FDP: A Frequency-Decomposition Preprocessing Pipeline for Unsupervised Anomaly Detection in Brain MRI

Hao Li[1*], Zhenfeng Zhuang[1*], Jingyu Lin[1*], Yu Liu[1], Yifei Chen[2], Qiong Peng[1], Lequan Yu[3†], Liansheng Wang[1†]

[1]Xiamen University
[2]Tencent
[3]The University of Hong Kong

Abstract

Due to the diversity of brain anatomy and the scarcity of annotated data, supervised anomaly detection for brain MRI remains challenging, driving the development of unsupervised anomaly detection (UAD) approaches. Current UAD methods typically utilize artificially generated noise perturbations on healthy MRIs to train generative models for normal anatomy reconstruction, enabling anomaly detection via residual maps. However, such simulated anomalies lack the biophysical fidelity and morphological complexity characteristic of true clinical lesions. To advance UAD in brain MRI, we conduct the first systematic frequency-domain analysis of pathological signatures, revealing two key properties: (1) anomalies exhibit unique frequency patterns distinguishable from normal anatomy, and (2) low-frequency signals maintain consistent representations across healthy scans. These insights motivate our Frequency-Decomposition Preprocessing (FDP) framework, the first UAD method to leverage frequency-domain reconstruction for simultaneous pathology suppression and anatomical preservation. FDP can integrate seamlessly with existing anomaly simulation techniques, consistently enhancing detection performance across diverse architectures while maintaining diagnostic fidelity. Experimental results demonstrate that FDP consistently improves anomaly detection performance when integrated with existing methods. Notably, FDP achieves a 17.63% increase in DICE score with LDM while maintaining robust improvements across multiple baselines.

Code — https://github.com/ls1rius/MRI_FDP.

Introduction

Magnetic resonance imaging (MRI) play a crucial role in medical imagine system to provide detailed tissue information without requiring invasive procedures or exposure to radiation for aiding radiologists with their diagnostic and decision-making processes. In typical clinical MRI systems, a strong static magnetic field aligns hydrogen nuclei in the body. The nuclei are then excited by radiofrequency (RF) pulses, and as they return to equilibrium they emit detectable signals. Spatial localization is achieved using gradient fields that modulate the phase and frequency of the signal. The measurements are acquired as samples in k-space, a spatial-frequency domain in which each point encodes a spatial-frequency component of the image (Sarracanie et al. 2015; McRobbie et al. 2017; Rinck 2018). The spatial-domain MRI is then reconstructed by applying an inverse Fourier transform to the k-space data.

Unsupervised anomaly detection (UAD) methods overcome the limitations of supervised approaches by eliminating the need for annotated pathological data. Typically, these methods learn representations of healthy anatomy through generative models trained on normal MRI scans with artificial noise. During inference, the models process input scans (which may contain anomalies) and attempt their reconstruction under the learned healthy representation. The pixel-wise residuals between original and reconstructed scans serve as anomaly scores, enabling detection of pathologies such as tumors or lesions without requiring abnormal training samples. This paradigm addresses key challenges including annotation scarcity, privacy constraints, and class imbalance, while being applicable to real-world clinical scenarios where abnormalities are rare and exhibit high variability (Bercea et al. 2025; Bao et al. 2024; Behrendt et al. 2024; Kascenas, Pugeault, and O'Neil 2022; Wyatt et al. 2022). However, such synthetic noise is limited in reflecting the complexity of real pathological variations, which restricts generalization to clinical cases.

Given the inherent frequency-domain nature of MRI generation (via k-space acquisition), recent studies have increasingly explored frequency-based approaches for reconstruction and analysis (Yi et al. 2023; Liu et al. 2024; Zou et al. 2025). However, UAD in brain MRI remains predominantly spatial-domain focused, neglecting the diagnostic potential of frequency-component analysis. This represents a significant oversight, as k-space data naturally encodes structural information through distinct frequency components - features that conventional image-space methods cannot directly access. The analytical potential of frequency-domain representations remains particularly underutilized for characterizing pathological deviations in unsupervised settings.

In this paper, we conduct the first systematic frequency-domain analysis of brain MRI anomalies, focusing on: (1) characterizing the intrinsic relationship between pathologi-

[*]These authors contributed equally.
[†]Corresponding authors.

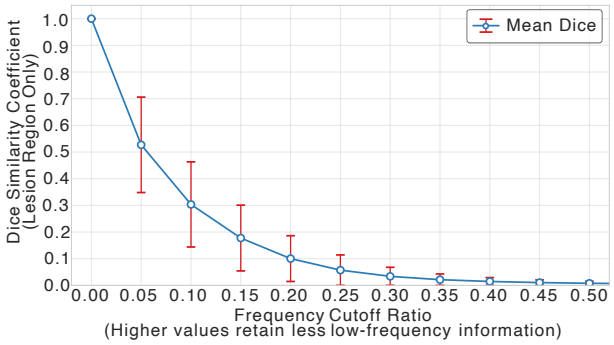

Figure 1: The effect of high-pass filter threshold on anomaly content in brain MRIs. The DICE coefficient between high-pass-filtered images and ground truth shows anomaly information primarily reside in low-frequency components.

cal features and their frequency signatures, (2) verifying the consistency of the low frequency representations, and (3) developing Frequency-Decomposition Preprocessing (FDP) - a modular component that mitigates pathological artifacts by low-frequency reconstruction while preserving diagnostically critical high-frequency details. This module could directly enhances existing anomaly simulation methods through its capability to separate and process pathological signatures in the frequency domain, significantly improving the quality of synthesized healthy MRIs without requiring modifications to downstream pipeline architectures.

