

# Rethinking Bias in Generative Data Augmentation for Medical AI: A Frequency Recalibration Method

Chi Liu<sup>1\*</sup>, Jincheng Liu<sup>1</sup>, Congcong Zhu<sup>1</sup>, Minghao Wang<sup>1</sup>, Sheng Shen<sup>2</sup>, Jia Gu<sup>1</sup>, Tianqing Zhu<sup>1</sup>, Wanlei Zhou<sup>1</sup>

<sup>1</sup>Faculty of Data Science, City University of Macau, Macao SAR, China

<sup>2</sup>Design and Creative Technology vertical, Torrens University Australia, Ultimo, 2007, NSW, Australia  
{chiliu,D21090103791,cczhu,mhwang,jiagu,tqzhu,wlzhou}@cityu.edu.mo, sheng.shen@torrens.edu.au

## Abstract

Developing Medical AI relies on large datasets and easily suffers from data scarcity. Generative data augmentation (GDA) using AI generative models offers a solution to synthesize realistic medical images. However, the bias in GDA is often underestimated in medical domains, with concerns about the risk of introducing detrimental features generated by AI and harming downstream tasks. This paper identifies the frequency misalignment between real and synthesized images as one of the key factors underlying unreliable GDA and proposes the Frequency Recalibration (FreRec) method to reduce the frequency distributional discrepancy and thus improve GDA. FreRec involves (1) Statistical High-frequency Replacement (SHR) to roughly align high-frequency components and (2) Reconstructive High-frequency Mapping (RHM) to enhance image quality and reconstruct high-frequency details. Extensive experiments were conducted in various medical datasets, including brain MRIs, chest X-rays, and fundus images. The results show that FreRec significantly improves downstream medical image classification performance compared to uncalibrated AI-synthesized samples. FreRec is a standalone post-processing step that is compatible with any generative model and can integrate seamlessly with common medical GDA pipelines.

## Introduction

Developing medical AI has become an essential practice for computer-aided disease diagnosis (Litjens et al. 2017; Li et al. 2018). This commonly involves a data-driven task that relies heavily on a large volume of medical images for model training. However, it often suffers from data scarcity due to privacy issues, high costs of collecting real-world medical images, and highly imbalanced class distributions (Rajpurkar et al. 2022; He et al. 2020; Wang et al. 2024b).

A promising solution to address the shortage of clinically available medical images is using generative AI models such as Generative Adversarial Networks (GANs) and Diffusion Models (DMs) to synthesize realistic images that augment existing training datasets (Dayarathna et al. 2024; Wang et al. 2024c). Known as generative data augmentation (GDA), this technique increasingly produces large-scale synthetic medical images that closely resemble real data,

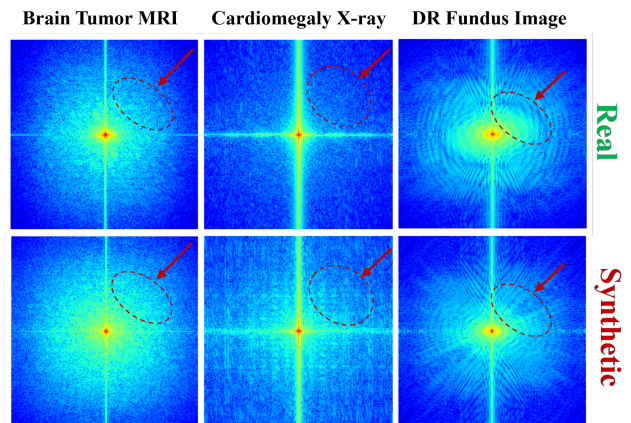


Figure 1: The average spectra of real and AI-synthetic images with different diseases and image types. The higher-frequency differences are discernible (e.g. as red circles indicate). Synthetic images are generated by different models: MRIs from FastGAN, fundus images from VC-Diffusion and X-rays from StyleGAN3.

thereby improving disease classification performance (Liu et al. 2020; Gao et al. 2023; Chen et al. 2024; Shang et al. 2024). Moreover, GDA can generate diverse pathological samples from healthy data, enhancing cross-domain generalization, supporting few-shot learning, and mitigating classifier bias (Ktena et al. 2024).

Despite being widely applied to developing medical AI, GDA faces growing concerns regarding its reliability. In general AI fields, AI-synthesized samples have proven to cause bias in certain tasks. For instance, language models can experience performance collapse after repeated training on generated content (Shumailov et al. 2024). Likewise, computer vision models trained with GDA often fail to achieve consistent improvement through trial and error (Singh et al. 2024). However, in the medical imaging domain, this phenomenon is generally overlooked. Alongside many medical AI studies having reported a positive effect of GDA, it remains unclear whether synthesized medical samples consistently benefit downstream tasks or whether they may introduce detrimental features generated by AI.

\*Corresponding Author (email: chiliu@cityu.edu.mo)

Intrigued by this gap, in this paper, we provide a close look at the bias of GDA when applying in medical image classification tasks. Our empirical results show that GDA can sometimes be detrimental for downstream training, mirroring observations in general AI fields. Unlike a recent study of natural image tasks that addressed this bias as a domain shift without exploring its underlying cause (Wang et al. 2024d), we investigate it from a frequency-domain perspective. Our analysis is inspired by recent research revealing high-frequency misalignment between real and AI-generated images (Durall, Keuper, and Keuper 2020; Dzanic, Shah, and Witherden 2020; Liu et al. 2021; Corvi et al. 2023; Frank et al. 2020). Since medical images such as MRIs and X-rays are especially sensitive to high-frequency variations due to their imaging processes and reliance on subtle pathological details, they may be particularly vulnerable to such frequency gaps.

