# Reconstructing degraded SPECT myocardial images via deep biophysical models: A modern computational approach

Nguyen Thanh Trung*

Medical Equipment Department, 108 Military Central Hospital, 1 Tran Hung Dao, Hai Ba Trung, Hanoi, Vietnam.

*Corresponding author: trung.ntc10@benhvien108.vn

*Received 28 Jul. 2025; Revised 9 Sep. 2025; Accepted 20 Sep. 2025; Published 2 Oct. 2025.*

**DOI:** https://doi.org/10.54939/1859-1043.j.mst.106.2025.55-62

## ABSTRACT

*Myocardial perfusion imaging (MPI) using single-photon emission computed tomography (SPECT) is a critical tool for diagnosing coronary artery disease, but it is often affected by signal degradation due to soft tissue attenuation. In this study, we utilize a publicly available SPECT MPI dataset to establish a benchmark for the task of attenuation correction (AC) by reconstructing AC images from non-attenuation corrected (NC) inputs in a 2D slice-to-slice manner. We implement and compare the performance of several advanced generative models, including generative adversarial networks (GANs) and diffusion models. These models are trained on both general-domain and medical-domain data to evaluate their reconstruction capabilities. The results show that modern deep learning approaches can effectively generate high-quality AC images, demonstrating promising potential for integration into computer-aided diagnosis (CAD) systems for SPECT MPI.*

**Keywords:** Single-photon emission computed tomography; Myocardial perfusion imaging; Computer-aided diagnosis; Generative adversarial network; Diffusion model.

## 1. INTRODUCTION

The identification of coronary artery disease (CAD) remains a cornerstone of cardiovascular care, with myocardial perfusion imaging (MPI) via SPECT serving as one of the most commonly applied non-invasive diagnostic techniques. Nevertheless, the diagnostic reliability of SPECT MPI is frequently compromised by attenuation artifacts, which arise from heterogeneous soft tissue absorption. Such effects are particularly pronounced in women (due to breast tissue), men (due to diaphragmatic interference), and patients with elevated body mass index, ultimately lowering clinical accuracy. The most established strategy to mitigate this issue involves hybrid SPECT/CT systems, where CT-derived attenuation maps are employed for attenuation correction (AC). Although effective, this solution entails increased financial burden, greater infrastructural demands, longer acquisition times, and additional radiation exposure—factors that hinder widespread adoption in many healthcare environments.

In recent years, deep learning has emerged as a promising alternative for attenuation correction, with generative frameworks such as GANs and diffusion models enabling direct generation of AC-equivalent images from NC inputs, thus reducing reliance on CT. While encouraging results have been reported, large-scale validations of robustness and clinical applicability remain limited. Several strategies have been proposed to enhance SPECT myocardial perfusion imaging. Liu et al. [1] introduced a post-reconstruction correction framework using deep learning–generated attenuation maps, improving quantitative accuracy for SPECT-only systems; Chen et al. [2] presented DuDoSS, a dual-domain sinogram synthesis method for reconstructing full-view projections from sparse data, enabling faster acquisitions with preserved quality; and Liu et al. [3] employed noise-to-noise training with a coupled U-Net, achieving superior denoising and improved perfusion defect detection over Gaussian smoothing. Despite these advances, clinical translation is still constrained by limited dataset diversity, scanner generalizability, and model

interpretability, underscoring the need for more robust and adaptable approaches. In this work, we utilize a publicly available, standardized SPECT MPI dataset that has been employed in prior investigations [4-6] to perform an extensive benchmarking study on slice-wise NC-to-AC translation. We implement and evaluate several state-of-the-art generative models for image-to-image translation, encompassing both general-purpose architectures and those tailored to the medical imaging domain. Our quantitative and qualitative analyses demonstrate the strong capability of deep learning-based approaches to address attenuation artifacts in SPECT MPI and provide a solid foundation for future CAD research.

## 2. RELATED WORK

The reconstruction of attenuation-corrected (AC) images from non-corrected (NC) SPECT data has strong clinical potential. Advances in generative models for image-to-image translation have produced two main research directions: (i) models developed on general-domain datasets, forming the basis of modern architectures, and (ii) specialized approaches in medical imaging. In natural image translation, Pix2Pix [7] pioneered conditional GANs with paired data, followed by CycleGAN [8] and UNIT [9] for unpaired settings. More recently, diffusion models such as DDPM [4], Palette [6], and Latent Diffusion Models (LDM) [5] demonstrated stable training and high-fidelity outputs, while BBDM [10] introduced bidirectional diffusion for direct domain mapping without paired data. Despite impressive results, these models are still largely restricted to natural images.