Our main contributions are summarized as follows:

- We present the first comprehensive frequency-domain analysis of pathological MRI anomalies, revealing their distinctive frequency characteristics that are specific to lesion pathology.
- We propose FDP, a novel preprocessing framework for unsupervised anomaly detection that enhances MRI reconstruction by selectively filtering lesion-related frequencies while preserving anatomical integrity.
- We demonstrate that FDP consistently improves the performance of existing anomaly detection methods through comprehensive experimental validation.

Related Work

Generative Models Recent advances in generative models have significantly impacted medical image analysis, with several architectures demonstrating particular promise. Among these, generative adversarial networks (GAN)-based (Goodfellow et al. 2020; Isola et al. 2017) and variational autoencoder (VAE)-based (Pinaya et al. 2021; Raad et al. 2023; Wijanarko et al. 2024) methods have gained attention for their success in image generation and translation. The emergence of Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020; Nichol and Dhariwal 2021) marked a significant breakthrough, achieving superior generation quality through iterative denoising processes that provide better distribution coverage and training stability compared to previous approaches. Building on these foundations, the Latent Diffusion Model (LDM) (Rombach et al. 2022) improves computational efficiency by performing the diffusion process in a compressed latent space, and has been applied to medical and heterogeneous imaging (Kebaili et al. 2025; Dong et al. 2025; Lin et al. 2024).

Unsupervised Anomaly Detection in Brain MRI Baur et al. (Baur et al. 2019) initiated UAD in brain MRI lesion analysis by training spatially variant autoencoders on normal data only, using spatial bottlenecks to preserve spatial details, and detected lesions by analyzing pixel-wise reconstruction errors. Subsequent methods such as F-AnoGAN (Schlegl et al. 2019) and AAE (Chen and Konukoglu 2018) introduced generative adversarial networks for UAD in MRI, assuming that healthy regions remain unchanged during reconstruction and that latent representations of lesion and non-lesion images are similar. However, this assumption often fails, as lesion intensities can distort latent projections (Wijanarko et al. 2024). With the rise of denoising-based techniques, denoising autoencoders (DAEs)(Kascenas, Pugeault, and O'Neil 2022) and related approaches simulate anomalies by adding noise, typically in the form of discrete points. Others (Iqbal et al. 2023; Behrendt et al. 2024) apply large regular masks and reconstruct the original image from the corrupted input. AnoDDPM (Wyatt et al. 2022) introduced a partial diffusion scheme that adds noise at a specific timestep and reconstructs from the corrupted image. Recently, Tri-VAE (Wijanarko et al. 2024) improved UAD by introducing noise while decoupling metric learning from latent sampling. It aligns images to a lesion-free distribution, incorporates a semantic-guided retrieval module, and uses structural similarity as an additional training objective. Despite their controllability, artificial noise methods often lack natural variability and anatomical structure, limiting their realism and effectiveness. mDDPM (Iqbal et al. 2023) applies masking in both the image and frequency domains during training to encourage a diffusion model to learn a prior over healthy anatomy. At inference, no masking is introduced, and the model is directly conditioned on the unhealthy input. As a result, the reconstruction is not explicitly constrained to suppress pathology-related cues, and the generated image may retain lesion information, which can weaken residual-based anomaly localization.

Fourier Transform Fourier Transform converts a signal from the spatial domain into the frequency domain by breaking it down into its sine and cosine components at different frequencies, revealing the signal's frequency characteristics. In image processing, the Fourier Transform is used for frequency domain analysis (Wang et al. 2024; Vaish, Wang, and Strisciuglio 2024; Dong et al. 2024). High-frequency components capture fine details, edges, and textures, while low-frequency components represent broader and regional structural features. Fourier Transform converts signals into the frequency domain, enabling the application of filters for tasks such as noise reduction, edge detection, and compression, enhancing image processing capabilities (Cochran et al. 1967; Ye et al. 2024; Zhou et al. 2024). Given a 2D image $I \in R^{H \times W}$, where $I(x, y)$ denotes the image signal located at position (x, y), the corresponding 2D Discrete

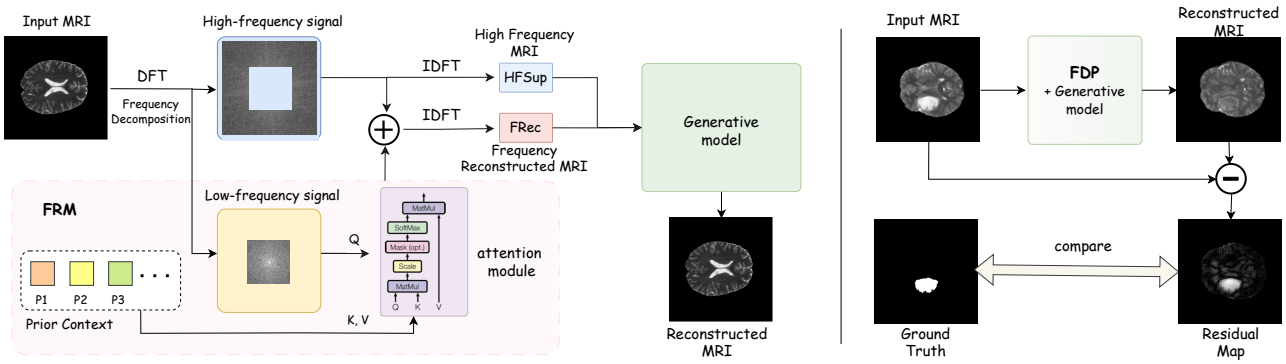

Figure 2: Training and inference pipeline of the proposed FDP method. **Left: Training Phase.** High-frequency signals are used for both frequency-domain reconstruction and as auxiliary input (HFSup) to enhance structural details. **Right: Inference Phase.** The input MRI with lesions is processed by FDP for frequency reconstruction, then fed into the generative model.

