International Journal of Innovation Studies

DIFFUSION MODELS FOR MEDICAL IMAGE RECONSTRUCTION AND DENOISING IN LOW-RESOURCE SETTINGS

Dr. Geetha K[1], Dr. M.Charles Arockiaraj[2], Prof.Gunasekaran K[3], Prof. Roshan Vegas[4], Dr.S.Dilliarasu[5]

[1]Associate professor St. Joseph University Palanjur, Nazarethpet p chennai – 600123
[2]Associate professor AMC Engineering College Bangalore-560083, Karnataka
[3]Assistant professor AMC Engineering College Bangalore-560083, Karnataka
[4]Assistant professor AMC Engineering College Bangalore-560083, Karnataka
[5]Assistant professor Dharmamurthi Rao Bahadur Calavala Cunnan Chetty's Hindu College, Chennai - 600 072

Abstract

Medical imaging plays a vital role in diagnosis and treatment planning, but access to high-quality imaging technologies and reliable data acquisition pipelines remains limited in low-resource settings. Recent advances in generative models, especially diffusion models, offer new possibilities for robust medical image reconstruction and denoising even under severe data constraints. In this paper, we investigate the application of diffusion probabilistic models for medical image reconstruction and denoising tasks, with a focus on scenarios involving low signal-to-noise ratios (SNR), sparse data acquisition, and limited computational infrastructure. We demonstrate that these models can produce high-fidelity reconstructions from degraded or incomplete inputs and outperform conventional methods and other deep learning baselines in multiple medical imaging modalities including MRI, CT, and ultrasound. We further propose lightweight and compressed diffusion model architectures tailored for deployment in low-resource clinical environments.

Keywords: Diffusion Models, Denoising Diffusion Probabilistic Models (DDPM), Medical Image Reconstruction, Image Denoising, Deep Generative Models, Inverse Problems in Imaging

1. Introduction

Medical imaging is indispensable in modern diagnostics, yet millions lack access to high-quality imaging due to resource constraints. These challenges are particularly acute in low- and middle-income countries (LMICs), where limitations in hardware availability, power supply, and trained radiologists create a pressing need for efficient image enhancement and reconstruction tools.

Traditional image reconstruction methods often fail under conditions of sparse or noisy data acquisition. Deep learning approaches, while promising, typically require extensive computational resources and large labeled datasets for training. Diffusion models—recently gaining traction in computer vision—have shown superior generative capacity and robustness in high-noise regimes. This work explores how diffusion models can bridge the quality-access gap in medical imaging under constrained settings.

2. Background and Related Work

2.1 Medical Image Reconstruction and Denoising

Reconstruction from undersampled or noisy measurements is a fundamental problem in modalities such as Magnetic Resonance Imaging (MRI), Computed Tomography (CT), and ultrasound. Conventional approaches include compressed sensing, total variation minimization, and iterative reconstruction algorithms. Deep learning-based solutions have also emerged, utilizing convolutional autoencoders, UNets, and GANs.

2.2 Diffusion Models

Diffusion models are a class of generative models that learn to reverse a Markovian noising process. They have shown strong performance in image synthesis, inpainting, and super-resolution. In medical imaging, preliminary works (e.g., Med-DDPM, CMDDPM) have applied diffusion models for MRI reconstruction, showing improved detail preservation and robustness to noise.

2.3 Imaging in Low-Resource Settings

LMICs often face poor SNR imaging due to substandard devices, lack of maintenance, and power fluctuations. This necessitates reconstruction models that generalize well and can operate with minimal computing power. However, current high-performing generative models are often too large for practical deployment in such settings.

3. Methodology

3.1 Problem Definition

Given an observed medical image x0x_0 degraded by noise or missing data (e.g., under sampling in k-space for MRI), our goal is to recover the clean image xx using a diffusion model DθD_\theta, trained to denoise progressively corrupted versions of the data.

3.2 Diffusion Process

We implement a standard denoising diffusion probabilistic model (DDPM) with a forward process defined by:

q(x/|xt−1)=N(xt;1−βtxt−1,βtI)q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I)

and a learned reverse process:

pθ(xt−1|xt)=N(xt−1;μθ(xt,t),Σθ(xt,t))p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))

3.3 Conditional Sampling for Reconstruction

For reconstruction tasks (e.g., from sparse or masked inputs), we employ conditional sampling by initializing the noising process with degraded images and guiding sampling using known measurements (e.g., k-space lines or CT projections) as constraints.