Hence, we are motivated to posit this frequency gap (as shown in Fig. 1) as one of the key factors underlying unreliable GDA. Following this, we propose Frequency Recalibration (FreRec), a two-step method enabling coarse-to-fine alignment of frequency distributions between synthesized and real images. The first step, Statistic High-frequency Replacement (SHR), roughly aligns distributions by replacing the high-frequency components of AI samples with statistically sampled counterparts from real images. The second step, Reconstructive High-frequency Mapping (RHM), refines the perturbed samples from Step 1, reconstructing their high-frequency details by mapping onto the frequency distribution of natural images while enhancing the image quality for downstream use. Experiments on diverse medical image classification tasks demonstrate that FreRec-calibrated GDA samples significantly and consistently improve downstream performance compared to uncalibrated AI-synthesized samples. Notably, unlike previous methods that require retraining generative models with frequency-domain regularization to reduce their frequency bias (Durall, Keuper, and Keuper 2020; Jiang et al. 2021), FreRec is a standalone post-processing step compatible with any generative model, including GANs and diffusion models, making it a practical, plug-and-play, and cost-effective solution for medical AI pipelines.

## Related Work

### Generative Data Augmentation for Medical AI

Using Generative AI models to synthesize medical data for dataset augmentation has become a common practice to address medical data scarcity in developing medical AI (Wang et al. 2024c). Successful applications can be identified in various medical domains, including X-ray scans (Liu et al. 2020; Gao et al. 2023), computed tomography (CT) (Chen et al. 2024), fundus images (Shang et al. 2024), MRI (Dayarathna et al. 2024), and even foundation models for multi-modality generation (Wang et al. 2024a). The usage of GDA can also be extended to serve specific purposes such as privacy-preserving training (Guillaudeux et al. 2023) or fairness enhancement (Ktena et al. 2024). However, although the reliability of AI-synthesized samples has recently been

scrutinized in general deep learning fields (Shumailov et al. 2024; Singh et al. 2024), in medical domains, this risk remains underestimated and underexplored. A recent study addressed the issue from a domain-shift perspective but evaluated only natural images and did not investigate its root cause (Wang et al. 2024d).

### Frequency Gap in AI-Synthesized Images

Previous studies have identified a frequency gap between AI-generated and real images. Existing work (Xu et al. 2019; Wang et al. 2020; Rahaman et al. 2019) has demonstrated that deep neural networks exhibit preference in learning information across different frequency bands, leading to frequency artifacts when applied to generating images. For example, Durall et al. (Durall, Keuper, and Keuper 2020) observed significant differences in the spectral distributions of real and GAN-generated images. Other research (Dzanic, Shah, and Witherden 2020; Liu et al. 2021; Corvi et al. 2023; Frank et al. 2020) has shown that this disparity is detectable in both phase and amplitude spectra and occurs in both GAN-generated and Diffusion-generated images. There are also studies proposing to address the generative frequency bias via a frequency-domain regularizer or a frequency loss (Durall, Keuper, and Keuper 2020; Jiang et al. 2021). But such methods require retraining the generative model, and thus are less practical for medical GDA considering cost-effectiveness and data availability.

## The Frequency Recalibration Method

This study is motivated by the intriguing generative frequency bias phenomenon, with a goal to answer the following questions: 1) *will the frequency gap between AI-generated and real images affect GDA in medical AI?* 2) *If so, is there a cost-effective, universal method to close the frequency gap to improve GDA for various medical image classification tasks?*

Since prior research shows that frequency abnormalities in AI-generated images often appear in higher frequency components, a pattern also observed in medical images (see Fig.1), a potential improvement on GDA is post-processing AI-synthesized samples by recalibrating their high-frequency distribution to match real images. To address this, We propose the Frequency Recalibration Method (FreRec) involving a two-phase manipulation: first, Statistic High-frequency Replacement (SHR) replaces the high-frequency components of synthetic images with the average from real images, achieving initial alignment; second, Reconstructive High-frequency Mapping (RHM) further mapping the high-frequency distribution to that of real images and restoring image quality through learnable unidirectional reconstruction. Together, these steps provide a coarse-to-fine calibration of frequency distributions for AI-synthesized medical images.

### Statistic High-Frequency Replacement

Since AI-synthesized samples implicitly differ from real samples in higher-frequency components (Durall, Keuper, and Keuper 2020), a rough distributional alignment can be

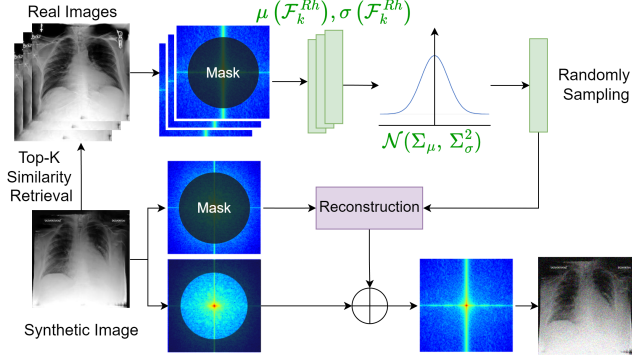


Figure 2: The workflow of Statistic High-frequency Replacement.

achieved via replacing the high-frequency components of AI samples with the counterparts of real samples by statistic.

Let  $x_i^S \in \mathbb{R}^{N \times N \times C}$  be a synthesized image, its Fourier transformation can be denoted as  $\mathcal{F}(x_i^S)(u, v, c) = \sum_{h=0}^{N-1} \sum_{w=0}^{N-1} x_i(h, w, c) e^{-j2\pi(\frac{h}{N}u + \frac{w}{N}v)}$ . A binary mask  $\mathcal{M} \in \mathbb{R}^{r \times r}$ , whose values are zero except for the center region with a fixed ratio  $r$ , is applied to the centered Fourier spectrum to separated it into the low-frequency component  $\mathcal{F}^l(x_i^S)$  and the high-frequency component  $\mathcal{F}^h(x_i^S)$ . The same operation can be applied to a real medical image  $x_i^R$  for  $\mathcal{F}^l(x_i^R)$  and  $\mathcal{F}^h(x_i^R)$ , and on this basis the high-frequency components can be exchanged.