Fourier Transform (2D-DFT) can be represented as follows:

$$\begin{aligned} f(u,v) &= \mathrm{DFT}(I(x,y)) \\ &= \sum_{x=0}^{H-1}\sum_{y=0}^{W-1} I(x,y)e^{-j2\pi(\frac{ux}{H}+\frac{vy}{W})}, \end{aligned} \tag{1}$$

here u and v serve as indices representing the horizontal and vertical spatial frequencies in the Fourier spectrum, and $f(u,v)$ denotes the corresponding frequency signal. The 2D-IDFT is defined as follows:

$$\begin{aligned} I(x,y) &= \mathrm{IDFT}(f(u,v)) \\ &= \frac{1}{HW}\sum_{u=0}^{H-1}\sum_{v=0}^{W-1} f(u,v)e^{j2\pi(\frac{ux}{H}+\frac{vy}{W})}. \end{aligned} \tag{2}$$

Both DFT and IDFT can be computed using Fast Fourier Transform (FFT) algorithm (Brigham and Morrow 1967).

Methodology

We start with a systematic analysis of pathological signatures in the frequency domain of brain MRI, identifying their distinct frequency signal patterns. Building on the findings, we propose a novel frequency-domain signal processing framework that selectively suppresses anomalies while preserving anatomical features. The overall structure of Frequency-Decomposition Preprocessing (FDP) pipeline is shown in Figure 2. FDP primarily consists of the Frequency Reconstruction Module (FRM) for core frequency signal reconstruction, supplemented by the High-Frequency component (HFSup) for enhanced structural detail refinement. MRIs are first transformed into the frequency domain and split into high- and low-frequency components using a frequency filter. The FRM reconstructs the low-frequency part, which is then merged with the preserved high-frequency signals to form a complete frequency-domain representation. This is converted back into the spatial domain to guide healthy MRI reconstruction using a generative model. In parallel, the high-frequency signals are also transformed into an image and used as an auxiliary structural prior, which helps preserve anatomical structures and sharpen edges.

Characteristics of Anomaly in Frequency Domain

From the right part of Figure 2, we observe that the lesion region in pathological MRI appears as a smooth, homogeneous area in the spatial domain. According to Fourier principles, such gradual intensity variations primarily contribute to low-frequency signal components (Additional visual comparisons are provided in the Appendix). To verify this assumption, we transform the MRI into the frequency domain and apply a high-pass filter with threshold m, retaining a proportion of $(1-m)$ of the high-frequency components. Meanwhile, we employ the DICE coefficient (Dice 1945) to quantitatively assess the correlation between low-frequency components and lesion-related information. As shown in Figure 1, the DICE coefficient decreases rapidly as the high-pass filtering threshold m increases, dropping below 0.1 when m reaches 0.2. This indicates that lesion-related signals are primarily contained in the low-frequency components and can be effectively removed through high-pass filtering. **Unlike general anomaly detection tasks where anomalies are mainly concentrated in high-frequency components, MRI lesions typically manifest as continuous, low-frequency regional signals.** Leveraging this property, high-pass filtering not only suppresses irrelevant smooth or uniform background signals but also enhances the visibility of structural features such as edges and textures. Since high-frequency components contain minimal lesion-related information, emphasizing them allows the model to better capture anatomical structural information and improves its reconstruction fidelity.

Consistency of Low Frequency Signals

However, simply removing low-frequency signals may lead to information loss. To address this, we investigate whether it was possible to recover the lost low-frequency signals without reintroducing lesion-related information as much as possible. Our systematic frequency-domain analysis reveals fundamental differences between normal and pathological MRIs in their real and imaginary components. As demonstrated in Figure 3, through log-amplitude analysis in the frequency domain, two principal findings emerge: (1) normal MRIs exhibit highly consistent real-component patterns

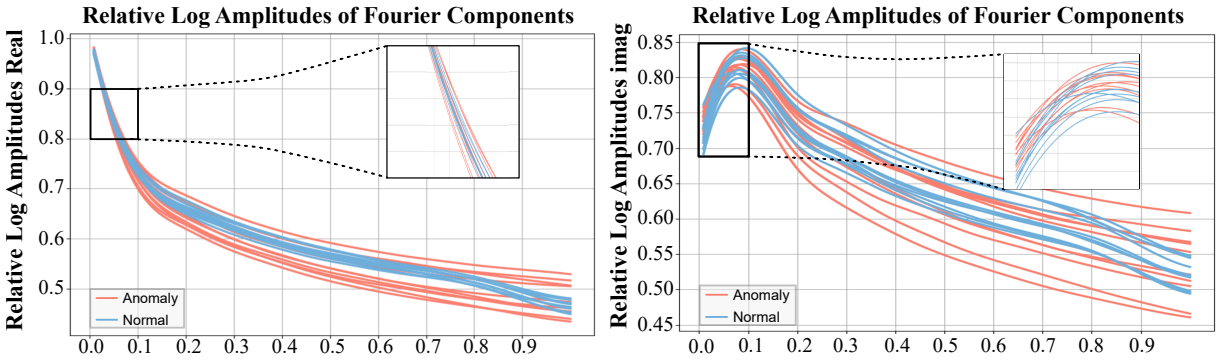

Figure 3: The horizontal and vertical axes represent the high-pass filtering threshold m and the normalized real and imaginary parts of the amplitude, respectively.

in low-frequency regions ($m \leq 0.1$), whereas pathological cases show substantially greater dispersion in the same bands; (2) both groups display comparable signal variability in high-frequency ranges ($m > 0.1$). These observations are consistent with basic frequency-domain principles: low-frequency components ($m \leq 0.1$) represent global anatomical structures, while high-frequency components ($m > 0.1$) capture local details. In normal MRIs, stable low-frequency patterns reflect structural homogeneity, while pathological alterations disrupt this coherence, manifesting as increased dispersion in low-frequency bands—consistent with lesion-induced anatomical discontinuities (More quantitative analysis is provided in the Appendix).

