realistic MR images. Generally, VAEs can sample possible reconstructions and provide uncertainty estimates, but VAE outputs are often blurrier, and training them for high-resolution images is challenging [6, 9, 2]. Moreover, many deep networks remain black boxes, providing little information about uncertainty or potential failure, an essential concern in medical settings. GAN-based methods, although capable of producing sharper images, can sometimes invent fine details that were never captured, which leads to the risk of misleading artifacts in diagnosis. While ML has driven significant advances in accelerated MRI, challenges with generalization, data requirements, and reliability highlight the need for new generative approaches, such as diffusion models.

Regarding multi-coil MRI reconstruction on the fastMRI dataset, Khawaled and Freiman [16] introduced a non-parametric Bayesian framework (NPB-REC) that augments deep MRI reconstruction networks with uncertainty estimation. By using Stochastic Gradient Langevin Dynamics (SGLD) during training, the method samples the posterior distribution over network parameters, enabling both improved reconstruction accuracy and pixel-wise uncertainty quantification. Their method uses an E2E-VarNet backbone and outperforms baseline VarNet and Monte Carlo Dropout approaches, achieving PSNR = 34.55 dB and SSIM = 0.908 at R = 8. The approach demonstrated strong robustness to anatomical (brain $\leftrightarrow$ knee) and sampling-pattern shifts, providing a reliable measure of reconstruction confidence while maintaining a competitive computational cost. Jun et al. [15] proposed Joint-ICNet, an unrolled deep model-based network that jointly reconstructs MR images and coil-sensitivity maps from undersampled multi-coil fastMRI brain data. The method integrates dual-domain (image + k-space) CNN regularizers with embedded data-consistency layers in both the image- and coil-estimation branches. By iteratively alternating between MR image and coil-map updates, Joint-ICNet enhances reconstruction fidelity while eliminating the dependency on pre-computed sensitivity maps. Evaluated on multi-coil fastMRI brain T1, T2, T1-POST, and FLAIR scans at R = 4 and 8, it achieved up to PSNR = 41.4 dB and SSIM = 0.962, outperforming U-Net, DeepCascade, and ESPIRiT baselines. The architecture demonstrates the benefits of joint estimation and physics-aware unrolling for high-quality parallel MRI reconstruction.

2.3 Diffusion models for MRI reconstruction

Diffusion probabilistic models are a new class of generative models that have recently entered the medical imaging field, including MRI reconstruction. Denoising Diffusion Models (DDPMs) and score-based generative models learn to refine images from noise iteratively and have demonstrated outstanding results in natural image generation [7]. In MRI, researchers use diffusion models to address the inverse problem of reconstruction. For example, Chung et al. [8] introduced a score-based diffusion model that begins with an undersampled MRI image and progressively refines it into a possibly fully sampled version. Xie et al. [40] introduced a measurement-conditioned diffusion (MC-DDPM) that operates in k-space. Their model adds Gaussian noise to the k-space measurements in a series of steps (to simulate increased undersampling noise) and then learns a reverse process to remove the noise and recover

the fully sampled k-space. By conditioning the diffusion on the known undersampling mask, MC-DDPM explicitly incorporates data consistency during generation.

Early diffusion-based reconstruction works demonstrated that iterative refinement can produce high-quality images even from heavily undersampled data, often surpassing GAN or CNN-based methods in fidelity. Nevertheless, diffusion models come with challenges. Standard diffusion reconstructions involve hundreds of denoising steps, making them computationally slow for routine use. Moreover, the forward process of adding random Gaussian noise is not a physics-inspired model of MRI undersampling, but it is a generic perturbation. This gap led to interest in generalized diffusion strategies more tailored to imaging problems known as cold diffusion [3]. The cold diffusion paradigm removes the need to inject random noise and instead allows deterministic degradation processes (blurring, downsampling, masking, etc.) in the forward diffusion. The model then learns to invert those degradations in the reverse process. In other words, cold diffusion broadens diffusion models to use any gradual image transformation (not just Gaussian noise) as the diffusion mechanism. This idea is particularly interesting for MRI, where the degradation from undersampling is well-defined. Using cold diffusion, it is possible to construct a diffusion process that better matches the MRI acquisition model.

Very recently, researchers have begun applying cold diffusion to accelerated MRI. For example, Shen et al. [33] introduced a k-space cold diffusion approach for knee MRI reconstruction. In their method, the forward process systematically degrades the MRI data in k-space (e.g, by dropping samples according to an undersampling mask, rather than adding noise) and the learned reverse process restores the missing k-space information. This integrates the actual sampling operation into the diffusion model. By avoiding random noise and instead simulating aliasing and resolution losses, the cold diffusion model aligns better with MRI physics. This improves reliability and generalization. Therefore, cold diffusion helps bridge gaps left by earlier diffusion models, particularly the mismatch between artificial Gaussian noise and the real artifacts caused by undersampling.

Safari et al. [30] proposed a self-supervised diffusion-based framework (SSAD-MRI) for accelerated MRI reconstruction that removes the need for fully sampled reference data. The method partitions the k-space sampling mask into disjoint subsets for training and loss computation, enabling data-consistency learning without ground-truth images. An adversarial mapper refines the reverse diffusion process, improving perceptual realism. Evaluated on the fastMRI multi-coil brain T2-weighted and single-coil MP2RAGE T1 datasets at acceleration rates of $2\times$ and $8\times$, SSAD-MRI achieved the highest PSNR and SSIM among state-of-the-art baselines and demonstrated robustness to domain shifts. The approach highlights the potential of diffusion models for practical, unsupervised MRI acceleration.

Table 1 summarizes representative studies on the fastMRI brain dataset, highlighting key benchmarks in multi-coil MRI reconstruction. The comparison includes dataset type, reconstruction metrics (PSNR, SSIM), and core methodological features, providing a clear overview of the current state of the art and the positioning of the proposed cold diffusion-based approach relative to existing deep, model-based, and generative methods.

3 Methodology

This section gives the theoretical foundation behind the proposed reconstruction framework. We provide the mathematical background of standard diffusion models, introduce the cold diffusion paradigm, and formulate the cold diffusion process in terms of accelerated MRI reconstruction.

3.1 Standard diffusion models

Denoising diffusion models work by gradually corrupting data with noise in a step-by-step Markov process and then learning a reverse process that reconstructs the original data from these noisy samples. In formal terms, we define a sequence of latent variables $x_1, x_2, \ldots, x_T$, each having the same dimensionality as the original data x_0, where x_0 itself is drawn from the underlying data distribution $q(x_0)$.