However, one-to-one replacement is prone to randomness that fails to ensure distributional alignment. Therefore, as shown in Fig. 2, for each synthetic image, we select a batch of real images for a statistic-based replacement. Additionally, to minimize perturbation from rough alignment and preserve the semantic features of the modified synthetic image, we retrieve the top-K real images  $\{x_k^R\}_{k=1}^K$  that are similar as  $x_i^S$  by Structural Similarity (SSIM) score for replacement. Then, given  $\{x_k^R\}$ , the channel-wise mean and standard deviation of high-frequency spectral distribution are computed as (for simplicity, we denote  $\mathcal{F}^h(x_k^R)$  as  $\mathcal{F}_k^{Rh}$ ):

$$\begin{aligned} \mu(\mathcal{F}_k^{Rh}) &= \frac{1}{H^2} \sum_{u=1}^H \sum_{v=1}^H \mathcal{F}_k^{Rh}(u, v, c) \\ \sigma(\mathcal{F}_k^{Rh}) &= \frac{1}{H^2} \sum_{u=1}^H \sum_{v=1}^H \sqrt{[\mathcal{F}_k^{Rh}(u, v, c) - \mu(\mathcal{F}_k^{Rh})]^2}. \end{aligned} \quad (1)$$

Assuming the frequency components at each spectral band in a batch of independent images follow a Gaussian distribution (see Supplement for justification), the statistical variances are calculated as follows:

$$\begin{aligned} \Sigma_\mu(\mathcal{F}_k^{Rh}) &= \sqrt{\frac{1}{K} \sum [\mu(\mathcal{F}_k^{Rh}) - \mathbb{E}[\mu(\mathcal{F}_k^{Rh})]]^2} \\ \Sigma_\sigma(\mathcal{F}_k^{Rh}) &= \sqrt{\frac{1}{K} \sum [\sigma(\mathcal{F}_k^{Rh}) - \mathbb{E}[\sigma(\mathcal{F}_k^{Rh})]]^2}. \end{aligned} \quad (2)$$

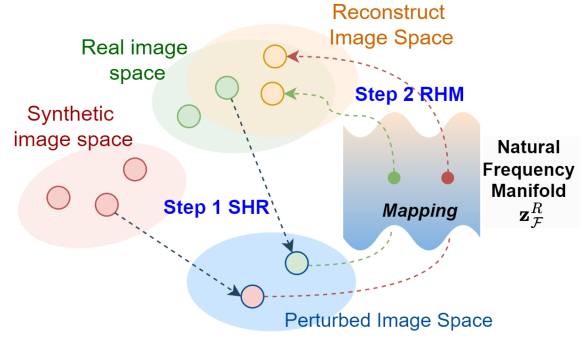


Figure 3: A conceptual explanation of the Reconstructive High-frequency Mapping. Images are first transformed into the same starting space during the initial alignment by SHR. Then the synthetic images can be further calibrated by mapping onto the natural frequency manifold following the same reconstruction path learned from real images.

Then the Gaussian distribution for probabilistic statistics of high-frequency components of  $\{\mathcal{F}_k^R\}_{k=1}^K$  can be modeled. With this Gaussian distribution, we can randomly sample new mean  $\hat{\mu}$  and standard deviation  $\hat{\sigma}$  to modify the original high-frequency component  $\mathcal{F}_i^{Sh}$  of the synthetic image  $x_i^S$ :

$$\begin{aligned} \hat{\mu}(\mathcal{F}_k^{Rh}) &= \mu(\mathcal{F}_k^{Rh}) + \epsilon_\mu \Sigma_\mu(\mathcal{F}_k^{Rh}), \epsilon_\mu \sim \mathcal{N}(0, 1), \\ \hat{\sigma}(\mathcal{F}_k^{Rh}) &= \sigma(\mathcal{F}_k^{Rh}) + \epsilon_\sigma \Sigma_\sigma(\mathcal{F}_k^{Rh}), \epsilon_\sigma \sim \mathcal{N}(0, 1), \\ \hat{\mathcal{F}}_i^{Sh} &= \hat{\sigma}(\mathcal{F}_k^{Rh}) \left( \frac{\mathcal{F}_i^{Sh} - \mu(\mathcal{F}_k^{Rh})}{\sigma(\mathcal{F}_k^{Rh})} \right) + \hat{\mu}(\mathcal{F}_k^{Rh}). \end{aligned} \quad (3)$$

By combining the above  $\hat{\mathcal{F}}_i^{Sh}$  with the original low-frequency component  $\mathcal{F}^{Sl}$ , followed by the inverse Fourier transformation, a lossy calibrated sample  $\hat{x}_i^S$  is obtained.

### Reconstructive High-Frequency Mapping

Statistic High-frequency Replacement provides coarse frequency alignment but causes spectral distortion and reduces image quality. Therefore, a further step is needed to restore quality and recover high-frequency details resembling real images from previous perturbed versions. A simple denoising reconstruction of the perturbed synthetic samples via a direct mapping  $\hat{x}_i^S \rightarrow x_i^S$  is ineffective, as it reverts back to the original, frequency-misaligned synthetic image. Instead, we address this using unidirectional manifold mapping, i.e., learning a latent natural frequency manifold ( $\mathbf{z}_{\mathcal{F}}^R$ ) from real images, and then projecting synthetic images onto it (see Fig. 3). Reconstruction from this shared latent space  $\mathbf{z}$  restores high-frequency details, aligning synthetic images with the natural frequency distribution of real images.

To capture the natural frequency manifold  $\mathbf{z}_{\mathcal{F}}^R$ , we first train a denoising auto-encoder  $\mathcal{A} : \hat{x}_i^R \rightarrow x_i^R$ , where  $\hat{x}_i^R$  is a perturbed version of  $x_i^R$ . Importantly,  $\mathcal{A}$  must be **trained exclusively on real images** to form the accurate frequency reconstruction direction towards the true frequency