3.4 Model Compression and Lightweight Deployment

To enable usage in low-resource devices, we compress the diffusion model using techniques such as:

- Knowledge distillation
- Weight pruning
- Quantization-aware training

We also design a shallow UNet-based noise predictor with fewer parameters and optimized for ARM-based edge devices.

Denoising diffusion probabilistic models

Denoising Diffusion Probabilistic Models (DDPMs)[11] typically contain two Markov chains: a forward Markov chain and a reverse Markov chain. Figure 2 shows the process of adding noise and denoising using DDPM. The forward chain adds noise to the images, while the inverse chain removes noise. The core work of DDPMs is to train a neural network to learn the data distribution of the training dataset and generate new data.

Process of adding noise and denoising using DDPM. The process of adding noise uses a forward Markov chain to gradually add noise to the original image until the original image becomes purely noisy. The process of denoising uses a reverse Markov chain to gradually denoise the image until the image is restored to its original state.

The forward process involves gradually adding Gaussian noise to the original data until the data structure is corrupted and becomes random noise. Specifically, given the original data x0x0, its distribution is denoted as x0∼q(x)x0∼q(x). Since the memoryless nature of the Markov chain, the probability distribution of the next state xtxt in the diffusion process can only be determined by the current state xt−1xt−1, i.e., xtxt and xt−1xt−1 satisfy the following relation:

(1)

xt=1−βt√xt−1+βt−−√εxt=1−βtxt−1+βtε

where βtβt increases with timestep tt. There can be several choices for the distribution of the noise εε, DDPMs use the standard normal distribution, denoted as ε∼N(0,1)ε∼N(0,1).

Since the noise adding process of DDPMs follows a Gaussian distribution, the process from xt−1xt−1 to xtxt can be described as

(2)

q(xt|xt−1)=N(xt;1−βt√xt−1,βtI)q(xt|xt−1)=N(xt;1−βtxt−1,βtI)

Equation (2) is likewise the most common choice for transition kernels, in which 1−βt√xt−11−βtxt−1 is the mean and βtβt is the variance. II represents the identity matrix, which is used to ensure that the noise is independent and has the same variance across all dimensions. In addition, if we make αt=1−βtαt=1−βt and α¯t=∏ti=1αiα¯t=∏i=1tαi, in Eq. (1) by the reparameterization trick, the noise data xtxt at timestep tt can be expressed as

(3)

xt=α¯t−−√x0+1−α¯t−√εxt=α¯tx0+1−α¯tε

This means that we can introduce the noise data xtxt for any step only if we have x0x0, and have determined βtβt for each step.

The reverse process means that the real samples are generated by gradually denoising the noised data until the original data structure is restored. Specifically, given the noise data xtxt, its distribution is denoted as xt∼N(0,1)xt∼N(0,1). Then, the DDPMs denoise the data using the reverse Markov chain.

4. Experiments

This section evaluates the proposed diffusion model framework for medical image reconstruction and denoising across multiple imaging modalities and degradation scenarios, with an emphasis on low-resource applicability. We benchmark against state-of-the-art methods and evaluate both quantitative performance and computational efficiency.

4.1 Datasets

We conduct experiments on publicly available datasets that simulate real-world low-resource conditions:

- **FastMRI (knee subset)** – T1-weighted MR images with retrospectively applied undersampling (4× and 8× Cartesian masks).
- **LoDoPaB-CT** – Low-dose CT dataset based on the LIDC/IDRI database with simulated 25% radiation dose.
- **POCUS Ultrasound Dataset** – Point-of-care ultrasound images annotated for pathology classification, modified with synthetic speckle noise to simulate handheld low-cost device outputs.

For all datasets, images are center-cropped and resized to 128×128 resolution to match low-resolution acquisition setups typical in under-resourced clinics.

4.2 Baselines

We compare our diffusion-based approach to several widely-used denoising and reconstruction methods:

- **Total Variation (TV) Minimization**
- **Compressed Sensing MRI (CS-MRI)** using L1-wavelet regularization
- **UNet-based Autoencoders** trained for denoising/reconstruction
- **GAN-based Image-to-Image Translation** models (e.g., pix2pix)
- **Score-based Generative Models** (e.g., Score-SDE)

For fairness, all deep learning baselines are trained using the same dataset splits and augmentations.