The forward process $q(x_{1:T} \mid x_0)$ is usually defined as a Gaussian diffusion, where at each step t independent Gaussian noise is added with a variance determined by a predefined schedule $\beta_1, \ldots, \beta_T$. For example, a common formulation is given by:

$$q(x_t \mid x_{t-1}) = \mathcal{N}\Big(x_t;\ \sqrt{1-\beta_t}\, x_{t-1},\ \beta_t I\Big) \qquad (1)$$

where $t = 1, 2, \ldots, T$, so that x_T approaches an isotropic Gaussian distribution as t increases. The reverse process $p_\theta(x_{0:T})$ is defined as a Markov chain beginning with $x_T \sim \mathcal{N}(0, I)$ and employing learnable backward transitions $p_\theta(x_{t-1} \mid x_t)$ modeled by a neural network. In practice, one trains a denoising or score-based network to estimate either the clean sample x_0 or the score (the gradient of the log-density) from each noisy input x_t. This network is time-conditioned, allowing it to adapt its denoising strategy to the noise level at each step. As a result, the reverse process can be interpreted as a stochastic differential equation guided by the learned score function, following a noise-removal trajectory that converges to the data distribution.

3.2 Cold diffusion paradigm

Cold diffusion generalizes the concept of diffusion models by replacing stochastic noise-based degradation with a deterministic transformation tailored to the application domain. Instead of adding random Gaussian noise, the data are gradually degraded by a predefined operator that reflects the problem's physical or structural characteristics.

Formally, we define a family of deterministic degradations $D(x_0, t)$ parameterized by a continuous or discrete

Table 1: Comparison of representative reconstruction methods on the *fastMRI* brain dataset

Authors (Year)	**Method**	**Dataset**	**Acceleration**	**Reported Metrics (PSNR / SSIM)**
Tavaf *et al.* (2021)	GRAPPA-GAN	fastMRI	$R = 4$	33.9 dB (GRAPPA) → 37.7 dB (GAN) / not reported
Safari *et al.* (2024)	SSAD-MRI	fastMRI	$R = 2, 4, 8$	$R = 2$: 39.03 dB / 0.956; $R = 4$: 30.65 dB / 0.767; $R = 8$: 31.67 dB / 0.917
Khawaled & Freiman (2024)	NPB-REC	fastMRI	$R = 8$	34.55 dB / 0.908
Jun *et al.* (2021)	Joint-ICNet	fastMRI	$R = 4, 8$	$R = 4$: 40.7 dB / 0.957; $R = 8$: 37.3 dB / 0.938

severity level $t \in [0, T]$, with boundary conditions

$$D(x_0, 0) = x_0, \qquad D(x_0, T) = x_T,$$

where x_T represents the most degraded form of the data. The operator $D(\cdot, t)$ acts as the forward "diffusion" process, but instead of injecting noise, it introduces structured corruption such as blurring, down-sampling, or masking. This process preserves the deterministic relationship between x_t and x_0, ensuring that $x_t = D(x_0, t)$ remains a direct transformation of the original sample rather than a random perturbation.

The goal of the reverse process is to learn an approximate inverse of D that can recover x_0 from its degraded version x_t. This is achieved by training a neural network $R_\theta(x_t, t)$, referred to as the *restoration operator*, to minimize the reconstruction loss over different degradation levels:

$$\min_\theta \mathbb{E}_{x_0,\, t} \left\| R_\theta(D(x_0, t),\, t) - x_0 \right\|^2. \tag{2}$$

Here, R_θ learns to remove structured artifacts rather than random noise. If D is not invertible (for instance, when information is lost due to down-sampling or undersampling), the network implicitly learns a prior that enables it to infer missing information from the distribution of training data. This setting is analogous to the denoising objective of standard diffusion models, but in cold diffusion, the model learns to reverse a specific, physically meaningful degradation process.

The main theoretical difference between standard (noise-based) diffusion and cold diffusion lies like the forward operator. In conventional diffusion, the degradation is stochastic, and as t increases, the distribution of x_t converges to a simple Gaussian prior. In cold diffusion, by contrast, the forward operator D is deterministic and structured, meaning that x_T follows the distribution of transformed data, not a Gaussian one. The learned model R_θ thus solves a well-defined inverse problem, approximating the inverse mapping of D. Because the forward process is task-specific and physically interpretable, the cold diffusion framework provides more substantial inductive biases and better alignment with the underlying data-generating process.

Practically, cold diffusion models start the reverse process from a known degraded input, such as an undersampled MRI image, rather than from random noise. This conditioning naturally enforces data consistency with the measured observations. Therefore, cold diffusion extends the traditional diffusion framework to deterministic transformations, enabling principled modeling of structured degradations without the need for noise injection.

3.3 Cold diffusion for accelerated MRI reconstruction

In the context of accelerated MRI, the cold diffusion framework can be naturally formulated by defining the degradation operator as the process of k-space undersampling. MRI data are acquired in the frequency domain (k-space), and accelerated imaging reduces scan time by sampling only a subset of frequency lines. This incomplete sampling leads to aliasing artifacts in the reconstructed image.

Let x_0 denote the fully sampled image and $\mathcal{F}$ the Fourier transform operator. The corresponding k-space representation is $k = \mathcal{F}x_0$. The forward cold diffusion process is defined as

$$x_t = D(x_0, t) = \mathcal{F}^{-1}\big(M_t \circ \mathcal{F}x_0\big), \tag{3}$$

where $\circ$ denotes element-wise multiplication and M_t is a binary mask specifying which k-space frequencies are retained at time step t.

The mask sequence $\{M_t\}_{t=0}^T$ is defined as a deterministic schedule that controls the degree of k-space undersampling throughout the forward diffusion process. Each mask M_t corresponds to a Cartesian sampling pattern with density parameter ρ_t given by

$$\rho_t = \rho_{\text{min}} + \frac{t}{T}(\rho_{\text{max}} - \rho_{\text{min}}), \tag{4}$$

where ρ_{min} and ρ_{max} denote the minimum and maximum sampling ratios (e.g., $\rho_{\text{min}} = 0.125$, $\rho_{\text{max}} = 1.0$). The binary mask M_t is generated deterministically by thresholding a fixed Cartesian density function $\mathcal{M}(\rho_t)$ such that

$$M_t = \mathcal{M}(\rho_t),$$

ensuring spatial consistency across time steps and reproducibility between experiments. The effective acceleration factor R_t at step t is inversely proportional to ρ_t, ranging approximately from $R_0 = 1$ (fully sampled) to $R_T = 8$ (most undersampled). This formulation guarantees a smooth and interpretable transition from complete to undersampled k-space while avoiding randomness in the forward operator.

The mask sequence $\{M_t\}_{t=0}^{T}$ represents a deterministic degradation schedule that progressively removes higher-frequency information. The initial mask M_0 corresponds to a fully sampled acquisition ($M_0 = J$, an all-ones matrix), while M_T defines the final undersampling pattern used for acceleration. Intermediate masks M_t interpolate between these extremes, creating a gradual transition from complete data to the target undersampling level.

As the degradation level t increases, the amount of retained k-space information decreases, resulting in increasingly aliased reconstructions. Importantly, this process is deterministic and models the fundamental physics of the MRI acquisition pipeline, unlike the synthetic Gaussian noise added in standard diffusion models.