Frequency Decomposition

The signal processing is performed in the frequency domain. Although direct access to native k-space data is ideal for this purpose, such data are often difficult to obtain in practice. This is due to several factors, including patient privacy considerations, the large storage requirements of raw frequency data, and substantial variability in acquisition protocols across MRI vendors. Moreover, most publicly available datasets provide only reconstructed images rather than raw k-space measurements. Therefore, we apply a Fourier transform to convert MRI images back into the frequency domain for subsequent processing.

For a MRI $I \in R^{H \times W}$, we apply 2D Discrete Fourier Transform (2D-DFT) to convert the image from the spatial domain to the frequency domain. We then use an ideal high-pass filter $\mathrm{H_{hp}}(u, v)$ to retain the high-frequency components and remove the low-frequency components. The filter $\mathrm{H_{hp}}(u, v)$ is defined as follows:

$$\mathrm{H_{hp}}(u, v) = \begin{cases} 0 & \text{if } \mathrm{D}(u, v) \leq \mathfrak{D}_0 \\ 1 & \text{if } \mathrm{D}(u, v) > \mathfrak{D}_0 \end{cases}, \tag{3}$$

where $\mathrm{D}(u, v)$ is the distance from the frequency origin and $\mathfrak{D}_0$ refers to the distance from the current point to the center of the frequency domain, given by:

$$\mathrm{D}(u, v) = \sqrt{(u - H/2)^2 + (v - W/2)^2}, \tag{4}$$

$$\mathfrak{D}_0 = \min(m * H, m * W), m \in [0, 1], \tag{5}$$

where m represents the threshold of the high-pass filter.

Finally, the 2D Inverse Discrete Fourier Transform (2D-IDFT) converts the filtered signal back from the frequency domain to the spatial domain, i.e.,

$$f(u, v) = \mathrm{DFT}(I(x, y)), \tag{6}$$

$$f_{\mathrm{h}}(u, v) = f(u, v) \odot \mathrm{H_{hp}}(u, v), \tag{7}$$

$$f_{\mathrm{l}}(u, v) = \mathrm{CROP}(f(u, v), f_{\mathrm{h}}(u, v)), \tag{8}$$

$$I_{\mathrm{h}} = \mathrm{IDFT}(f_{\mathrm{h}}(u, v)), \tag{9}$$

here $\mathrm{CROP}(a, b)$ means to delete off the high-frequency signals b from the complete frequency signals a. In this way, we decompose an MRI I into low-frequency signals f_{l} and high-frequency signals f_{h} using a high-pass filter. In addition, the high-frequency image I_h can later be used to enhance image detail quality.

Frequency Reconstruction Module

Through comprehensive quantitative analyses (e.g., PCA, t-SNE, and maximum likelihood estimation) presented in the Appendix, we demonstrate that these low-frequency signals exhibit low variance and lie approximately on a low-dimensional manifold. This motivates us to model the distribution of low-frequency signals using a prior context bank $\boldsymbol{P}$, estimated from healthy MRIs, such that the low-frequency component of any healthy MRI can be approximated via latent sampling and mapping.

We therefore propose the Frequency Reconstruction Module (FRM), which serves as a decoder or a linear mapping based on learned dictionaries. This is analogous to the reparameterization trick used in variational inference, enabling us to sample smooth, consistent low-frequency signals in a differentiable manner. We adopt an attention-based retrieval strategy from a learnable set of prior context $\boldsymbol{P} = [p_1, p_2, \ldots p_k]$, where k denotes the number of learnable prior context elements and each $p_i \in R^{(m*H)\times(m*W)}$. These contexts are initialized using k-means++ (Arthur and Vassilvitskii 2006) clustering over the training set to promote faster and more stable convergence. Given a low-frequency query f_{l}, we reconstruct the MRI via:

$$\hat{f}_{\mathrm{l}} = \mathrm{ATTN}(f_{\mathrm{l}}, P, P), \tag{10}$$

$$\hat{f} = \mathrm{MERGE}(\hat{f}_{\mathrm{l}}, f_{\mathrm{h}}), \tag{11}$$

$$\hat{I} = \mathrm{IDFT}(\hat{f}), \tag{12}$$

where ATTN represents an attention module, and $\mathrm{MERGE}(a, b)$ merges the low-frequency component a and the high-frequency component b into a complete frequency representation. In this way, the reconstructed low-frequency signals $\hat{f}\mathrm{l}$ are integrated with the original high-frequency signals $f\mathrm{h}$ to obtain the reconstructed frequency signals $\hat{f}$, which are subsequently transformed back to the spatial domain to produce the frequency-reconstructed MRI $\hat{I}$. We supervise the reconstruction of the low-frequency signals using the L1 loss, formulated as $\mathrm{L1}(\hat{f}\mathrm{l}, f\mathrm{l})$. This design enables the model to generate plausible low-frequency content conditioned on the preserved high-frequency signal and the prior context, thereby minimizing the chance that lesion-related information is reintroduced during reconstruction.

Subsequently, $\hat{I}$ and I_{h} can be used for MRI reconstruction using generative models such as VAEs and LDMs. Additionally, the obtained I_{h} can serve as auxiliary structural information, which we term High-Frequency Supplement (HFSup), to enhance anatomical detail preservation.