In medical imaging, GAN-based methods like MedGAN [11] and RegGAN [12] have shown promise in PET-to-CT, CT-to-MRI, and related translation tasks, with RegGAN incorporating registration for improved realism. Diffusion-based methods are emerging, such as CPDM [15], which performs CT-to-PET translation guided by attention and attenuation maps. While these approaches achieve encouraging results, limitations remain in their dependence on paired datasets and challenges in producing clinically reliable, diagnostically accurate reconstructions.

## 3. PROBLEM

### 3.1. Theoretical foundation

Attenuation correction addresses artifacts in medical imaging caused by non-uniform absorption of radioactive photons in soft tissues, which in SPECT MPI for coronary artery disease can compromise image quality and diagnostic accuracy, particularly near the diaphragm, breast tissue, or in certain patients. The task aims to reconstruct attenuation-corrected (AC) images from non-corrected (NC) inputs, improving clarity and reliability without requiring additional CT hardware. Existing methods fall into two categories: traditional iterative reconstruction (IR-based) approaches and deep learning-based methods, the latter dominated by GAN- and diffusion-based frameworks. In this study, we focus on the deep learning paradigm and introduce a comprehensive benchmark to evaluate state-of-the-art generative models for attenuation correction.

### 3.2. GAN-based Models

GAN-based models form a powerful framework for image generation by establishing a min-max game between two competing networks with opposing objectives. This architecture comprises two main components: a generator (G) and a discriminator (D). In the context of attenuation correction, generator G aims to learn the data distribution that maps non-attenuation corrected images $I_{NC}$ to attenuation-corrected images $I_{AC}$. These two networks are trained simultaneously: the discriminator seeks to maximize its ability to distinguish real images from those generated by G, while the generator attempts to produce synthetic images that are indistinguishable from real ones. The standard objective function for GAN-based models is defined in equation 1.

$$L_{GAN} = E[logD(I_{AC})] + E\left[\log\left(1 - D\left(G(I_{NC})\right)\right)\right] \quad (1)$$

Several representative image-to-image translation models are included in our benchmark. In the general domain, Pix2Pix [4], CycleGAN [5], and UNIT [8] are selected, representing supervised, unsupervised, and hybrid translation paradigms. For medical imaging, MedGAN [6] and RegGAN [9] are considered due to their effectiveness in cross-modality translation (e.g., MRI ↔ CT, PET ↔ MRI) and image enhancement tasks, where they have demonstrated the ability to generate high-quality synthetic images that preserve structural fidelity and support diagnostic performance.

### 3.3. Diffusion-based models

Diffusion-based models simulate a forward noise process and learn to reverse it for data generation. Foundational methods such as DDPM [4] and score-based models have shown strong performance in high-fidelity synthesis. Building on this, the Brownian Bridge Diffusion Model (BBDM) [10] leverages the Brownian Bridge process to achieve effective image-to-image translation conditioned on input images.

*3.3.1. Forward process (Brownian Bridge)*

The forward process employed here is modeled as a Brownian Bridge, which interpolates between the target image $x_0 := I_{AC}$, we perform a diffusion process toward the source image (non-corrected image) $y := I_{NC}$ over T step. At step t, the intermediate state $x_t$ is computed as follows in equation 2. At each diffusion step t ∈{1,…,T}, the intermediate latent variable $x_t$ is constructed as a convex combination of $x_0$ and y, with added Gaussian noise, shows in equation 2.

$$x_t = (1 - m_t)x_0 + m_t y + \sqrt{\delta_t}\epsilon_t \quad (2)$$

where $m_t = \frac{t}{T}$, T is the total number of diffusion steps, $\delta_t$ is the variance at step t and $\epsilon_t \sim N(0, 1)$. Thus, the forward process is defined as:

$$q_{BB}(x_t|x_0, y) = N(x_t; (1 - m_t)x_0 + m_t y, \delta_t I) \quad (3)$$

which corresponds to a Brownian Bridge with fixed endpoints $x_0$ and y. This design enables smooth, controllable interpolation in the latent space, providing a structured forward process that guides reverse generation or reconstruction. Accordingly, the marginal distribution of $x_t$ , conditioned on both endpoints, is given by:

$$q_{BB}(x_t|x_{t-1}, y) = N\left(x_t; \frac{1 - m_t}{1 - m_{t-1}}x_{t-1} + \left(m_t - \frac{1 - m_t}{1 - m_{t-1}}m_{t-1}\right)y, \delta_{t|t-1}I\right) \quad (4)$$

where với $\delta_{t|t-1} = \delta_t - \frac{\delta_{t-1}(1-m_t)^2}{(1-m_{t-1})^2}$ . This expression accounts for the decrease in uncertainty due to conditioning on the previous state and reflects the structure of the Brownian Bridge, where variance shrinks as the trajectory is anchored at both endpoints. This conditional variance plays a key role in defining the forward transition kernel and informs the design of the reverse denoising process during training and inference.