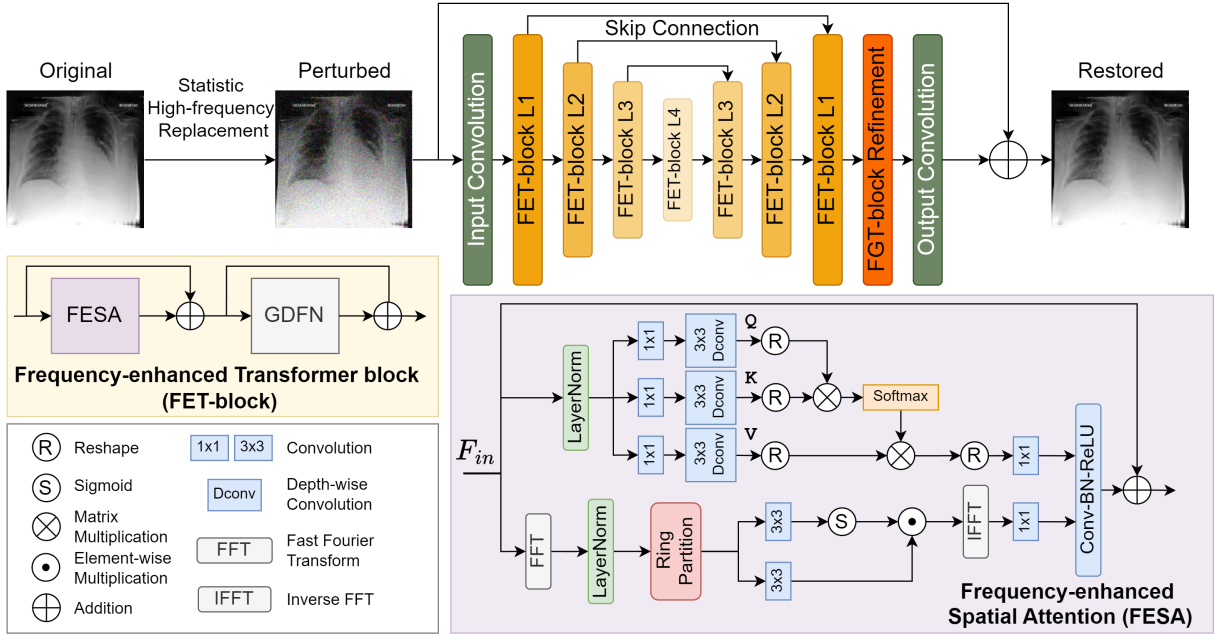


Figure 4: The details of the denoising auto-encoder used in Reconstructive High-frequency Mapping.

distribution of real images. Additionally, the input real image  $\hat{x}_i^R$  is perturbed using the same SHR method (Eq.3) as well, to ensure a consistent projection path where both real and synthetic images are reconstructed from the same input space, as explained in Fig 3. After training, we apply the well-trained auto-encoder  $\mathcal{A}^*$  to synthetic images for high-frequency reconstruction, during which their frequency distribution is calibrated to real images’.

**The Backbone of  $\mathcal{A}$**  Designing  $\mathcal{A}$  is challenging, as it must accurately capture the frequency and visual details of real images without introducing additional frequency distortion by itself. We address this through both model architecture and loss function design. Specifically, we propose a transformer network based on the Restoration Transformer (Restormer) (Zamir et al. 2022), as shown in Fig 4, replacing its transformer blocks with our Frequency-enhanced Transformer blocks (FET-blocks) to better learn high-frequency details. Each FET-block incorporates a novel Frequency-enhanced Spatial Attention (FESA) module and the original Gated-Dconv Feed-forward Network (GDFN) from Restormer. The network contains 2, 4, 6, and 8 FET-blocks in levels 1 through 4, respectively, with two additional FET-blocks in the refinement stage.

FESA is built on a cross-attention mechanism that fuses the global spatial self-attention branch with a local frequency self-attention branch. The global spatial self-attention branch takes the RGB feature as input. The layer-normalized feature  $F_{rgb} \in \mathbb{R}^{H \times W \times C}$  is split into query ( $Q$ ), key ( $K$ ) and value ( $V$ ) projections using  $1 \times 1$  convolutions to aggregate pixel-wise local context followed by  $3 \times 3$  depth-wise convolutions to encode channel-wise spatial context (Zamir et al. 2022). The query and key projections are then reshaped for dot-product, yielding a transposed-

attention map of size  $\mathbb{R}^{C \times C}$ . Overall, the global spatial self-attention can be denoted as:

$$\hat{\mathbf{F}}_{rgb} = C_1(\text{softmax}(Q \cdot K/a) \cdot V) \quad (4)$$

where  $\hat{\mathbf{F}}_{rgb}$  is the output feature map;  $Q \in \mathbb{R}^{HW \times C}$ ,  $K \in \mathbb{R}^{C \times HW}$  and  $V \in \mathbb{R}^{HW \times C}$  reshaped tensors.  $a$  is a learnable scaling parameter to control the magnitude of the dot product.  $C_n$  indicates  $n \times n$  convolution.

The local frequency self-attention branch enriches the global learning by embedding additional local frequency information into the spatial attention. The RGB feature is transformed into the frequency domain by Fast Fourier Transform (FFT) to obtain its amplitude spectrum as input. To further disentangle frequency components, the average-pooled spectrum feature  $\mathbf{F}_{spe} \in \mathbb{R}^{2R \times 2R \times 1}$  is then partitioned into many non-overlapping circular rings with a width of  $d$  along the radius dimension, where  $d$  defines the range of frequency components. We group the features located in the same circular ring into one channel and obtain  $\mathbf{F}_{fre} \in \mathbb{R}^{P \times R/d}$ , where  $P$  denotes the number of frequency components on each channel. Then, we use two  $3 \times 3$  convolutions. The first convolution outputs features into a sigmoid function to obtain a local attention map, as sigmoid gates are often used to control the local flow of information without adding a large computational overhead (Hochreiter 1997). The second convolution is used to refine the spectrum features, followed by a point-wise multiplication with the local attention map and inverse FFT. The local frequency self-attention can be denoted as:

$$\hat{\mathbf{F}}_{fre} = C_1(\text{iFFT}(C_3(\delta(\mathbf{F}_{fre})) \cdot C_3(\mathbf{F}_{fre}))) \quad (5)$$

Finally, the global RGB feature and local frequency feature are fused by concatenation and a  $1 \times 1$ -Conv-BN-ReLU

Datasets		GenAIs		Training Set				Test Set	
Task	Dataset	Type	Source	Real		Synthetic		Real	
				Positive	Negative	Positive	Negative	Positive	Negative
Brain tumor MRI	Kaggle-Brain tumor	FastGAN	Pre-trained	1000	1000	1500	1500	300	300
DR Fundus image	Kaggle-DR	VC-Diffusion	Pre-trained	7000	20000	3000	-	3000	6000
Cardiomegaly X-ray	MIMIC-CXR	StyleGAN3	From scratch	800	350	800	1250	100	100

Table 1: Details of datasets and settings.

block. After that, the original RGB features  $\mathbf{F}_{rgb}$  are added, resulting in the refined features:

$$\hat{\mathbf{F}}_{final} = \mathbf{F}_{rgb} + \text{ReLU-BN-Conv}(\hat{\mathbf{F}}_{rgb} \parallel \hat{\mathbf{F}}_{fre}), \quad (6)$$

where  $\parallel$  indicated feature concatenation.