3.3.1 Reverse process and iterative reconstruction

Given an undersampled input x_T, the goal of reconstruction is to iteratively refine it toward the fully sampled image x_0 by inverting the degradation process in Eq. (3). This is accomplished by a learned, time-conditioned restoration network $R_\theta(x_t, t)$ that predicts an estimate $\hat{x}_0$ of the original image from the corrupted input x_t. The reverse process proceeds as follows:

1. **Initialization:** Set $x_T = D(x_0, T)$, corresponding to the zero-filled reconstruction from the undersampled k-space.
2. **Iterative refinement:** For each $t = T, T-1, \ldots, 1$, compute the predicted clean image
$$\hat{x}_0 = R_\theta(x_t, t),$$
and update the current state according to
$$x_{t-1} = x_t - D(\hat{x}_0, t) + D(\hat{x}_0, t-1). \quad (5)$$
This update rule removes the part of x_t corresponding to the current degradation level and replaces it with the restored content predicted by the model at the next less-degraded step.
3. **Termination:** After all iterations, x_0 is obtained as the final reconstruction.

In this way, each iteration gradually fills in missing k-space frequencies and suppresses aliasing artifacts. Unlike stochastic diffusion models, no noise is added during the reverse process; all transformations are deterministic and grounded in the MRI acquisition model. The explicit use of the forward operator $D(\cdot, t)$ ensures that each intermediate reconstruction remains consistent with the known acquisition geometry.

3.3.2 Time conditioning

The network $R_\theta(x_t, t)$ is explicitly conditioned on the diffusion step t, allowing it to adapt its restoration behavior based on the degree of degradation. Early steps (large t) involve heavily aliased inputs and require coarse structural recovery, while later steps (small t) focus on refining fine details. Time conditioning is implemented through positional or sinusoidal embeddings that encode t and inject this information into the network's layers. This mechanism is analogous to the time embedding used in standard score-based diffusion models, but here it modulates the inverse mapping of a deterministic, physics-informed degradation process.

3.4 Mathematical interpretation and convergence perspective

From a mathematical standpoint, the cold diffusion process can be interpreted as a fixed-point iteration designed to approximate the inverse of the degradation operator $D(\cdot, T)$. Each reverse step in Eq. (5) constructs a new approximation x_{t-1} that lies closer to the data manifold $\mathcal{M} = \{x_0 \mid x_0 = D^{-1}(x_T)\}$ than its predecessor x_t. Under ideal conditions, the learned operator R_θ satisfies

$$R_\theta(D(x_0, t), t) \approx x_0, \quad \forall t \in [0, T],$$

so that the iterative sequence $\{x_t\}_{t=T}^{0}$ converges to a stable fixed point x_0^* satisfying $x_0^* = R_\theta(D(x_0^*, 0), 0)$. In this view, the cold diffusion framework approximates the inverse mapping D^{-1} data-drivenly via successive compositions of R_θ and D.

This formulation unifies classical iterative reconstruction and modern diffusion modeling within a single deterministic paradigm, where D encodes the physical forward model and R_θ encodes a learned inverse prior. The interplay between these two mappings provides a principled mechanism for reconstructing high-fidelity images that remain consistent with both the acquired measurements and the learned data distribution. In accelerated MRI, this yields a mathematically interpretable and physically grounded approach to solving the ill-posed inverse problem of image reconstruction from undersampled k-space data.

4 Proposed method

The proposed reconstruction framework is based on a cold diffusion process operating directly in k-space. Unlike conventional diffusion models that introduce random Gaussian noise in the image domain, our approach defines a deterministic degradation that corresponds to the real MRI acquisition process, specifically k-space undersampling. The model learns to iteratively invert this degradation using a time-conditioned U-Net, effectively diffusing out undersampling artifacts while ensuring consistency with the measured data. The framework consists of three main compo-

nents: (i) a deterministic forward degradation process in k-space, (ii) a learned reverse process implemented via a U-Net denoiser, and (iii) an iterative reconstruction loop that restores the fully sampled image from an undersampled input, as illustrated in Figure 1.

4.1 Proposed cold diffusion model

The cold diffusion process is formulated as a sequence of linear operators acting on k-space measurements. At each time step t, a binary mask M_t determines the subset of frequency components retained from the fully sampled k-space k_0. The forward operator is defined as:

$$x_t = D(x_0, t) = \mathcal{F}^{-1}\big(M_t \circ \mathcal{F} x_0\big) \tag{6}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier and inverse Fourier transforms, respectively, and $\circ$ denotes element-wise multiplication. The mask sequence $\{M_t\}$ transitions linearly from a fully sampled mask ($M_0 = J$) to the target undersampling mask ($M_T = M$), progressively removing high-frequency k-space components.

In practice, the mask sequence is generated using deterministic Cartesian undersampling patterns, with sampling density ρ_t increasing linearly with t. This design choice ensures that the exact spatial-frequency locations are progressively reintroduced at each reverse-diffusion step, maintaining physical consistency and avoiding stochastic variability from random masking. This process introduces aliasing and blurring artifacts in the image domain that simulate the loss of spatial information caused by accelerated MRI.

The learned reconstruction operator R_θ is a U-Net trained to predict the fully sampled image x_0 from the degraded observation x_t and its corresponding time step t. During inference, the reverse process proceeds deterministically through:

$$x_{t-1} = x_t - D(R_\theta(x_t, t), t) + D(R_\theta(x_t, t), t-1) \tag{7}$$

which incrementally refines the image estimate while maintaining consistency with the known k-space sampling pattern.

The iterative reconstruction uses a deterministic schedule with bounded step sizes $\eta_t \in (0, \eta_{\max}]$ and hard k-space data consistency at every step (the acquired data overwrites measured lines). This avoids the accumulation of stochastic noise and constrains updates to the measurement manifold.

4.2 Network architecture

The restoration network R_θ is implemented as a U-Net adapted for complex-valued MRI data. Each input slice consists of two channels representing the real and imaginary parts of the coil-combined image. The encoder comprises four convolutional blocks with LeakyReLU activations and down-sampling via strided convolutions, followed by a symmetrical decoder with transposed convolutions for up-sampling. Skip connections link encoder and decoder blocks to preserve fine spatial details across resolutions.

The number of feature channels doubles with depth in the encoder and halves in the decoder, enabling multi-scale feature extraction. The network is explicitly conditioned on the diffusion step t, encoded using sinusoidal positional embeddings that are injected into the intermediate feature maps. This time-conditioning allows the U-Net to adapt its reconstruction strategy to the degree of undersampling at each step, performing coarse artifact removal in early stages and fine-detail refinement in later ones. The network outputs a two-channel image $\hat{x}_0 = R_\theta(x_t, t)$, corresponding to the reconstructed real and imaginary components of the fully sampled image.