*3.3.2. The reverse process*

The reverse process is designed to gradually reconstruct the target image from a noisy latent representation $x_t = y$, starting from the source image. It begins by setting the initial latent variable at the final timestep equal to the source image, which serves as the starting condition. From this point, the model iteratively moves backward through the diffusion steps, progressively denoising the image. At each timestep, the model assumes that the latent variable follows a conditional Gaussian distribution $p_\theta(x_{t-1}|x_t, y) = N(x_{t-1}; \mu_\theta(x_t, t), \tilde{\delta}_t I$, where the mean $\mu_\theta(x_t, t)$ is predicted by a neural network that takes the current latent state and the timestep as input. This predicted mean reflects the model's estimate of the clean image content at that step. The variance $\tilde{\delta}_t$ at each step is predefined and controls the level of stochasticity during the denoising process.

*3.3.3. Training objective*

The training process is carried out by minimizing the Evidence Lower Bound (ELBO), which provides a variational approximation to the true data likelihood. The ELBO is formulated as:

$$ELBO = -E_q(D_{KL}(q_{BB}(x_T|x_0,y)||p(x_T|y)) + \sum_{t=2}^{T} D_{KL}(q_{BB}(x_{t-1}|x_t,x_0,y)||p_\theta(x_{t-1}|x_t,y)) - \log p_{\theta(x_0|x_1,y)}$$

The training objective minimizes the Evidence Lower Bound (ELBO), balancing alignment of latent representations with a prior, consistency between forward and reverse transitions, and accurate reconstruction of target images. Through this optimization, the model learns a probabilistic reverse process that transforms non-attenuation-corrected images into attenuation-corrected counterparts, guided by both the structural prior and paired supervision.

*3.3.4. Sampling process*

In this study, the fast sampling strategy of DDIM is employed to accelerate inference by reducing diffusion steps while preserving consistency with the original distribution. Following the CPDM design [12] is adopted in the context of image translation from non-attenuation-corrected images $I_{NC}$ to attenuation-corrected images $I_{AC}$.

## 4. RESULTS AND DISCUSSION

### 4.1. Dataset

The dataset was constructed by the research team at the Department of Nuclear Medicine, 108 Military Central Hospital, and comprises 203 myocardial perfusion SPECT cases. The data were collected from routine clinical practice, reviewed by multiple board-certified nuclear medicine physicians, and served as the basis for treatment decision-making. The study was conducted with formal approval from the Department of Nuclear Medicine, ensuring compliance with ethical standards and patient data confidentiality regulations.

Each case in the dataset includes two phases (rest and stress) and two types of images: non-attenuation-corrected SPECT (NC) and CT-based attenuation-corrected SPECT (CT-AC). The data are oriented along three standard cardiac planes: short axis (SA), vertical long axis (VLA), and horizontal long axis (HLA), each comprising 29 slices to capture detailed myocardial perfusion distribution and left ventricular motion with clinically consistent resolution. This design supports both quantitative and qualitative analysis, enabling comparison between NC and AC images, the development of deep learning–based attenuation correction without CT, and evaluation of diagnostic performance under real-world conditions. Figures 1 and 2 illustrate representative reconstructions: CT-AC images (figure 1) allow high-resolution assessment of perfusion defects and ventricular function, while NC images (figure 2) reflect raw perfusion data affected by attenuation artifacts, underscoring both the value of correction and the interpretive challenges posed by uncorrected scans.