**Loss Function** In addition to the standard pixel-level similarity loss for image reconstruction, we introduce a spectral similarity loss to further ensure that  $\mathcal{A}$  can accurately learn frequency and visual details of real images, resulting in a joint loss function:

$$\min_{\mathcal{A}} \mathcal{L} := \underbrace{\|x_i^R - \mathcal{A}(\hat{x}_i^R)\|^2}_{\text{Pixel similarity}} + \underbrace{\|\mathcal{F}(x_i^R) - \mathcal{F}(\mathcal{A}(\hat{x}_i^R))\|^2}_{\text{Frequency similarity}} \quad (7)$$

### Incorporating FreRec with GDA

Due to the inevitable sub-optimality of  $\mathcal{A}^*$ , subtle frequency differences may remain between its reconstructions and the original images, causing the reconstructed image space to shift slightly from the real image space (see Fig.3). Thus, to achieve the best practice, we recommend implementing FreRec as a unified, plug-and-play pre-processing module in both training and inference phases of downstream disease classifiers. That is, all training and testing samples, synthetic and real, should be processed with FreRec to make the final frequency distributions unified throughout the disease classification workflow.

## Experiments

**Datasets** To maintain experimental diversity, we evaluated three medical classification tasks based on different public image datasets: brain tumor detection (brain MRIs) (Bhuvaji et al. 2020), cardiomegaly diagnosis (chest X-rays) (Johnson et al. 2019), and diabetic retinopathy (DR) classification (fundus photography) (Dugas et al. 2015). The generative models employed for generative data augmentation (GDA) are varied by model type and source: a pretrained GAN (FastGAN (Liu et al. 2020)) for the synthesis of brain tumor MRI images, a pretrained diffusion model (VC-Diffusion (Ilanchezian et al. 2023)) for DR fundus images, and a from-scratch GAN (StyleGAN3 (Karras et al. 2021)) for cardiomegaly X-rays. Both pretrained models have demonstrated GDA effectiveness in their original studies. Figure 5 presents examples of real and synthetic samples from three datasets, while Table 1 provides a summary of the dataset configurations. We perform GDA to enhance the training sample size while striving to balance the distribution of disease classes as effectively as possible.




Figure 5: Examples of original real images (above) and synthetic images (bottom) from three datasets. The synthetic images maintain high visual quality and fidelity.

**Baselines and Settings** We compared no augmentation (RAW), GDA without frequency alignment (GDA), GDA with FreRec (GDA+FreRec). In addition, we also assessed three image processing-based augmentation methods, AutoAug (Cubuk et al. 2019), Mix-up (Psroudakis and Kollias 2022), and Fourier-basis Augmentation (AFA) (Vaish, Wang, and Strisciuglio 2024), plus Domain Gap Embeddings (DoGE), a recent domain bias reduction method for AI-synthesized images (Wang et al. 2024d). Different classifier backbones, including ResNet50, DenseNet, and ViT-B-16, were evaluated for disease classification. Classification accuracy (Acc.), F1 and AUC scores were the evaluation metrics, with each classifier evaluated five times and the average results are reported to reduce randomness. The mask ratio  $r$  and the sample number  $k$  of SHR was set to 0.5 and 200 respectively, considering the trade-off between alignment effect, image quality and running time (see Supplement for the decision process).

**Disease Classification Results** Table 2 shows the classification results of three diseases under various augmentation strategies. Normal augmentations (AutoAug, Mix-up, and AFA) slightly improve performance over raw training, where AutoAug and AFA are more reliable than Mix-up, which occasionally degrades performance. The plain generative data augmentation, GDA, also suffers from unreliability, sometimes having a negative impact on classifiers despite significantly expanding the training set with synthetic images. For instance, GDA decreases accuracy and F1 scores of cardiomegaly-ResNet50 and all brain tumor classifiers compared to raw training. This observation aligns with prior findings that AI-synthesized samples do not always improve downstream tasks (Singh et al. 2024). In comparison, recalibrating frequency with FreRec (GDA+FreRec) significantly improves the augmentation effect of GDA, mak-

ing GDA more reliable. GDA+FreRec consistently improves raw classifiers across all tests and metrics, demonstrating its effectiveness as an augmentation strategy. This supports the hypothesis that the frequency distributional gap between real and synthetic samples is a key cause of GDA instability, and closing this gap via frequency calibration improves reliability. In general, GDA+FreRec outperforms baseline augmentation strategies in all tasks except for DoGE in DR classification. This is because DoGE directly reduces domain bias in synthetic images via source-to-target domain adaptation. However, it heavily relies on a fixed, known synthetic image source, rendering DoGE only applicable to specific generation models. In contrast, FreRec is trained solely on real images, enabling a unidirectional mapping of synthetic images from unknown and arbitrary generation models to real ones, suggesting greater adaptability to unseen synthetic images.

**Ablation Study** We compare the complete FreRec with sole SHR and sole RHM, with and without FESA (i.e., the original Restormer), in terms of classification performance and reconstruction quality. As shown in Table 3, FreRec consistently outperforms its partial variants in all classification tasks. While RHM without FESA yields the best image quality due to its 1:1 reconstruction, the quality difference between it and FreRec is minor, indicating FreRec achieves frequency alignment at a negligible cost of image quality. Notably, sole SHR does not improve classification performance despite roughly aligning frequency distributions, as the alignment process significantly distorts image quality. Therefore, the RHM step is essential for fine-grained calibration and restoration of image details. Reconstructed synthetic samples can be visualized in the Supplement.