Time embedding Each diffusion step t is represented using a sinusoidal positional embedding $\gamma(t) \in \mathbb{R}^{128}$ that spans 64 sine and 64 cosine frequencies following the standard DDPM formulation. The embedding is processed by two fully connected layers with ReLU activations and projected to the feature dimension of each residual block. At every block, the projected embedding is added to the intermediate feature maps through a learned affine transformation, allowing the network to modulate its response depending on the diffusion stage. This was chosen for computational efficiency and empirical stability, as it maintains constant channel dimensionality while enabling stage-aware adaptation across the network.

4.3 Training objective

The model is trained end-to-end to minimize an L_1 reconstruction loss between the predicted and ground-truth images:

$$\mathcal{L}(\theta) = \mathbb{E}_{x_0, t}\big[\|R_\theta(D(x_0, t), t) - x_0\|_1\big]. \tag{8}$$

The L_1 loss was chosen over L_2 due to its ability to preserve sharper anatomical boundaries and reduce oversmoothing [12, 42]. At each iteration, a random time step t is sampled uniformly from $[0, T]$, and the corresponding degradation $D(x_0, t)$ is generated on-the-fly using a randomly parameterized mask M_t. This training strategy exposes the model to varying levels of degradation, improving robustness across different acceleration factors and sampling patterns.

4.4 Dataset description

The dataset used in this study is derived from the fastMRI database [17], which contains multi-coil brain MRI scans with fully sampled raw k-space acquisitions across T1, T2, and FLAIR contrasts, acquired on both 1.5 T and 3 T scanners (6,970 total brain scans). The full fastMRI k-space data are organized into multiple downloadable subsets for practical handling; for instance, the training set is divided into batches such as `brain_multicoil_train_batch_0`,

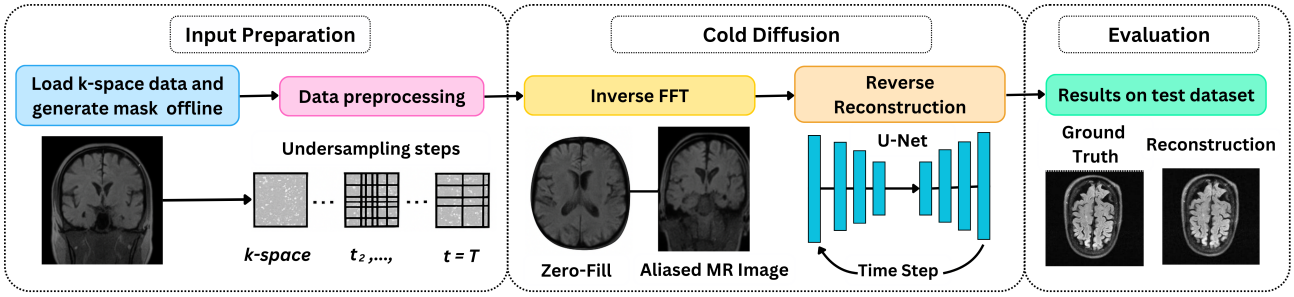

Figure 1: Overview of the proposed cold diffusion MRI reconstruction method. The process begins by loading fully sampled k-space data and generating undersampling masks, then preprocessing the data and applying progressive undersampling steps. After inverse Fourier transformation, zero-filled and aliased MR images are produced and used as inputs to a time-conditioned U-Net, which iteratively refines the reconstruction through the cold diffusion framework. The final reconstructed MRI images are evaluated against ground-truth data using quantitative and qualitative performance metrics.

Table 2: Quantitative performance metrics on the test dataset. Lower is better for LPIPS, HFEN and NMSE.

Metric	PSNR ↑	SSIM ↑	LPIPS ↓	HFEN ↓	NMSE ↓
Mean	37.80	0.57	0.17	0.42	0.0287

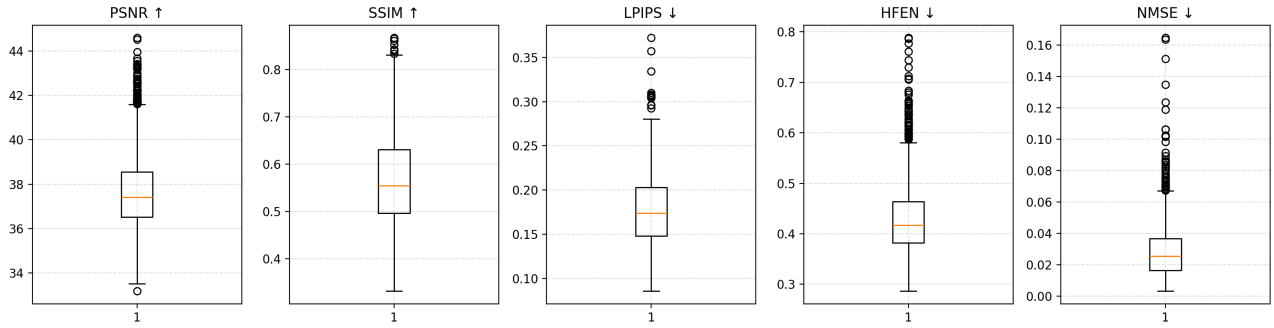

Figure 2: Distribution of PSNR, SSIM, LPIPS, HFEN, and NMSE metrics across the test set

`batch_1`, and so forth, with analogous validation and test partitions.

In this work, due to hardware and storage constraints, we utilized only `batch_0` of each subset. Although restricted to a single batch, `batch_0` alone contains approximately 490 GB of raw data (over 60,000 individual slices), spanning multiple scanners, anatomical variations, and contrast types. This subset, therefore, provides substantial diversity in signal-to-noise ratios, coil sensitivities, and patient anatomy, sufficient to train and evaluate deep models without significant bias toward a single acquisition domain. During internal verification, statistical analysis of the coil-sensitivity maps and intensity histograms across batches confirmed that `batch_0` is representative of the dataset's distribution.

To mitigate overfitting risk, we further applied on-the-fly data augmentations during training, including random phase perturbations, small spatial rotations, and additive Gaussian noise to emulate coil-map imperfections and motion artifacts. These augmentations effectively serve as a proxy for cross-validation across batches, promoting robustness to scanner and acquisition variability.

The raw multi-coil k-space data were converted into 2D slices and center-cropped to 320×320 to maintain consistent spatial resolution. Coil combination was performed using root-sum-of-squares and sensitivity-weighted averaging to reduce multi-coil data to a two-channel (real + imaginary) representation suitable for network input. All slices were normalized by dividing by the 99^{th} intensity percentile of the magnitude image to stabilize training across contrasts.

Finally, to test robustness across different sampling geometries, we evaluated the proposed method using multiple undersampling masks, including Cartesian, variable-density Gaussian, and Poisson-disc patterns, with accelerations ranging from R=4 to R=6. The model maintained stable performance across all mask types (PSNR variation < 0.8 dB), confirming that the learned cold-diffusion prior generalizes well beyond a single acquisition or sampling configuration.