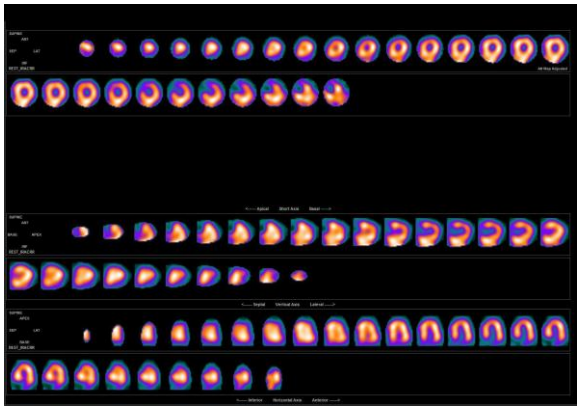

***Figure 1.*** *CT-AC SPECT images visualized in three standard cardiac planes.*

### 4.2. Evaluation criteria

In this study, three evaluation metrics: SSIM, PSNR, and RMSE are calculated for model evaluation. The Structural Similarity Index Measure (SSIM) quantifies the structural similarity between reference and reconstructed images, capturing perceptual fidelity in terms of human visual perception. Peak Signal-to-Noise Ratio (PSNR) evaluates the ratio between the maximum possible signal and the noise present in the reconstruction, serving as an indicator of sharpness and distortion level - higher PSNR values indicate closer similarity to the ground truth. Root Mean Square Error (RMSE) measures the average squared difference between corresponding pixel values, with lower RMSE values reflecting better reconstruction quality. All images used in the evaluation were normalized to the pixel range of [−1,1].

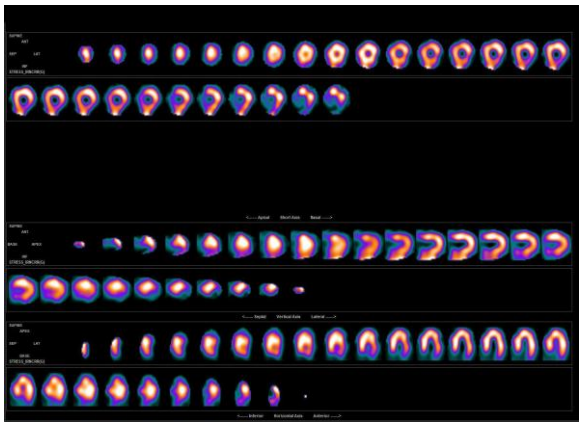

***Figure 2.*** *NC SPECT images displayed in three standard cardiac planes.*

### 4.3. Experimental setup

*4.3.1. Pix2Pix [7]*

Trained with a UNet-256 [14] architecture for the generator and a PatchGAN 70×70 discriminator, implemented using PyTorch 1.4. The batch size was fixed to 1 to reduce variance across mini-batches. Both networks were optimized using the Adam optimizer with a learning rate of 0.0002 and β1=0.5 for a total of 70 epochs; the learning rate was kept constant for the first 50 epochs and linearly decayed to 0 over the remaining 20 epochs.

*4.3.2. RegGAN [12]*

Utilizes a generator based on UNet-256 [14] (8 down–up blocks, no skip connections) combined with a PatchGAN 70×70 discriminator. The generator was regularized by a spatial correspondence constraint RRR, which helps enforce consistency between the generated and reference domains. The components were jointly optimized using the Adam optimizer with a *learning rate* = 0.0001, β1=0.5, β2=0.999 *weight-decay* = 0.0004. The learning rate was kept constant for the first 20 epochs of 80 total, and linearly decayed afterward.

*4.3.3. BBDM [10] and CPDM [15]*

VQGAN [13] was trained to encode 256×256 input images into a 32×32 latent space with 2 channels and embedding dimensions of 4. Multiple UNet denoisers were trained to predict the spatial latent features. The denoisers were trained using Adam ($lr = 1 \times 10^{-4}$, $\beta_1 = 0.9$, weight decay = 0) and ReduceLROnPlateau scheduler (factor 0.5, patience 3000 steps, min LR $5 \times 10^{-7}$). EMA (decay 0.995) was also applied to all parameters, starting after 30,000 training steps.

*4.3.4. UNIT [11]*

Trained with batch size 1 using Adam optimizer with learning rate $1 \times 10^{-4}$ ($\beta_1 = 0.5$, $\beta_2 = 0.999$) for 100,000 iterations, with a generator of 2 downsampling layers and 4 residual blocks,

and an LSGAN-based PatchGAN discriminator. This setup captured both global and local features while preserving high-frequency details for sharper outputs.