### Frequency Recalibration Effectiveness

**Distribution Visualization** To confirm the effectiveness of FreRec in aligning frequency distributions of synthetic and real images, we present a visualization of their frequency distributions before and after recalibration. The frequency distribution is computed as a one-dimensional profile by azimuthally integrating the spectral magnitudes over radial frequencies  $\theta$  (Durall, Keuper, and Keuper 2020). The details of computing the one-dimensional profile can be found in Supplement.

Figure 6 illustrates the average frequency distributional comparisons across three datasets. Prior to applying FreRec, substantial frequency discrepancies were observed between real and synthetic images in all datasets, despite the synthetic images spanning three distinct modalities and being generated by different AI models. After recalibration, these distributional gaps were dynamically reduced across all datasets, leading to improved alignment in frequency distributions. Notably, the alignments in the Brain Tumor MRI and Cardiomegaly X-ray datasets are more thorough compared to those in DR Fundus images. This discrepancy may be attributed to the fact that DR Fundus images are color photographs, whereas the others are grayscale. Color images inherently contain richer pixel information, making the learning of Reconstructive High-frequency Mapping more challenging for the denoising auto-encoder  $\mathcal{A}$ .

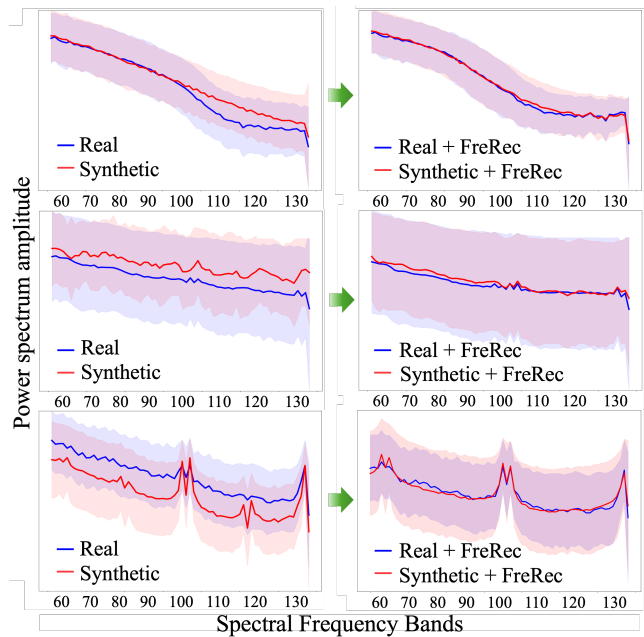


Figure 6: Frequency distributions of real and synthetic images before (the left column) and after (the right column) frequency recalibration in Brain Tumor MRI, Cardiomegaly X-ray, and DR Fundus image datasets (from top to bottom).




Figure 7: Feature-space visualization of real and synthetic images before (the left column) and after (the right column) frequency recalibration in Brain Tumor MRI, Cardiomegaly X-ray, and DR Fundus image datasets (from top to bottom). Orange dots: Real samples; Blue dots: synthetic samples.

		DenseNet			ResNet50			ViT		
		AUC	Acc.	F1	AUC	Acc.	F1	AUC	Acc.	F1
Card	RAW	0.842	0.803	0.787	0.834	0.792	0.788	0.832	0.794	0.779
	AutoAug	0.853	0.814	0.798	0.847	0.803	0.797	0.843	0.804	0.787
	Mix-up	0.833	0.812	0.795	0.828	0.804	0.787	0.823	0.808	0.792
	AFA	0.855	0.808	0.800	0.831	0.822	0.815	0.859	0.819	0.817
	DoGE	0.880	0.833	0.833	0.869	0.821	<b>0.836</b>	0.851	0.818	0.822
	GDA	0.871	0.803	0.804	0.834	0.782	0.783	0.848	0.813	0.814
	GDA+FreRec	<b>0.899</b>	<b>0.848</b>	<b>0.834</b>	<b>0.888</b>	<b>0.832</b>	<u>0.834</u>	<b>0.888</b>	<b>0.838</b>	<b>0.832</b>
DR	RAW	0.840	0.803	0.793	0.843	0.792	0.783	0.834	0.794	0.782
	AutoAug	0.857	0.821	0.809	0.861	0.808	0.799	0.848	0.811	0.798
	Mix-up	0.852	0.791	0.804	0.831	0.780	0.794	0.847	0.783	0.793
	AFA	0.861	0.805	0.810	0.855	0.801	0.788	0.837	0.801	0.800
	DoGE	<b>0.880</b>	<b>0.841</b>	0.816	0.871	<b>0.825</b>	<b>0.819</b>	<b>0.862</b>	0.811	<b>0.810</b>
	GDA	0.863	0.813	0.804	0.848	0.803	0.792	0.834	0.783	0.773
	GDA+FreRec	<u>0.879</u>	<u>0.834</u>	<b>0.823</b>	<b>0.878</b>	<u>0.823</u>	<u>0.813</u>	<u>0.852</u>	<b>0.813</b>	<u>0.803</u>
BT	RAW	0.840	0.742	0.733	0.793	0.783	0.772	0.753	0.723	0.713
	AutoAug	0.833	0.759	0.740	0.811	0.799	0.779	0.771	0.731	0.729
	Mix-up	0.853	0.731	0.744	0.782	0.794	0.761	0.764	0.734	0.702
	AFA	0.841	0.824	0.730	0.807	0.814	0.766	0.771	0.728	0.739
	DoGE	<b>0.859</b>	0.833	0.742	0.811	0.822	0.780	0.785	0.733	0.747
	GDA	0.794	0.732	0.713	0.783	0.723	0.723	0.758	0.732	0.743
	GDA+FreRec	<u>0.855</u>	<b>0.834</b>	<b>0.753</b>	<b>0.843</b>	<b>0.843</b>	<b>0.783</b>	<b>0.787</b>	<b>0.737</b>	<b>0.752</b>

Table 2: Disease classification results in three datasets using different augmentation strategies. Bold and underline indicate the best and second best results in each group, respectively. Card = cardiomegaly; DR = Diabetic Retinopathy; BT = Brain Tumor.