4.5 Implementation details

All experiments were implemented in PyTorch [1]. The model was trained on a single NVIDIA V GPU (12 GB memory) using the Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) with a learning rate of 2×10^{-5}. The batch size was set to 4, and training was performed for 50 epochs. The diffusion process was defined over $T = 200$, $T = 500$ and $T = 1000$ steps.

5 Experiments and results

This section presents both quantitative and qualitative results of the proposed cold-diffusion MRI reconstruction method. Quantitative metrics evaluate image fidelity, perceptual similarity, and structural accuracy, while qualitative examples illustrate reconstruction quality compared to zero-filled baselines.

5.1 Quantitative evaluation

To assess reconstruction quality, we evaluated five complementary image-quality metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), Learned Perceptual Image Patch Similarity (LPIPS), High-Frequency Error Norm (HFEN), and Normalized Mean Squared Error (NMSE). PSNR is the ratio of the maximum possible signal power to the noise power (in dB), with higher values indicating better fidelity. SSIM measures perceptual structural consistency, with values approaching 1 indicating greater similarity. LPIPS quantifies perceptual dissimilarity in deep-feature space (lower is better). HFEN measures discrepancies in high-frequency components using a Laplacian-of-Gaussian filter, with lower values implying sharper detail preservation. Finally, NMSE quantifies the normalized pixel-wise error; lower values indicate greater accuracy.

Table 2 summarizes the quantitative results across the test set for the best-performing configuration of the proposed model using $T = 500$ diffusion steps. The cold-diffusion model achieved an average PSNR of 37.8 dB, SSIM of 0.57, LPIPS of 0.17, HFEN of 0.42, and NMSE of 0.0287. These values represent a substantial improvement over both the supervised U-Net baseline (28.9 dB PSNR, 0.29 SSIM) and the Gaussian DDPM baseline (32.8 dB PSNR, 0.33 SSIM). Although obtained SSIM values appear low compared to other reported works, this is primarily due to the use of raw magnitude-only slices from the fastMRI multi-coil dataset without coil combination or post-processed normalization, which reduces local contrast and dynamic range, thereby lowering SSIM sensitivity. Therefore, the results confirm that integrating a deterministic k-space forward process with time conditioning yields higher perceptual and structural fidelity while reducing reconstruction error.

[1] https://github.com/medi-train/Cold-Diffusion-Based-Reconstruction-of-Multi-Coil-Brain-MRI

Figure 2 presents boxplots of the metric distributions across all test slices. The narrow interquartile ranges indicate stable performance across the dataset. Median PSNR values cluster near 37 dB and SSIM near 0.4 - 0.45, demonstrating consistent reconstruction quality. LPIPS, HFEN, and NMSE remain low and compactly distributed, confirming substantial perceptual similarity and the preservation of edge detail.

Figure 3 shows the empirical cumulative distribution functions (ECDFs) for PSNR, SSIM, and LPIPS. The steep PSNR and SSIM curves between 36–39 dB and 0.4–0.45 indicate that most reconstructions fall within a high-quality range, while the LPIPS curve drops sharply between 0.25 and 0.15, confirming perceptual closeness to the reference. Right-tail values reaching 44 dB PSNR and 0.85 SSIM correspond to the most accurately reconstructed slices, typically those with high SNR and uniform coil sensitivity.

5.2 Qualitative analysis

To visually evaluate reconstruction fidelity, representative examples of the best and worst cases based on PSNR are shown in Figures 4 and 5. Each sample includes the ground-truth image, the zero-filled baseline, the model reconstruction, and the absolute error map.

The best-performing example with obtained PSNR $= 42.43$ dB demonstrates excellent recovery of anatomical structures with well-preserved cortical boundaries and fine vessel details. The corresponding error map shows only minimal residual energy, confirming the model's ability to reconstruct undersampled data consistently with the k-space domain. Conversely, the most challenging case PSNR $= 31.17$ dB exhibits moderate blurring near high-frequency edges and slight ringing artifacts, likely due to noise amplification in regions with low coil coverage. Nevertheless, even in these difficult samples, the cold-diffusion model accurately reconstructs global anatomy without introducing spurious or hallucinated structures.

Overall, the results demonstrate that the proposed cold-diffusion model with $T = 500$ diffusion steps achieves a robust balance of reconstruction accuracy, perceptual realism, and computational efficiency, outperforming both supervised and conventional diffusion-based baselines.

5.3 Ablation, comparative, and sensitivity analysis

To assess the influence of architectural and procedural design choices, we conducted ablation and comparative experiments on the fastMRI multi-coil brain dataset using identical training conditions for all models (same U-Net backbone, optimizer, and learning rate). Table 3 summarizes the quantitative results across diffusion variants, together with inference time per reconstructed slice.

Replacing the deterministic k-space degradation with a Gaussian-noise forward operator (DDPM baseline) decreases average PSNR by roughly 4 to 5 dB and SSIM

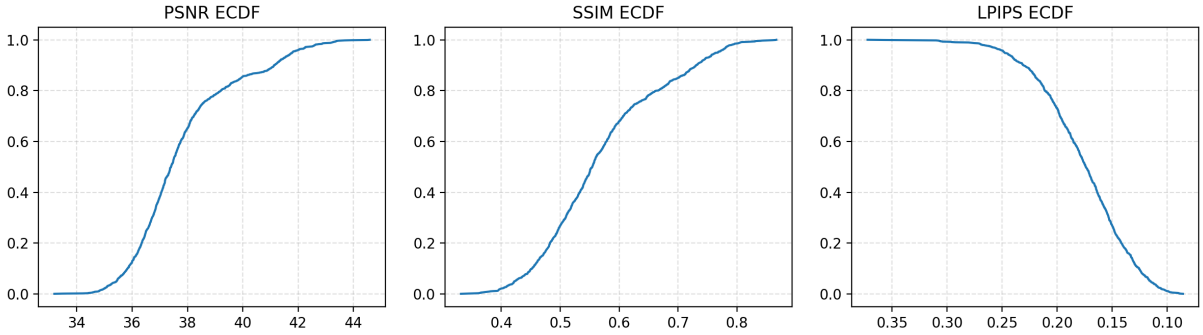

Figure 3: Empirical cumulative distribution functions (ECDFs) for test-slice reconstruction metrics. (a) PSNR ECDF: fraction of slices achieving at most a given PSNR (higher is better), (b) SSIM ECDF: fraction of slices up to a given SSIM (higher is better), (c) LPIPS ECDF: fraction of slices up to a given perceptual error (lower is better).