### 4.4. Results

***Table 1.*** *Quantitative results of attenuation-corrected image reconstruction.*

| **Model** | **SSIM ↑** | **PSNR ↑** | **RMSE ↓** |
|---|---|---|---|
| Pix2Pix [7] | 0.8170 | <u>29.69</u> | 0.1450 |
| RegGAN [12] | 0.8242 | 25.64 | <u>0.1278</u> |
| UNIT [11] | 0.8100 | **33.17** | 0.1300 |
| CPDM [15] | <u>0.8407</u> | 27.60 | 0.1285 |
| BBDM [10] | **0.8529** | 28.33 | **0.1177** |

*We use **bold** and underlined text to denote the best and second-best scores, respectively.*
*↑/↓ indicates higher/lower is better.*

Table 1 summarizes the performance of the five models in reconstructing attenuation-corrected images, evaluated using **SSIM, PSNR**, and **RMSE** with pixel values normalized to the range [-1, 1]. Among these, the diffusion-based models BBDM and CPDM outperform GAN-based models in terms of SSIM and RMSE. Specifically, BBDM achieves the highest SSIM score (0.8529), which is 3.5% higher than RegGAN, the best GAN-based model. It also attains the lowest RMSE (0.1177), outperforming RegGAN by 7.9%.

CPDM, another diffusion-based model, achieves a high SSIM score (0.8407), 2.9% higher than RegGAN, though with slightly lower PSNR, highlighting the strength of diffusion models in preserving spatial structures and fine details through iterative denoising. In contrast, GAN-based approaches such as Pix2Pix and RegGAN, while yielding higher PSNR (e.g., Pix2Pix at 29.69 dB), generally show lower SSIM and higher RMSE, reflecting reduced perceptual quality and structural consistency. Notably, UNIT obtains the highest PSNR (33.17 dB), but in SSIM (0.8100) and RMSE (0.1300), it even underperforms RegGAN, underscoring that high PSNR alone does not guarantee superior image quality. These results illustrate the common over-smoothing tendency of GANs, where pixel-wise accuracy comes at the cost of fine anatomical details, whereas diffusion models better balance global and local consistency, ultimately offering superior perceptual fidelity and stronger potential for clinical SPECT attenuation correction.

### 4.5. Discussion about potential improvement

Diffusion-based approaches outperform GANs in reconstruction quality and preservation of structural details, while their stochastic denoising process enables uncertainty quantification via confidence maps. When combined with interpretability techniques such as attention visualization or saliency attribution, these models enhance transparency, reveal anatomical features guiding correction, and strengthen clinical trust. To ensure robustness, however, broader validation is required across multiple scanners, acquisition protocols, noise levels, and radiotracer doses. Future work should integrate uncertainty estimation with interpretability into a unified framework, complementing strong reconstruction performance with clinically actionable insights to support safe and effective adoption of deep learning–based attenuation correction in CAD workflows.

## 5. CONCLUSIONS

This study introduced a comprehensive benchmark comparing GAN-based and diffusion-based generative models for reconstructing attenuation-corrected myocardial perfusion SPECT images from non-corrected inputs, with evaluations across SSIM, PSNR, and RMSE consistently demonstrating the superior performance of diffusion models in preserving anatomical fidelity and achieving higher perceptual and quantitative accuracy. These findings highlight the advantages of iterative denoising mechanisms in medical image reconstruction and provide a reproducible

foundation for future NC-to-AC methodologies. Moving forward, expanding the dataset to include more diverse patient cohorts and acquisition conditions will be essential to strengthen robustness and generalizability, while the exploration of more efficient diffusion-based architectures could address current limitations in inference time without sacrificing reconstruction quality. In parallel, incorporating uncertainty quantification and explainability frameworks will be critical to foster clinical trust, and prospective validation under real-world conditions will ultimately determine the feasibility of integrating NC-to-AC reconstruction directly into SPECT imaging workflows, thereby reducing or eliminating reliance on CT-based attenuation correction.

## REFERENCES


[1]. Liu, H., Wu, J., Shi, L., Liu, Y., Miller, E., Sinusas, A., ... & Liu, C. *"Post-reconstruction attenuation correction for SPECT myocardium perfusion imaging facilitated by deep learning-based attenuation map generation."* Journal of Nuclear Cardiology, 29(6), 2881–2892, (2022).

[2]. Chen, X., Zhou, B., Xie, H., Miao, T., Liu, H., Holler, W., ... & Liu, C. *"DuDoSS: Deep-learning-based dual-domain sinogram synthesis from sparsely sampled projections of cardiac SPECT."* Medical Physics, 50(1), 89–103, (2023).