		AUC	Acc.	PSNR	SSIM
Card	SHR	0.81	0.79	25.10	0.76
	RHM w.o. FESA	0.85 $\uparrow$ 0.04	0.79 $\uparrow$ 0.01	36.44 $\uparrow$ 11.34	0.98 $\uparrow$ 0.21
	RHM w. FESA	0.87 $\uparrow$ 0.02	0.82 $\uparrow$ 0.03	35.51 $\downarrow$ 0.93	0.96 $\downarrow$ 0.02
	FreRec	0.89 $\uparrow$ 0.03	0.84 $\uparrow$ 0.02	35.62 $\uparrow$ 0.11	0.95 $\uparrow$ 0.00
DR	SHR	0.81	0.77	25.22	0.76
	RHM w.o. FESA	0.85 $\uparrow$ 0.04	0.81 $\uparrow$ 0.05	36.33 $\uparrow$ 11.11	0.97 $\uparrow$ 0.22
	RHM w. FESA	0.86 $\uparrow$ 0.01	0.82 $\uparrow$ 0.01	34.23 $\downarrow$ 2.10	0.94 $\downarrow$ 0.03
	FreRec	0.87 $\uparrow$ 0.01	0.82 $\uparrow$ 0.01	34.29 $\uparrow$ 0.06	0.95 $\uparrow$ 0.01
BT	SHR	0.78	0.73	25.95	0.77
	RHM w.o. FESA	0.79 $\uparrow$ 0.01	0.75 $\uparrow$ 0.02	41.23 $\uparrow$ 15.28	0.99 $\uparrow$ 0.22
	RHM w. FESA	0.80 $\uparrow$ 0.01	0.80 $\uparrow$ 0.05	40.21 $\downarrow$ 1.02	0.98 $\downarrow$ 0.01
	FreRec	0.83 $\uparrow$ 0.02	0.81 $\uparrow$ 0.01	41.15 $\uparrow$ 0.94	0.98 $\uparrow$ 0.00

Table 3: Classification performances and image quality of different frequency calibration methods.

**Feature-Space Visualization** To verify that frequency discrepancies between synthetic and real images contribute as a main cause for GDA’s instability in downstream disease classification and that frequency alignment improves its reliability in enriching useful features, we provide a feature-space visualization using T-SNE. For each disease, a ResNet50 classifier is independently trained on 80% of the real images from the training set described in Table 1. The classification head is then removed, turning the classifier into a pre-trained feature encoder to extract features from the remaining 20% of real images and an equal number of randomly selected synthetic images. To ensure unbiased evaluation, this process is performed without regard to class labels. The extracted features are visualized in a two-dimensional space using T-SNE.

Figure 7 presents the clustering results. In the left column, the features of synthetic and real samples are clearly

separated across all datasets, indicating a domain bias in the synthetic samples that can compromise GDA reliability. In the right column, frequency recalibration shifts the synthetic features closer to the real ones. For the Brain Tumor MRI and Cardiomegaly X-ray datasets, the synthetic and real features completely overlap, eliminating domain bias. In the DR dataset, the shift is not complete, likely due to color images have richer pixel information that may complicate the high-frequency reconstruction and induce other contextual biases that frequency recalibration cannot fully address.

## Conclusion

Generative data augmentation is widely used for medical image classification tasks. It can synthesize realistic medical image samples to complement the original dataset. However, the reliability of AI-synthesized samples should be carefully investigated to avoid any domain bias and negative features brought by generative AI models. In this study, we explored this problem from a frequency perspective. We identified that the frequency misalignment between real and synthetic images is a main cause of the instability of GDA. To address this issue, we proposed a novel frequency recalibration method, which consists of two steps, Statistical High-frequency Replacement and Reconstructive High-frequency Mapping, to formulate a coarse-to-fine alignment of the frequency distributions between real and synthetic images. Our extensive experiments performed on various datasets and on different generative AI models demonstrate that the calibrated synthetic samples have a much closer distribution to the real ones, and can significantly improve downstream medical image classifications.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62402009), and the Science and Technology Development Fund of Macao under Grants 0002/2024/RIA1, 0013-2024-ITP1 and 0069/2025/ITP2.