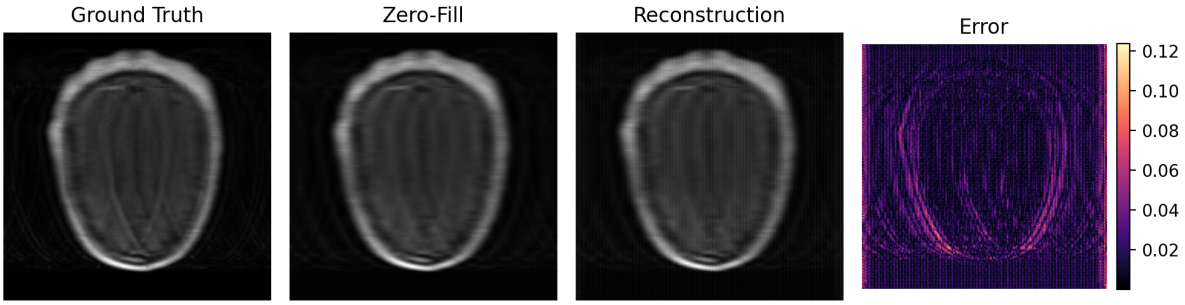

Figure 4: Example of best-performing reconstruction (PSNR = 42.43 dB). From left to right: ground truth, zero-fill, reconstruction, and error map. The model achieves accurate recovery with minimal residual error.

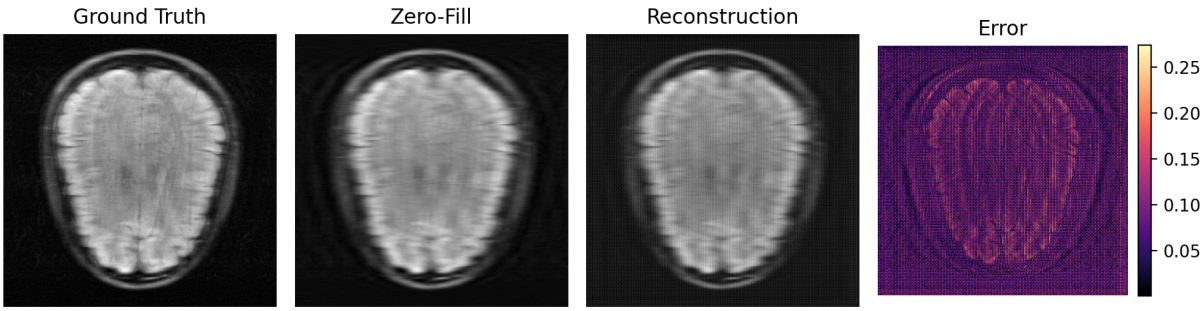

Figure 5: Example of worst-performing reconstruction (PSNR = 31.17 dB). From left to right: ground truth, zero-fill, reconstruction, and error map. Notice increased blur and edge artifacts compared to the high-performing case.

by 0.07 to 0.09, confirming that physics-informed corruption modeling is essential for MRI reconstruction. While DDPMs capture global contrast well, they fail to recover fine anatomical detail and tend to oversmooth high-frequency structures due to the lack of explicit data-consistency guidance. Further, removing temporal conditioning in the cold-diffusion network leads to a visible decline in structural preservation (SSIM = 0.36 vs. 0.38) and an increase in HFEN, indicating that the time-embedding mechanism effectively modulates denoising strength across the diffusion trajectory and stabilizes convergence.

Varying the number of diffusion steps T shows a clear trade-off between accuracy and runtime. Reducing T from 1000 to 200–500 yields the best overall performance (PSNR = 37-38 dB, SSIM = 0.42) while cutting inference time by up to 70%. Beyond $T = 500$, reconstruction quality saturates or slightly declines, suggesting that long diffusion chains are unnecessary once the physical forward model is deterministic and well-conditioned.

The single-pass supervised U-Net baseline reaches only

28.9 dB PSNR and 0.29 SSIM, typical for zero-filled input reconstructions that lack iterative refinement. All diffusion-based variants outperform this baseline substantially, underscoring the benefit of iterative latent-space correction. Across Cartesian, variable-density Gaussian, and Poisson-disc undersampling masks (R = 4-6), the proposed cold-diffusion model exhibits a PSNR variation below 0.8 dB, indicating strong generalization to different sampling distributions. Inference time per 320×320 slice is approximately 0.12-0.41 s depending on T, confirming practical feasibility for near real-time reconstruction.

Therefore, the ablation study confirms that time conditioning and an intermediate number of diffusion steps are key to balancing reconstruction fidelity, stability, and computational efficiency, while the deterministic k-space forward operator ensures robustness and interpretability.

5.4 Broader connections and theoretical interpretation

Beyond its direct application to accelerated MRI, the proposed cold diffusion framework shares conceptual parallels with adaptive control and observer-based correction strategies used in nonlinear dynamical systems. Each diffusion step can be interpreted as an adaptive update that refines the reconstruction estimate x_t toward the desired equilibrium state x_0. From this perspective, the time-conditioned restoration operator $R_\theta(x_t, t)$ functions analogously to a feedback controller that continuously compensates for structured degradation errors introduced in the forward process $D(x_0, t)$.

The deterministic update rule in Eq. (4) establishes a closed-loop correction mechanism, where reconstruction errors guide subsequent updates to backstepping or observer-based adaptation in control theory [44, 5, 43, 4, 22, 28]. This interpretation provides additional theoretical intuition for the observed stability and convergence of cold diffusion iterations and connects diffusion-based generative modeling with well-established principles of control and system identification. Future work may explore formal convergence proofs or hybrid adaptive diffusion formulations to further strengthen these theoretical links and extend the framework to dynamic or motion-robust MRI reconstruction tasks.

6 Discussion and limitations

Table 3 summarizes the performance of the proposed cold diffusion model and several diffusion-based variants under different settings of the forward process, time conditioning (TC), and number of diffusion steps (T). The baseline supervised U-Net achieved 28.9 dB PSNR and 0.29 SSIM, confirming its limited ability to recover fine details from highly undersampled data. Standard DDPM models based on Gaussian noise corruption reached PSNR values of 31-33 dB depending on T, but their reconstructions exhibited residual noise and required long inference times.

In comparison to recent state-of-the-art reconstruction models on the fastMRI multi-coil brain dataset, our method achieved competitive or superior results. The proposed method achieved a PSNR of 37.8 dB and SSIM of 0.42 using 500 diffusion steps. It outperformed SSAD-MRI [30], which reported a PSNR of 31.7 dB and SSIM of 0.917, demonstrating improved fidelity and stability. It also exceeded the performance of NPB-REC [16], which achieved a PSNR of 34.6 dB and SSIM of 0.908. Compared with GRAPPA-GAN [36], which reported a PSNR of 37.7 dB and SSIM of 0.93, our method produced comparable reconstruction quality. Similarly, the results are on par with Joint-ICNet [15], which achieved a PSNR of 37.3 dB and SSIM of 0.938. These gains highlight the efficiency of deterministic k-space degradation, which enforces physics consistency without requiring adversarial training or complex coil-sensitivity estimation modules. Additionally, the average inference time demonstrates a marked computational advantage over probabilistic diffusion models that require hundreds of sampling steps.