4.3 Evaluation Metrics

We use standard metrics for evaluating reconstruction quality and model efficiency:

- **PSNR (Peak Signal-to-Noise Ratio)** – Measures image fidelity.
- **SSIM (Structural Similarity Index Measure)** – Assesses perceptual similarity.
- **MAE (Mean Absolute Error)** – Pixel-wise difference.
- **Inference Time (seconds/image)** – Measures runtime efficiency.
- **Model Size (MB)** – Measures suitability for edge deployment.

4.4 Implementation Details

- **Training:** All models are trained using the AdamW optimizer with an initial learning rate of 10−410^{-4}10−4, a batch size of 32, and a cosine decay schedule over 100,000 steps.
- **Diffusion Steps:** We use 100 steps during training and 25 steps for inference with DDIM acceleration.
- **Hardware:** Training is conducted on an NVIDIA RTX 3090 GPU. Low-resource benchmarks are performed on a Raspberry Pi 4 (8 GB) and NVIDIA Jetson Nano.

4.4 Results

Our approach consistently outperforms baseline models in terms of PSNR/SSIM under noisy and incomplete inputs. We demonstrate:

- **MRI:** 3–5 dB PSNR gain over compressed sensing with 4× undersampling
- **CT:** Enhanced edge preservation under 25% radiation dose conditions
- **Ultrasound:** Robust despeckling with minimal feature distortion

Compressed versions of our models achieved inference speeds under 1s on Raspberry Pi 4 and similar ARM-based edge devices.

Method	PSNR (4×)	SSIM (4×)	PSNR (8×)	SSIM (8×)	Inference Time (s)
TV	28.3 dB	0.812	24.1 dB	0.722	1.20
CS-MRI	30.1 dB	0.865	26.7 dB	0.798	1.35
UNet	32.5 dB	0.901	28.3 dB	0.840	0.15
Score-SDE	33.0 dB	0.918	29.1 dB	0.856	2.40
Ours (Diffusion)	**33.6 dB**	**0.926**	**29.7 dB**	**0.866**	0.95 (Jetson Nano)

CT Denoising (LoDoPaB)

Method	PSNR	SSIM	Inference Time	Notes
TV	30.8	0.81	1.5 s	Edge blurring
UNet	33.5	0.88	0.2 s	Over-smoothing
GAN	34.1	0.89	0.3 s	Artifacts under noise
Ours	**35.4**	**0.91**	1.1 s	Sharp structures retained

Ultrasound Denoising (POCUS)

Method	PSNR	SSIM	Inference Time	Visual Quality
Median Filter	27.5	0.73	0.01 s	Blurry details
UNet	29.4	0.81	0.2 s	Acceptable
Ours	**30.8**	**0.86**	0.9 s	Preserved speckle structure

5. Discussion

The results show diffusion models' strength in capturing fine structural details and resilience to noise. Their iterative nature aligns well with reconstruction from partial measurements. However, training remains compute-intensive, and fast sampling remains a bottleneck. Further work on accelerating inference and learning from fewer examples is needed.

5.1 Reconstruction Quality

Across all tested modalities—MRI, CT, and ultrasound—our diffusion-based approach consistently outperformed traditional and deep learning baselines in both quantitative metrics (PSNR, SSIM) and visual fidelity. Notably:

- **MRI reconstruction** at 4× and 8× undersampling rates demonstrated clear superiority in preserving fine anatomical structures without introducing artifacts common in GAN-based methods.
- **CT denoising** at low-dose levels showed enhanced clarity of soft tissue boundaries and reduced streak artifacts, validating the robustness of the model to extreme noise.
- **Ultrasound despeckling** retained clinically relevant textures that were often oversmoothed by CNNs or removed by conventional filters.

These improvements reflect diffusion models' ability to learn complex noise distributions and high-dimensional manifolds characteristic of medical images.

5.2 Robustness in Low-Resource Scenarios

A key contribution of this work is demonstrating the feasibility of deploying advanced generative models in resource-constrained settings:

- **Model Compression** techniques (quantization, pruning, and distillation) reduced model size by over 75% with minimal loss in accuracy, enabling real-time inference on devices such as Jetson Nano and Raspberry Pi 4.
- **Reduced training data** experiments showed that diffusion models maintained high performance even when trained on a fraction (30%) of the full dataset, highlighting their data efficiency—a critical trait for settings with limited labeled examples.

These results suggest that diffusion models can operate reliably where compute and data are both scarce, making them well-suited for rural clinics and mobile diagnostic units.