## References

- Bhuvaji, S.; Kadam, A.; Bhumkar, P.; Dedge, S.; and Kanchan, S. 2020. Brain Tumor Classification (MRI).
- Chen, Q.; Chen, X.; Song, H.; Xiong, Z.; Yuille, A.; Wei, C.; and Zhou, Z. 2024. Towards generalizable tumor synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11147–11158.
- Corvi, R.; Cozzolino, D.; Poggi, G.; Nagano, K.; and Verdoliva, L. 2023. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 973–982.
- Cubuk, E. D.; Zoph, B.; Mane, D.; Vasudevan, V.; and Le, Q. V. 2019. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 113–123.
- Dayarathna, S.; Islam, K. T.; Uribe, S.; Yang, G.; Hayat, M.; and Chen, Z. 2024. Deep learning based synthesis of MRI, CT and PET: Review and analysis. *Medical image analysis*, 92: 103046.
- Dugas, E.; Jared, J.; Jorge, J.; and Cukierski, W. 2015. Diabetic Retinopathy Detection. <https://kaggle.com/competitions/diabetic-retinopathy-detection>. Kaggle.
- Durall, R.; Keuper, M.; and Keuper, J. 2020. Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7890–7899.
- Dzanic, T.; Shah, K.; and Witherden, F. 2020. Fourier spectrum discrepancies in deep network generated images. *Advances in neural information processing systems*, 33: 3022–3032.
- Frank, J.; Eisenhofer, T.; Schönherr, L.; Fischer, A.; Kolossa, D.; and Holz, T. 2020. Leveraging frequency analysis for deep fake image recognition. In *International conference on machine learning*, 3247–3258. PMLR.
- Gao, C.; Killeen, B. D.; Hu, Y.; Grupp, R. B.; Taylor, R. H.; Armand, M.; and Unberath, M. 2023. Synthetic data accelerates the development of generalizable learning-based algorithms for X-ray image analysis. *Nature Machine Intelligence*, 5(3): 294–308.
- Guillaudeux, M.; Rousseau, O.; Petot, J.; Bennis, Z.; Dein, C.-A.; Goronflot, T.; Vince, N.; Limou, S.; Karakachoff, M.; Wargny, M.; et al. 2023. Patient-centric synthetic data generation, no reason to risk re-identification in biomedical data analysis. *NPJ Digital Medicine*, 6(1): 37.
- He, M.; Li, Z.; Liu, C.; Shi, D.; and Tan, Z. 2020. Deployment of artificial intelligence in real-world practice: opportunity and challenge. *Asia-Pacific Journal of Ophthalmology*, 9(4): 299–307.
- Hochreiter, S. 1997. Long Short-term Memory. *Neural Computation MIT-Press*.
- Ilanchezian, I.; Boreiko, V.; Kühlewein, L.; Huang, Z.; Ayyhan, M. S.; Hein, M.; Koch, L.; and Berens, P. 2023. Generating Realistic Counterfactuals for Retinal Fundus and OCT Images using Diffusion Models. arXiv:2311.11629.
- Jiang, L.; Dai, B.; Wu, W.; and Loy, C. C. 2021. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF international conference on computer vision*, 13919–13929.
- Johnson, A. E. W.; Pollard, T. J.; Berkowitz, S. J.; Greenbaum, N. R.; Lungren, M. P.; Deng, C.-y.; Mark, R. G.; and Horng, S. 2019. MIMIC-CXR, a de-identified Publicly Available Database of Chest Radiographs with Free-Text Reports. *Scientific Data*, 6(1): 317.
- Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-Free Generative Adversarial Networks. In *Proc. NeurIPS*.
- Ktena, I.; Wiles, O.; Albuquerque, I.; Rebuffi, S.-A.; Tanno, R.; Roy, A. G.; Azizi, S.; Belgrave, D.; Kohli, P.; Cemgil, T.; et al. 2024. Generative models improve fairness of medical classifiers under distribution shifts. *Nature Medicine*, 1–8.
- Li, Z.; Keel, S.; Liu, C.; He, Y.; Meng, W.; Scheetz, J.; Lee, P. Y.; Shaw, J.; Ting, D.; Wong, T. Y.; et al. 2018. An automated grading system for detection of vision-threatening referable diabetic retinopathy on the basis of color fundus photographs. *Diabetes care*, 41(12): 2509–2516.
- Litjens, G.; Kooi, T.; Bejnordi, B. E.; Setio, A. A. A.; Ciompi, F.; Ghafoorian, M.; Van Der Laak, J. A.; Van Ginneken, B.; and Sánchez, C. I. 2017. A survey on deep learning in medical image analysis. *Medical image analysis*, 42: 60–88.
- Liu, B.; Zhu, Y.; Song, K.; and Elgammal, A. 2020. Towards faster and stabilized gan training for high-fidelity few-shot image synthesis. In *International conference on learning representations*.
- Liu, H.; Li, X.; Zhou, W.; Chen, Y.; He, Y.; Xue, H.; Zhang, W.; and Yu, N. 2021. Spatial-Phase Shallow Learning: Rethinking Face Forgery Detection in Frequency Domain. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 772–781. Los Alamitos, CA, USA: IEEE Computer Society.
- Psaroudakis, A.; and Kollias, D. 2022. MixAugment & Mixup: Augmentation Methods for Facial Expression Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2367–2375.
- Rahaman, N.; Baratin, A.; Arpit, D.; Draxler, F.; Lin, M.; Hamprecht, F.; Bengio, Y.; and Courville, A. 2019. On the spectral bias of neural networks. In *International conference on machine learning*, 5301–5310. PMLR.
- Rajpurkar, P.; Chen, E.; Banerjee, O.; and Topol, E. J. 2022. AI in health and medicine. *Nature medicine*, 28(1): 31–38.
- Shang, F.; Fu, J.; Yang, Y.; Huang, H.; Liu, J.; and Ma, L. 2024. SynFundus-1M: A High-quality Million-scale Synthetic fundus images Dataset with Fifteen Types of Annotation. arXiv:2312.00377.

- Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Papernot, N.; Anderson, R.; and Gal, Y. 2024. AI models collapse when trained on recursively generated data. *Nature*, 631(8022): 755–759.
- Singh, K.; Navaratnam, T.; Holmer, J.; Schaub-Meyer, S.; and Roth, S. 2024. Is Synthetic Data all We Need? Benchmarking the Robustness of Models Trained with Synthetic Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2505–2515.
- Vaish, P.; Wang, S.; and Strisciuglio, N. 2024. Fourier-Basis Functions to Bridge Augmentation Gap: Rethinking Frequency Augmentation in Image Classification. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17763–17772.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8684–8694.
- Wang, J.; Wang, K.; Yu, Y.; Lu, Y.; Xiao, W.; Sun, Z.; Liu, F.; Zou, Z.; Gao, Y.; Yang, L.; et al. 2024a. Self-improving generative foundation model for synthetic medical image generation and clinical applications. *Nature Medicine*, 1–9.
- Wang, Y.; Liu, C.; Hu, W.; Luo, L.; Shi, D.; Zhang, J.; Yin, Q.; Zhang, L.; Han, X.; and He, M. 2024b. Economic evaluation for medical artificial intelligence: accuracy vs. cost-effectiveness in a diabetic retinopathy screening case. *NPJ Digital Medicine*, 7(1): 43.
- Wang, Y.; Liu, C.; Zhou, K.; Zhu, T.; and Han, X. 2024c. Towards regulatory generative AI in ophthalmology healthcare: a security and privacy perspective. *British Journal of Ophthalmology*, 108(10): 1349–1353.
- Wang, Y. O.; Chung, Y.; Wu, C. H.; and De la Torre, F. 2024d. Domain Gap Embeddings for Generative Dataset Augmentation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 28684–28694.
- Xu, Z.-Q. J.; Zhang, Y.; Luo, T.; Xiao, Y.; and Ma, Z. 2019. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5728–5739.