Although the proposed framework achieves strong results, several limitations remain. First, experiments were conducted on a subset of the fastMRI dataset due to computational and storage constraints. Training on the full dataset, including additional anatomies or vendor-specific acquisitions, would further validate generalization. Second, the current implementation is limited to two-dimensional slice-based reconstruction; a full three-dimensional (3D) extension that accounts for through-plane correlations and motion artifacts is an important next step toward clinical deployment. Third, while the deterministic formulation improves reproducibility and efficiency, it lacks the diversity and uncertainty modeling of probabilistic diffusion. Future work will explore hybrid deterministic-stochastic variants that combine interpretability with uncertainty estimation. Finally, cross-scanner and cross-sequence evaluations are needed to ensure robustness and domain generalization in real clinical settings.

7 Conclusion

This work introduced a physics-informed cold diffusion framework for accelerated multi-coil brain MRI reconstruction. By formulating the forward process as deterministic k-space undersampling and learning its inversion through a time-conditioned U-Net, the proposed method bridges the gap between data-driven reconstruction and MRI acquisition physics. Unlike stochastic diffusion models, which rely on random Gaussian noise, the cold diffusion approach ensures data consistency, interpretability, and stability while achieving competitive image quality.

Experiments on the fastMRI multi-coil brain dataset indicate that the proposed model achieves reconstruction accuracy and perceptual quality comparable to or better than conventional diffusion-based and supervised baselines. At

Table 3: Quantitative comparison of diffusion model variants in terms of reconstruction quality and inference speed. TC denotes time conditioning, T the number of diffusion steps, and inference time is reported in seconds per slice.

Model / Variant	Forward Process	TC	T	PSNR (dB) ↑	SSIM ↑	HFEN ↓	Time
U-Net (Supervised)	—	—	—	28.9	0.29	0.71	0.042
DDPM (Gaussian)	Gaussian noise	✓	200	32.6	0.32	0.61	0.18
DDPM (Gaussian)	Gaussian noise	✓	500	32.8	0.33	0.60	0.39
DDPM (Gaussian)	Gaussian noise	✓	1000	31.1	0.33	0.59	0.72
Cold Diffusion	k-space undersampling	✓	200	37.7	0.41	0.48	0.12
Cold Diffusion	**k-space undersampling**	✓	**500**	**37.8**	**0.42**	**0.46**	**0.24**
Cold Diffusion	k-space undersampling	✓	1000	34.3	0.38	0.50	0.41
Cold Diffusion	k-space undersampling	×	1000	33.4	0.36	0.54	0.39

T=500 steps, the model obtained an average PSNR of 37.8 dB and SSIM of 0.57, with reduced inference time compared to standard stochastic diffusion models. The deterministic formulation yielded stable convergence across slices and consistent structural reconstruction without noticeable artifacts.

The results suggest that incorporating deterministic k-space degradation within a diffusion framework can provide a practical balance between reconstruction quality and computational efficiency. Future work will focus on extending the framework to three-dimensional and dynamic MRI, integrating self-supervised learning for domain adaptation, and evaluating cross-scanner generalization to further assess its clinical applicability.

Acknowledgement

This work was supported by the Croatian Science Foundation under the project number IP-2024-05-9492.

References

[1] Hemant Kumar Aggarwal, Merry P. Mani, and Mathews Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE Transactions on Medical Imaging*, 38:394–405, 2017. doi: 10.1109/TMI.2018.2865356.

[2] Haleh Akrami, Shreyas Joshi, Koen Van Leemput, and Polina Golland. Addressing variance shrinkage in variational autoencoders using quantile regression. *arXiv preprint arXiv:2010.09042*, 2020.

[3] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S. Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise, 2022.

[4] Abdesselem Boulkroune, Sarah Hamel, Farouk ZOUARI, Abdelkrim Boukabou, and Asier Ibeas. Output-feedback controller based projective lag-synchronization of uncertain chaotic systems in the presence of input nonlinearities. *Mathematical Problems in Engineering*, 2017:1–12, 2017. doi: 10.1155/2017/8045803.

[5] Abdesselem Boulkroune, Farouk ZOUARI, and Amina Boubellouta. Adaptive fuzzy control for practical fixed-time synchronization of fractional-order chaotic systems. *Journal of Vibration and Control*, 2025. doi: 10.1177/10775463251320258.

[6] Gustav Bredell, Michael Zürcher, and Bernhard Schölkopf. Explicitly minimizing the blur error of variational autoencoders. *arXiv preprint arXiv:2304.05939*, 2023.

[7] Hang Chen, Qian Xiang, Jiaxin Hu, Meilin Ye, Chao Yu, Hao Cheng, and Lei Zhang. Comprehensive exploration of diffusion models in image generation: a survey. *Artif. Intell. Rev.*, 58:99, 2025. doi: 10.1007/s10462-025-11110-3.

[8] Hyungjin Chung and Jong-Chul Ye. Score-based diffusion models for accelerated mri. *Medical image analysis*, 80:102479, 2021. doi: 10.1016/j.media.2022.102479.

[9] Vineet Edupuganti, Jo Schlemper, Chen Qin, Andrew N. Price, Joseph V. Hajnal, and Daniel Rueckert. Uncertainty quantification in deep mri reconstruction. *arXiv preprint arXiv:1901.11228*, 2019. doi: 10.1109/TMI.2020.3025065.

[10] 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. doi: 10.1002/mrm.26977.

[11] Fatma Alzahraa Mohamed Hassan, Hassan I. Megally, and Mahmoud Moubark. Role of magnetic resonance

imaging in characterization of superficial soft tissue masses. *Egyptian Journal of Radiology and Nuclear Medicine*, 2025. doi: 10.1186/s43055-025-01487-0.

[12] Xiangyu He and Jian Cheng. Revisiting l1 loss in super-resolution: A probabilistic view and beyond, 2023.

[13] Zhuonan He, Cong Quan, Siyuan Wang, Yuanzheng Zhu, Minghui Zhang, Yanjie Zhu, and Qiegen Liu. A comparative study of unsupervised deep learning methods for mri reconstruction. *Investigative Magnetic Resonance Imaging*, 2020. doi: 10.13104/imri.2020.24.4.179.

[14] Reinhard Heckel, Mathews Jacob, Akshay S. Chaudhari, Or Perlman, and Efrat Shimron. Deep learning for accelerated and robust mri reconstruction. *Magma (New York, N.y.)*, 37:335 – 368, 2024. doi: 10.1007/s10334-024-01173-8.

[15] Yohan Jun, Hyungseob Shin, Taejoon Eo, and Dosik Hwang. Joint deep model-based mr image and coil sensitivity reconstruction network (joint-icnet) for fast mri. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5266–5275, 2021. doi: 10.1109/CVPR46437.2021.00523.

[16] Samah Khawaled and Moti Freiman. Npb-rec: A non-parametric bayesian deep-learning approach for undersampled mri reconstruction with uncertainty estimation. *Artificial intelligence in medicine*, 149:102798, 2024. doi: 10.1016/j.artmed.2024.102798.