Experiments

DataSets

We adopted T2-weighted MRI as the primary modality because it is widely available and provides strong tissue contrast. In comparison, T1-weighted images primarily capture anatomical detail but often show weaker conspicuity for many lesions. FLAIR acquisition protocols are less standardized across sites, and its advantage can be region dependent.

For training, we used the publicly available IXI dataset, which includes 560 T1- and T2-weighted brain MRI scans from three clinical centers. The dataset comprises only healthy subjects, providing a clean distribution for modeling normal anatomy. For evaluation, we conduct 10 evaluation runs, each using 32 cases randomly sampled from BraTS 2020 dataset (Bakas et al. 2018; Menze et al. 2014), which includes 369 annotated brain MRI scans across four modalities (T1, T1-CE, T2, FLAIR) and corresponding tumor segmentations. Each scan contains roughly 155 slices, with T2 images acquired at 0.9375×0.9375 mm^2 in-plane resolution, 0.125 mm slice thickness, and 240×240 image size.

In addition, to further validate the generalization capability of our model, we conducted experiments on the T2-weighted modality across several additional MRI brain datasets, including the Multimodal Brain Tumor Segmentation Challenge 2021 (BraTS21) (Baid et al. 2021), the multiple sclerosis dataset (MSLUB) (Lesjak et al. 2018), and the Multiple Sclerosis Lesion (MSSEG-2) (Commowick et al. 2021). These datasets use similar storage formats as BraTS20 but cover more diverse brain pathologies, offering broader scenarios to evaluate model generalization.

Implementation Details

We applied skull stripping with HD-BET (Isensee et al. 2019) to filter out the regions belonging to the foreground area so that the masking block can only be applied to the foreground pixel patches. During the experiments, all slices are resized to a uniform resolution of 256×256 after normalization. Unless otherwise specified, the LDM was chosen as the default generative model, the number of prior context was set to 128 by default, and the high-pass filtering threshold m used for FRM and HFSup are both set to 0.10 as default. The model was trained on 4 NVIDIA V100 GPUs (32GB) using the Adam optimizer, with a learning rate of 2e-5 and a batch size of 32 for 800 epochs. The details of evaluation metrics (DICE, AUPRC, AUROC) and post-processing can be found in Appendix. All of our experiments adopted the unified post-processing method.

Results and Analysis

We evaluated our method's compatibility with other open-source approaches. As presented in Table 1, integrating FDP with the base LDM results in a 17.63% improvement in the DICE score and a 20.92% improvement in the AUPRC score. Furthermore, when FDP is integrated with other methods, the results improve by at least 6.16% in DICE and 11.37% in AUPRC. In terms of AUROC, our improvements are limited, with some methods exhibiting slight performance drops after integration. We attribute this to the task's complexity and the trade-offs in anomaly removal. As shown in Figure 4, our predictions contain fewer errors, while the predicted lesion regions tend to be slightly smaller than those of other methods and the ground truth, indicating that our method achieves higher precision at the cost of slightly lower recall. To intuitively demonstrate FDP's effectiveness in removing MRI anomalies, we include visualizations of results at each pipeline stage in the Appendix.

Furthermore, our method exhibits strong generalization across multiple MRI datasets, achieving substantial improvements over the baseline LDM. As shown in Table 2, it yields significant performance gains of 15.53% and 14.48% DICE points on standard benchmarks (BraTS21 and MSSEG-2, respectively). Notably, even for the more challenging MSLUB dataset, it maintains a robust improvement of 4.71% DICE points. These results collectively demonstrate the consistent effectiveness of our method across diverse MRI data domains.

Ablation Studies

In this section, we analyze the impact of m_{FRM} (m used for FRM), m_{HFSup} (m used for HFSup), and the number of prior context on model performance. We further investigate how different hyperparameter settings for each component affect the overall performance. By default, we employ LDM without noise-adding as our generative model, with m_{FRM} set to 0.10, m_{HFSup} also set to 0.10, and the number of prior context set to 128.

As shown in Table 2 and Table 3, high-frequency information alone does not substantially improve performance. Instead, it serves as a supplementary feature that enhances

Model	DICE	AUPRC	AUROC
Base Models			
VAE (Chen et al. 2020)	34.90 ± 2.1	29.95 ± 3.3	94.46 ± 0.5
LDM (Rombach et al. 2022)	35.02 ± 1.8	30.75 ± 2.7	91.62 ± 1.1
AnoDDPM (Wyatt et al. 2022)	36.19 ± 1.5	32.01 ± 2.4	91.37 ± 1.3
F-AnoGAN (Schlegl et al. 2019)	37.68 ± 1.2	35.05 ± 1.9	91.88 ± 0.9
DAE (simplex) (Kascenas, Pugeault, and O'Neil 2022)	37.38 ± 1.4	32.93 ± 2.8	94.65 ± 0.6
DAE (coarse)	56.87 ± 3.2	43.23 ± 2.5	95.71 ± 0.5
pDDPM (Behrendt et al. 2024)	46.15 ± 2.4	45.67 ± 2.9	92.01 ± 0.9
FDP-enhanced Models			
FDP + VAE	46.32 ± 1.9 (+11.42)	41.32 ± 2.6 (+11.37)	92.16 ± 0.8 (-2.30)
FDP + LDM	52.66 ± 2.4 (+17.63)	51.67 ± 3.1 (+20.92)	93.12 ± 0.9 (+1.50)
FDP + AnoDDPM	48.24 ± 1.7 (+12.05)	47.56 ± 2.9 (+15.55)	92.97 ± 1.0 (+1.60)
FDP + F-AnoGAN	52.96 ± 2.1 (+15.28)	50.70 ± 2.7 (+15.65)	93.84 ± 0.7 (+1.96)
FDP + DAE (simplex)	54.77 ± 2.3 (+17.39)	51.88 ± 3.0 (+18.95)	94.69 ± 0.6 (+0.04)
FDP + DAE (coarse)	63.03 ± 3.5 (+6.16)	61.41 ± 3.3 (+18.18)	93.95 ± 1.2 (-1.76)
FDP + pDDPM	54.09 ± 1.9 (+8.04)	51.03 ± 3.0 (+5.36)	93.72 ± 1.0 (+1.62)

Table 1: Comparing anomaly detection performance on the BraTS20 dataset. Values in parentheses indicate performance changes after integrating FDP. Red indicates improved performance, while blue indicates reduced performance.