5.3 Sampling Efficiency

One commonly cited drawback of diffusion models is their high inference cost due to iterative sampling. However, we mitigated this through:

- **DDIM-based accelerated sampling**, reducing the number of denoising steps from 100 to 25 with negligible quality degradation.
- **Conditional guidance** using measured inputs (e.g., k-space in MRI) to converge faster and more accurately toward valid reconstructions.

These optimizations strike a balance between performance and runtime, with inference times approaching 1 second per image on modest hardware.

5.4 Generalization and Modality Adaptability

Unlike many supervised models that require modality-specific retraining, diffusion models generalize well to multiple tasks with modest architectural changes. In our tests:

- The same core model generalized across MRI, CT, and ultrasound with only minor modifications in conditioning and data augmentation strategies.
- The model’s ability to implicitly model diverse image priors made it robust to various noise types—Gaussian, Poisson, and speckle.

This adaptability supports broader use cases without the need for modality-specific models, reducing maintenance complexity in clinical deployments.

5.5 Limitations and Future Work

Despite encouraging results, several limitations remain:

- **Training cost remains high**, requiring GPU acceleration for initial model training—even if inference is lightweight. Future work should explore training on sparse or federated datasets across multiple sites.

- **Sampling speed**, while improved, is still slower than feedforward models. Further architectural innovations (e.g., diffusion transformers, learned samplers) may help close this gap.
- **Clinical validation**: While image quality metrics are useful, downstream impact on diagnostic accuracy must be evaluated through clinical studies involving expert radiologists.

6. Conclusion

Diffusion models offer a promising paradigm for medical image reconstruction and denoising in low-resource settings. With model compression and guided sampling, they can deliver high-quality imaging even when acquisition is limited or degraded. Our work highlights a path forward for democratizing advanced medical imaging through AI.

In this work, we have presented a diffusion model framework tailored for medical image reconstruction and denoising in low-resource settings. Through extensive evaluations on MRI, CT, and ultrasound modalities, our approach has demonstrated state-of-the-art performance in recovering high-fidelity images from severely degraded inputs, surpassing traditional and deep learning-based baselines.

Key strengths of the proposed method include its ability to:

- Operate effectively with limited data and compute resources,
- Generalize across imaging modalities with minimal architectural changes,
- Retain clinically relevant structures while removing complex noise patterns,
- Be deployed efficiently on low-cost hardware via quantization and model compression.

By combining the generative power of diffusion models with practical optimization strategies, this research bridges the gap between high-performance medical imaging and real-world clinical constraints. Our findings open the door for deploying advanced AI-driven imaging tools in underserved environments, rural clinics, and mobile diagnostic setups.

Future work will explore clinical validation, 3D volume reconstruction, and integration with real-time acquisition systems to further enhance applicability and impact in global health contexts.

References

[1].Goodfellow I. et al. Generative adversarial nets. Advances in neural information processing systems 27 (2014).

[2].Karras T. et al. Analyzing and improving the image quality of stylegan, 8110–8119 (2020).

[3].Kingma D. P., Welling M. et al. An introduction to variational autoencoders. Foundations and Trends® in Machine Learning 12, 307–392 (2019).

[4].Ho J., Jain A. & Abbeel P. Denoising diffusion probabilistic models. Advances in neural information processing systems 33, 6840–6851 (2020).

[5].Nichol A. Q. & Dhariwal P. Improved denoising diffusion probabilistic models, 8162–8171 (PMLR, 2021).

[6].Rombach R., Blattmann A., Lorenz D., Esser P. & Ommer B. High-resolution image synthesis with latent diffusion models, 10684–10695 (2022).

[7].Ktena I. et al. Generative models improve fairness of medical classifiers under distribution shifts. Nature Medicine 1–8 (2024). [DOI] [PMC free article] [PubMed]

[8].El Emam K., Mosquera L., Jonker E. & Sood H. Evaluating the utility of synthetic covid-19 case data. JAMIA open 4, ooab012 (2021). [DOI] [PMC free article] [PubMed]

[9].Rajotte J.-F. et al. Synthetic data as an enabler for machine learning applications in medicine. Iscience 25 (2022). [DOI] [PMC free article] [PubMed]
[10].Levine A. B. et al. Synthesis of diagnostic quality cancer pathology images by generative adversarial networks. The Journal of pathology 252, 178–188 (2020). [DOI] [PubMed]
[11].Dhariwal P. & Nichol A. Diffusion models beat gans on image synthesis. Advances in neural information processing systems 34, 8780–8794 (2021).
[12].Ramesh A., Dhariwal P., Nichol A., Chu C. & Chen M. Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1, 3 (2022).