[3]. Liu, J., Yang, Y., Wernick, M. N., Pretorius, P. H., & King, M. A. *"Deep learning with noise-to-noise training for denoising in SPECT myocardial perfusion imaging."* Medical Physics, 48(1), 156–168, (2021).

[4]. Ho, J., Jain, A., & Abbeel, P. *"Denoising diffusion probabilistic models."* Advances in Neural Information Processing Systems, 33, 6840–6851, (2020).

[5]. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. *"High-resolution image synthesis with latent diffusion models."* Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10684–10695, (2022).

[6]. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., & Norouzi, M. *"Palette: Image-to-image diffusion models."* Proceedings of the ACM SIGGRAPH Conference, 1–10, (2022).

[7]. Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Tao, A., Kautz, J., & Catanzaro, B. *"High-resolution image synthesis and semantic manipulation with conditional GANs."* (2018).

[8]. Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. *"Unpaired image-to-image translation using cycle-consistent adversarial networks."* Proceedings of the IEEE International Conference on Computer Vision, 2223–2232, (2017).

[9]. Liu, M.-Y., Breuel, T., & Kautz, J. *"Unsupervised image-to-image translation networks."* Advances in Neural Information Processing Systems, 30, (2017).

[10]. Li, B., Xue, K., Liu, B., & Lai, Y.-K. *"BBDM: Image-to-image translation with Brownian bridge diffusion models."* Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1952–1961, (2023).

[11]. Armanious, K., Jiang, C., Fischer, M., Küstner, T., Hepp, T., Nikolaou, K., Gatidis, S., & Yang, B. *"MedGAN: Medical image translation using GANs."* Computerized Medical Imaging and Graphics, 79, 101684, (2020).

[12]. Kong, L., Lian, C., Huang, D., Li, Z., Hu, Y., & Zhou, Q. *"Breaking the dilemma of medical image-to-image translation."* arXiv preprint arXiv:2110.06465, (2021).

[13]. Esser, P., Rombach, R., & Ommer, B. *"Taming transformers for high-resolution image synthesis."* Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 12873–12883, (2021).

[14]. Ronneberger, O., Fischer, P., & Brox, T. *"U-Net: Convolutional networks for biomedical image segmentation."* Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, 234–241, (2015).

[15]. Nguyen, D. T., Nguyen, T. T., Nguyen, H. T., Nguyen, T. T., Pham, H. H., Nguyen, T. H., ... & Le Nguyen, P. *"CT to PET Translation: A Large-scale Dataset and Domain-Knowledge-Guided Diffusion Approach."* Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 1498–1507, (2025).

## TÓM TẮT

### Tái tạo ảnh SPECT tim chuẩn hóa suy giảm bằng mô hình sinh học sâu: Hướng tiếp cận hiện đại và đánh giá thực nghiệm

*Chụp xạ hình tưới máu cơ tim (myocardial perfusion imaging - MPI) bằng máy chụp cắt lớp phát xạ đơn photon (single-photon emission computed tomography - SPECT) là công cụ quan trọng trong chẩn đoán bệnh động mạch vành nhưng thường bị ảnh hưởng bởi suy giảm tín hiệu do hấp thụ mô mềm. Trong nghiên cứu này, chúng tôi sử dụng bộ dữ liệu SPECT MPI công bố rộng rãi để xây dựng bộ benchmark đánh giá bài toán tái tạo ảnh hiệu chỉnh suy giảm (attenuation correction - AC) từ ảnh gốc chưa hiệu chỉnh (non-attenuation correction - NC) theo hướng 2D lát cắt (slice-to-slice). Chúng tôi triển khai và so sánh hiệu quả nhiều mô hình sinh ảnh tiên tiến, bao gồm mạng sinh đối nghịch (generative adversarial network - GAN) và mô hình khuếch tán (diffusion model). Các mô hình sinh được thiết kế để huấn luyện trên miền dữ liệu tổng quát (general domain) lẫn miền y học (medical domain). Kết quả cho thấy các phương pháp học sâu hiện đại có thể tái tạo ảnh AC với chất lượng cao, mở ra tiềm năng ứng dụng thực tiễn cho hệ thống hỗ trợ chẩn đoán bằng máy tính (computer-aided diagnosis - CAD) trên dữ liệu SPECT MPI.*


**Từ khoá:** Chụp cắt lớp phát xạ đơn photon (SPECT); Xạ hình tưới máu cơ tim (MPI); Hỗ trợ chẩn đoán bằng máy tính (CAD); Mạng sinh đối nghịch (GAN); Mô hình khuếch tán (Diffusion model).