Dataset	Model	DICE	AUPRC	AUROC
BraTS21	LDM	29.53	27.52	92.26
	LDM+FDP	45.06	37.17	93.13
MSLUB	LDM	8.35	8.62	84.98
	LDM+FDP	13.06	14.28	87.14
MSSEG-2	LDM	20.15	25.67	85.22
	LDM+FDP	34.63	37.29	88.32

Table 2: Comparison of models across different datasets.

FRM	HFSup	DICE	AUPRC	AUROC
		35.02	30.75	91.62
	✓	42.18	40.55	92.25
✓		50.00	45.86	92.97
✓	✓	52.66	51.67	93.12

Table 3: Ablation results on the impact of FRM and HFSup.

the baseline model. Relying exclusively on high-frequency information may lead to the loss of essential details in the input image, providing insufficient information for accurate MRI reconstruction. The FRM plays a pivotal role in enhancing overall model performance by effectively suppressing lesion-related information while preserving structural details in high-frequency components, thereby improving image reconstruction quality. The synergistic integration of these modules leads to significant performance gains.

High Frequency Filtering Threshold To determine the appropriate m_{FRM} and m_{HFSup} values for FRM and HFSup, we evaluated different mask ratios on each to assess their respective optimal settings. The results are presented in Table 4a and Table 4b, respectively. In Table 4a, m_{FRM} is

m_{FRM}	DICE	AUPRC	AUROC
0.01	39.45	37.13	90.54
0.05	50.00	44.86	92.97
0.10	52.66	51.67	93.12
0.15	45.24	43.56	91.78
0.20	43.89	41.87	93.08
0.25	43.75	41.68	92.73
0.30	41.82	39.02	92.45

(a) Effects of m_{FRM} (with $m_{\text{HFSup}} = 0.10$).

m_{HFSup}	DICE	AUPRC	AUROC
0.01	38.27	38.48	90.01
0.05	50.13	49.04	92.88
0.10	52.66	51.67	93.12
0.15	51.31	42.78	93.82
0.20	49.44	49.47	92.67
0.25	49.93	48.64	93.85
0.30	50.76	49.49	93.75

(b) Effects of m_{HFSup} (with $m_{\text{FRM}} = 0.10$).

Table 4: Comparison of hyperparameter sensitivity for FRM.

the variable parameter, while m_{HFSup} is fixed at 0.10. Conversely, in Table 4b, m_{HFSup} is the variable parameter, and m_{FRM} is fixed at 0.10.

As shown in Table 4a, when m_{FRM} is set to 0.10, the model achieves optimal performance by effectively removing lesions while retaining high-frequency MRI information and supplementing low-frequency signals through reconstruction. When m_{FRM} is set to 0.01, the low-frequency signals primarily capture general image characteristics (e.g., illumination intensity, etc.), which are common across similar images. In this case, the method fails to remove lesions effectively during inference, rendering it ineffective. When m_{FRM} is too large, the frequency representation mainly encodes subject-specific features, causing the signals to become overly dispersed and difficult to simulate using the prior context. This leads to substantial deviations between the generated and original images, resulting in suboptimal performance.

As shown in the Table 4b, when m_{HFSup} is set to 0.01, lesion information is directly introduced due to the overly low threshold. The retention of excessive information shifts the model's focus toward high-frequency signals contain-