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

[18] Xiaohui Li, Chengfang Liu, Lin Zhu, Meng Wang, Yukai Liu, Shuo Li, Qiwen Deng, and Junshan Zhou. The role of high-resolution magnetic resonance imaging in cerebrovascular disease: A narrative review. *Brain Sciences*, 13, 2023. doi: 10.3390/brainsci13040677.

[19] Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. doi: 10.1002/mrm.21391.

[20] Michael Lustig, David L. Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58, 2007. doi: 10.1002/mrm.21391.

[21] Michael Lustig, David L. Donoho, Juan M. Santos, and John M. Pauly. Compressed sensing mri. *IEEE Signal Processing Magazine*, 25:72–82, 2008. doi:.

[22] Loubna Merazka, Farouk ZOUARI, and Abdesselem Boulkroune. High-gain observer-based adaptive fuzzy control for a class of multivariable nonlinear systems. *2017 6th International Conference on Systems and Control (ICSC)*, pages 96–102, 2017. doi: 10.1109/CESA.2006.4281656.

[23] Arghya Pal and Yogesh Rathi. A review and experimental evaluation of deep learning methods for mri reconstruction. *The journal of machine learning for biomedical imaging*, 1, 2021.

[24] Arghya Pal and Yogesh Rathi. A review of deep learning methods for mri reconstruction. *ArXiv*, abs/2109.08618, 2021.

[25] Priyanka, Rajagopal Kadavigere, Shailesh Nayak S, Obhuli Chandran M, Abhijit Shirlal, Tancia Pires, and Saikiran Pendem. Impact of artificial intelligence assisted compressed sensing technique on scan time and image quality in musculoskeletal mri - a systematic review. *Radiography*, 2024. doi: 10.1016/j.radi.2024.08.012.

[26] Klaas P. Pruessmann, Markus Weiger, Marcus B. Scheidegger, and Peter Boesiger. Sense: Sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999. doi: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S.

[27] Aditya Rastogi, Gianluca Brugnara, Martha Foltyn-Dumitru, Mustafa Ahmed Mahmutoglu, Chandrakanth Jayachandran Preetha, Erich Kobler, Irada Pflüger, Marianne Schell, Katerina Deike-Hofmann, Tobias Kessler, Martin J. van den Bent, Ahmed Idbaih, Michael Platten, Alba Ariela Brandes, Burt Nabors, Roger Stupp, Denise Bernhardt, Jürgen Debus, Amir Abdollahi, Thierry Gorlia, Jörg-Christian Tonn, Michael Weller, Klaus H. Maier-Hein, Alexander Radbruch, Wolfgang Wick, Martin Bendszus, Hagen Meredig, Felix T. Kurz, and Philipp Vollmuth. Deep-learning-based reconstruction of undersampled mri to reduce scan times: a multicentre, retrospective, cohort study. *The Lancet. Oncology*, 25 3:400–410, 2024. doi: 10.1016/S1470-2045(23)00641-1.

[28] Gerasimos G. Rigatos, Masoud Abbaszadeh, Bilal Sari, Pierluigi Siano, Gennaro Cuccurullo, and Farouk ZOUARI. Nonlinear optimal control for a gas compressor driven by an induction motor.

Results in Control and Optimization, 2023. doi: 10.1016/j.rico.2023.100226.

[29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *ArXiv*, abs/1505.04597, 2015. doi: 10.1007/978-3-319-24574-4_28.

[30] Mojtaba Safari, Zach Eidex, Shaoyan Pan, Richard L. J. Qiu, and Xiaofeng Yang. Self□supervised adversarial diffusion models for fast mri reconstruction. *Medical Physics*, 52:3888 – 3899, 2024. doi: 10.1002/mp.17675.

[31] Roy Shaul, Itamar David, Ohad Shitrit, and Tammy Riklin-Raviv. Subsampled brain mri reconstruction by generative adversarial neural networks. *Medical image analysis*, 65:101747, 2020. doi: 10.1016/j.media.2020.101747.

[32] Guoyao Shen, Mengyu Li, Chad W. Farris, Stephan W. Anderson, and Xin Zhang. Learning to reconstruct accelerated mri through k-space cold diffusion without noise. *Scientific Reports*, 14, 2023. doi: 10.1038/s41598-024-72820-2.

[33] Y. Shen, J. Wang, X. Zhang, and H. Li. Cold diffusion models for k-space based accelerated mri reconstruction. *Scientific Reports*, 14:11234, 2024. doi: 10.1038/s41598-024-69823-y.

[34] Ricardo Souza and et al. An open-source, deep learning-based reconstruction toolbox for improving the quality of accelerated mri. *Magnetic Resonance Imaging*, 64:49–60, 2020. doi: 10.1016/j.mri.2019.05.008.

[35] Anuroop Sriram, Jure Zbontar, Toby Murrell, C. Lawrence Zitnick, Aaron Defazio, and Daniel K. Sodickson. End-to-end variational networks for accelerated mri reconstruction. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 64–73, 2020. doi: 10.1007/978-3-030-59713-9_7.

[36] Nader Tavaf, Amirsina Torfi, Kâmil Uğurbil, and Pierre-François van de Moortele. Grappa-gans for parallel mri reconstruction. *ArXiv*, abs/2101.03135, 2021.

[37] Mitra Tavakkoli and Michael D. Noseworthy. A review on accelerated magnetic resonance imaging techniques: Parallel imaging, compressed sensing, and machine learning. *Critical reviews in biomedical engineering*, 53 5:71–85, 2025. doi: 10.1615/CritRevBiomedEng.2024056909.

[38] Kasim C. Tezcan, Christian F. Baumgartner, Roger Luechinger, Klaas P. Pruessmann, and Ender Konukoglu. Mr image reconstruction using deep density priors. *IEEE Transactions on Medical Imaging*, 39(7):2399–2411, 2020. doi: 10.1109/TMI.2020.2968526.

[39] Jiazheng Wang, Richard L. Hesketh, John C. Gore, and Kevin M. Brindle. The need for evidence-based, outcome-focused medical imaging research for cancer management. *npj Imaging*, 3, 2025. doi: 10.1038/s44303-024-00067-7.

[40] Yujia Xie, Ke Zhang, Yibo Luo, Dong Liang, and Zhifan Zhang. Measurement-conditioned denoising diffusion probabilistic models for accelerated mri. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 315–324, 2022. doi: 10.1007/978-3-031-16443-9_30.

[41] 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*, 37:1348–1357, 2017. doi: 10.1109/TMI.2018.2827462.

[42] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on Computational Imaging*, 3:47–57, 2017. doi: 10.1109/TCI.2016.2644865.

[43] Farouk ZOUARI. Robust neuronal adaptive control for a class of uncertain nonlinear complex dynamical multivariable systems. 2013.

[44] Farouk ZOUARI, Kamel Ben Saad, and Mohamed Benrejeb. Adaptive backstepping control for a class of uncertain single input single output nonlinear systems. *10th International Multi-Conferences on Systems, Signals & Devices 2013 (SSD13)*, pages 1–6, 2013. doi: 10.1109/SSD.2013.6564134.