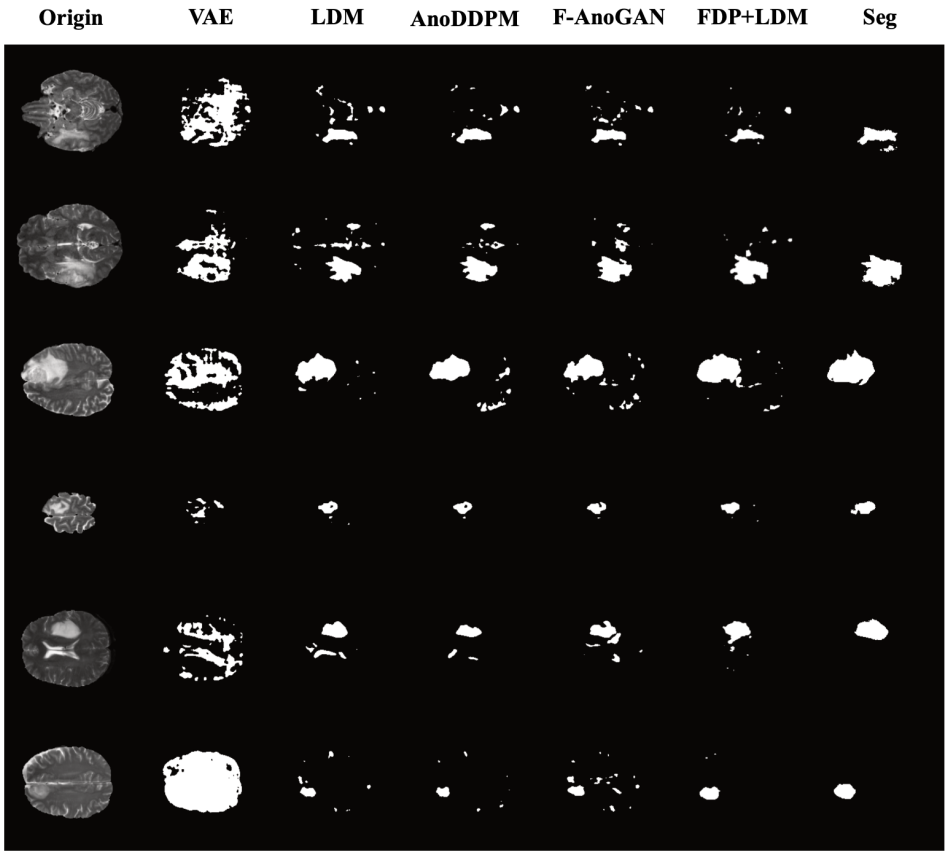

Figure 4: Visual comparison of results with other methods. Origin means the input MRIs, and Seg denotes the ground truth.

ing lesion details, leading to degraded performance. Conversely, when m_{HFSup} exceeds 0.15, the high-frequency structural information becomes insufficient, markedly reducing its contribution to model performance.

Prior Context Number We conducted experiments with varying amounts of prior context, evaluating five configurations (16, 32, 64, 128, and 256) to identify the optimal configuration for reconstructing low-frequency signals of healthy MRI images. As shown in Table 5, increasing the amount of prior context improves performance, with gains plateauing around 128. The initial addition of 16 contexts yields about a 2% improvement, but further increases bring limited benefit due to empty boundary slices and high inter-slice similarity. Given that typical MRI scans contain approximately 155 slices, 128 contexts effectively capture the volume's low-frequency structure without incurring unnecessary computational overhead.

Prior Context Number	DICE	AUPRC	AUROC
16	45.94	45.23	92.56
32	47.55	47.29	94.13
64	50.31	49.70	92.78
128	52.66	51.67	93.12
256	52.74	50.56	93.01

Table 5: Effects of different numbers of prior context.

Conclusion

In this paper, we present the first systematic frequency-domain analysis of pathological signatures in brain MRI, revealing two key properties: (1) the frequency separability between lesions and normal anatomy, and (2) the consistent low-frequency representations in healthy brain MRIs. Motivated by these insights, we propose Frequency-Decomposition Preprocessing (FDP), the preprocessing framework for UAD that leverages frequency-domain reconstruction to suppress pathologies and preserve anatomical structures. FDP integrates seamlessly with existing methods, without requiring any architectural changes to the original model. Through extensive experiments, we show that frequency-domain processing effectively suppresses pathological information, providing a reliable solution for unsupervised anomaly detection in brain MRI. We hope that our method will inspire new research directions for UAD in MRI and contribute to advancing the field.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. 62371409) and Fujian Provincial Natural Science Foundation of China (Grant No. 2023J01005).

References

Abdi, H.; and Williams, L. J. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4): 433–459.

Arthur, D.; and Vassilvitskii, S. 2006. k-means++: The advantages of careful seeding. Technical report, Stanford.

Baid, U.; Ghodasara, S.; Mohan, S.; Bilello, M.; Calabrese, E.; Colak, E.; Farahani, K.; Kalpathy-Cramer, J.; Kitamura, F. C.; Pati, S.; et al. 2021. The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314*.

Bakas, S.; Reyes, M.; Jakab, A.; Bauer, S.; Rempfler, M.; Crimi, A.; Shinohara, R. T.; Berger, C.; Ha, S. M.; Rozycki, M.; et al. 2018. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv preprint arXiv:1811.02629*.

Bao, J.; Sun, H.; Deng, H.; He, Y.; Zhang, Z.; and Li, X. 2024. Bmad: Benchmarks for medical anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4042–4053.

Baur, C.; Wiestler, B.; Albarqouni, S.; and Navab, N. 2019. Deep autoencoding models for unsupervised anomaly segmentation in brain MR images. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part I 4*, 161–169. Springer.

Behrendt, F.; Bhattacharya, D.; Krüger, J.; Opfer, R.; and Schlaefer, A. 2024. Patched diffusion models for unsupervised anomaly detection in brain mri. In *Medical Imaging with Deep Learning*, 1019–1032. PMLR.

Bercea, C. I.; Wiestler, B.; Rueckert, D.; and Schnabel, J. A. 2025. Evaluating normative representation learning in generative AI for robust anomaly detection in brain imaging. *Nature Communications*, 16(1): 1624.

Brigham, E. O.; and Morrow, R. 1967. The fast Fourier transform. *IEEE spectrum*, 4(12): 63–70.

Chen, X.; and Konukoglu, E. 2018. Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. *arXiv preprint arXiv:1806.04972*.

Chen, X.; You, S.; Tezcan, K. C.; and Konukoglu, E. 2020. Unsupervised lesion detection via image restoration with a normative prior. *Medical image analysis*, 64: 101713.

Cochran, W. T.; Cooley, J. W.; Favin, D. L.; Helms, H. D.; Kaenel, R. A.; Lang, W. W.; Maling, G. C.; Nelson, D. E.; Rader, C. M.; and Welch, P. D. 1967. What is the fast Fourier transform? *Proceedings of the IEEE*, 55(10): 1664–1674.

Commowick, O.; Cervenansky, F.; Cotton, F.; and Dojat, M. 2021. MSSEG-2 challenge proceedings: Multiple sclerosis new lesions segmentation challenge using a data management and processing infrastructure. In *MICCAI 2021-24th international conference on medical image computing and computer assisted intervention*, 126.

Dice, L. R. 1945. Measures of the amount of ecologic association between species. *Ecology*, 26(3): 297–302.

Dong, K.; Cheng, Y.; He, K.; and Suo, J. 2025. A generative model uses healthy and diseased image pairs for pixel-level chest X-ray pathology localization. *Nature Biomedical Engineering*, 1–13.

Dong, W.; Zhou, H.; Tian, Y.; Sun, J.; Liu, X.; Zhai, G.; and Chen, J. 2024. ShadowRefiner: Towards mask-free shadow removal via fast fourier transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6208–6217.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.

Iqbal, H.; Khalid, U.; Chen, C.; and Hua, J. 2023. Unsupervised anomaly detection in medical images using masked diffusion model. In *International Workshop on Machine Learning in Medical Imaging*, 372–381. Springer.

Isensee, F.; Schell, M.; Pflueger, I.; Brugnara, G.; Bonekamp, D.; Neuberger, U.; Wick, A.; Schlemmer, H.-P.; Heiland, S.; Wick, W.; et al. 2019. Automated brain extraction of multisequence MRI using artificial neural networks. *Human brain mapping*, 40(17): 4952–4964.

Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.

Kascenas, A.; Pugeault, N.; and O'Neil, A. Q. 2022. Denoising autoencoders for unsupervised anomaly detection in brain MRI. In *International Conference on Medical Imaging with Deep Learning*, 653–664. PMLR.

Kebaili, A.; Lapuyade-Lahorgue, J.; Vera, P.; and Ruan, S. 2025. Multi-modal MRI synthesis with conditional latent diffusion models for data augmentation in tumor segmentation. *Computerized Medical Imaging and Graphics*, 123: 102532.

Lesjak, Ž.; Galimzianova, A.; Koren, A.; Lukin, M.; Pernuš, F.; Likar, B.; and Špiclin, Ž. 2018. A novel public MR image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus. *Neuroinformatics*, 16: 51–63.

Levina, E.; and Bickel, P. 2004. Maximum likelihood estimation of intrinsic dimension. *Advances in neural information processing systems*, 17.

Lin, J.; Wu, Y.; Wang, Z.; Liu, X.; and Guo, Y. 2024. PairID: A dual modal framework for identity preserving image generation. *IEEE Signal Processing Letters*.

Liu, C.; Li, Z.; Chen, Z.; Zhao, B.; Zheng, Z.; and Song, X. 2024. Highly-accelerated CEST MRI using frequency-offset-dependent k-space sampling and deep-learning reconstruction. *Magnetic Resonance in Medicine*, 92(2): 688–701.

Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.

McRobbie, D. W.; Moore, E. A.; Graves, M. J.; and Prince, M. R. 2017. *MRI from Picture to Proton*. Cambridge university press.

Menze, B. H.; Jakab, A.; Bauer, S.; Kalpathy-Cramer, J.; Farahani, K.; Kirby, J.; Burren, Y.; Porz, N.; Slotboom, J.; Wiest, R.; et al. 2014. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging*, 34(10): 1993–2024.

Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.

Pinaya, W. H. L.; Tudosiu, P.-D.; Gray, R.; Rees, G.; Nachev, P.; Ourselin, S.; and Cardoso, M. J. 2021. Unsupervised brain anomaly detection and segmentation with transformers. *arXiv preprint arXiv:2102.11650*.

Raad, J. D.; Chinnam, R. B.; Arslanturk, S.; Tan, S.; Jeong, J.-W.; and Mody, S. 2023. Unsupervised abnormality detection in neonatal MRI brain scans using deep learning. *Scientific Reports*, 13(1): 11489.

Rinck, P. A. 2018. *Magnetic resonance in medicine: a critical introduction*. BoD–Books on Demand.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Sarracanie, M.; LaPierre, C. D.; Salameh, N.; Waddington, D. E.; Witzel, T.; and Rosen, M. S. 2015. Low-cost high-performance MRI. *Scientific reports*, 5(1): 15177.

Schlegl, T.; Seeböck, P.; Waldstein, S. M.; Langs, G.; and Schmidt-Erfurth, U. 2019. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54: 30–44.

Vaish, P.; Wang, S.; and Strisciuglio, N. 2024. Fourier-basis functions to bridge augmentation gap: Rethinking frequency augmentation in image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17763–17772.

Wang, W.; Wang, J.; Chen, C.; Jiao, J.; Cai, Y.; Song, S.; and Li, J. 2024. Fremim: Fourier transform meets masked image modeling for medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 7860–7870.

Wijanarko, H.; Calista, E.; Chen, L.-F.; and Chen, Y.-S. 2024. Tri-VAE: Triplet Variational Autoencoder for Unsupervised Anomaly Detection in Brain Tumor MRI. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3930–3939.

Wyatt, J.; Leach, A.; Schmon, S. M.; and Willcocks, C. G. 2022. Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 650–656.

Ye, Y.; Xu, K.; Huang, Y.; Yi, R.; and Cai, Z. 2024. Diffusionedge: Diffusion probabilistic model for crisp edge detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 6675–6683.

Yi, Q.; Fang, F.; Zhang, G.; and Zeng, T. 2023. Frequency learning via multi-scale fourier transformer for mri reconstruction. *IEEE Journal of Biomedical and Health Informatics*, 27(11): 5506–5517.

Zhou, C.; Huang, Y.; Pu, M.; Guan, Q.; Deng, R.; and Ling, H. 2024. Muge: Multiple granularity edge detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25952–25962.

Zou, J.; Liu, L.; Chen, Q.; Wang, S.; Hu, Z.; Xing, X.; and Qin, J. 2025. MMR-Mamba: Multi-modal MRI reconstruction with Mamba and spatial-frequency information fusion. *Medical Image Analysis*, 102: 103549.