

CT to PET Translation: A Large-scale Dataset and Domain-Knowledge-Guided Diffusion Approach

Dac Thai Nguyen¹, Trung Thanh Nguyen², Huu Tien Nguyen¹, Thanh Trung Nguyen³,
 Huy Hieu Pham⁴, Thanh Hung Nguyen¹, Thao Nguyen Truong⁵, and Phi Le Nguyen¹

¹Hanoi University of Science and Technology, Vietnam; ²Nagoya Univeristy, Japan

³108 Military Central Hospital, Vietnam; ⁴VinUniversity, Vietnam

⁵National Institute of Advanced Industrial Science and Technology, Japan

¹{thai.nd180172@sis, tien.nh205033@sis, hungnt@soict, lenp@soict}.hust.edu.vn

²nguyent@cs.is.i.nagoya-u.ac.jp, ³thanhrungys@yahoo.com

⁴hieu.ph@vinuni.edu.vn, ⁵nguyen.truong@aist.go.jp

Abstract

Positron Emission Tomography (PET) and Computed Tomography (CT) are essential for diagnosing, staging, and monitoring various diseases, particularly cancer. Despite their importance, the use of PET/CT systems is limited by the necessity for radioactive materials, the scarcity of PET scanners, and the high cost associated with PET imaging. In contrast, CT scanners are more widely available and significantly less expensive. In response to these challenges, our study addresses the issue of generating PET images from CT images, aiming to reduce both the medical examination cost and the associated health risks for patients. Our contributions are twofold: First, we introduce a conditional diffusion model named **CPDM**, which, to our knowledge, is one of the initial attempts to employ a diffusion model for translating from CT to PET images. Second, we provide the largest CT-PET dataset to date, comprising 2,028,628 paired CT-PET images, which facilitates the training and evaluation of CT-to-PET translation models. For the CPDM model, we incorporate domain knowledge to develop two conditional maps: the Attention map and the Attenuation map. The former helps the diffusion process focus on areas of interest, while the latter improves PET data correction and ensures accurate diagnostic information. Experimental evaluations across various benchmarks demonstrate that CPDM surpasses existing methods in generating high-quality PET images in terms of multiple metrics. The source code and data samples are available at <https://github.com/thanhff/CPDM>.

1. Introduction

The combination of PET and CT has markedly enhanced oncology imaging by integrating PET's functional imaging




Figure 1. Examples of synthetic PET images produced by natural image generative models (BBDM [18], DDIB [28], Palette [26], and Pix2Pix [35]). The generated images do not accurately replicate the ground truth.

capabilities with the anatomical detail of CT scans. This integration boosts accuracy in diagnosis, staging, and evaluation of cancer treatments [33, 36]. However, the application of PET in medical examination and treatment faces several key challenges: (1) PET imaging requires the injection of a small amount of radioactive material, which may pose health risks [4]; (2) PET scanners are limited in availability, particularly in underdeveloped regions [2, 6]; and (3) the high cost of PET imaging imposes a financial burden on patients [4]. In contrast, CT imaging mitigates many of these issues, primarily because of its lower cost, typically around half that of PET imaging, thereby reducing the financial burden on patients. Therefore, leveraging artificial intelligence to generate PET images from CT scans offers a viable solution to address the aforementioned challenges. This approach also holds substantial practical significance and high applicability for patient care.

The challenge of converting CT to PET images falls within the domain of Image-to-Image (I2I) translation, a prominent area of research in computer vision. Numerous methods have been proposed to address this task, with generative models such as Generative Adversarial Network (GAN) [12] emerging as the leading approach. Based on GAN, Pix2Pix [14] made significant strides in I2I translation tasks with its model consisting of a generator and a

discriminator. Since then, numerous GAN variants, such as CycleGAN [38], UVCGAN [30], and SelectionGAN [29], have been introduced. Despite their ability to generate synthetic images that closely resemble real ones, GAN-based models are often hindered by training instability and sensitivity to hyperparameters. Recently, diffusion model [7] has emerged as a promising alternative, noted for their stability during training and ability to generate highly realistic images [8]. Diffusion models have been applied to address the I2I translation problem in several studies [18, 21, 26]. However, as shown in Figure 1, their research primarily focuses on natural images, leading to suboptimal performance when applied to medical images, particularly in the CT-to-PET translation task. Several studies have concentrated on addressing the I2I translation problem for medical images, such as those in [3, 16, 20, 39]. However, current research encounters certain challenges, including: (1) A considerable amount of current I2I translation research in medical imaging is dedicated to transforming between two akin domains, such as MRI T1 to MRI T2 [16, 34, 37], CBCT to CT [20]. There is a noticeable shortage of research tackling translation between entirely distinct modalities, especially, from CT images to PET; (2) The dominant approach based on generative models entails generating a random image that conforms to the distribution of the target domain. While these models utilize the source image to guide the generation process, thereby enhancing the resemblance of the generated image to the ground truth, they fail to mitigate the inherent stochasticity of the generated images; (3) While the natural I2I translation benefits from large datasets [18], the domain of medical I2I, and CT-to-PET in particular, remains significantly data-constrained. This scarcity of data has hindered advancement in these fields.

To tackle the aforementioned challenges, this study focuses on a significant yet underexplored issue: CT-to-PET translation. Alongside providing an extensive dataset, we aim to develop a CT-to-PET translation model with the following objectives:

- **Enhancing the deterministic nature of the generated PET image:** Rather than relying on conventional generative models that aim to generate random images, we employ a diffusion model directly mapping from the source domain to the target domain. In essence, this diffusion model learns a process to directly convert a CT image into a corresponding PET image.
- **Improving the accuracy of the generated PET image:** We integrate domain-specific knowledge to guide the diffusion process. Specifically, we utilize two conditional maps, namely the *Attention map* and the *Attenuation map*, to steer the diffusion process. The former essentially highlights pixels that are likely to be within the PET region of interest. Meanwhile, the latter, vital

for accurate PET image quantification, computes photon attenuation at 511 keV from the CT scan, improving PET data correction and securing precise diagnostic information.

The main contributions of this paper include:

- We propose **CPDM (CT-to-PET Diffusion Model)**, a novel approach that leverages a Brownian Bridge process-based diffusion model to directly synthesize PET images from CT scans. To refine the quality of the generated PET images, we guide the diffusion process with the domain knowledge from our additional conditional maps, i.e., Attention and Attenuation maps.
- We offer a comprehensive large-scale PET/CT medical dataset comprising 2,028,628 paired CT-PET images. To our knowledge, this represents the largest CT-PET dataset available. This dataset is curated to support developing and evaluating advanced medical image translation methods, providing a valuable resource for the research community.
- We conduct comprehensive experiments to evaluate the performance of CPDM against state-of-the-art (SOTA) image translation methods. The results demonstrate that CPDM outperforms SOTA methods across multiple metrics for PET image quality.

2. Related Work

In this section, we review existing CT and PET datasets in Section 2.1, followed by a discussion of recent research on I2I translation in Section 2.2.

2.1. CT and PET Datasets

The advancement of cancer imaging research has been significantly enhanced by the availability of large-scale 3D PET/CT datasets. The RIDER Lung PET-CT [23] (2015) and Lung-PET-CT-Dx [19] (2020) datasets, both focus on lung cancer, provide PET and CT scans that have been instrumental in developing multimodal imaging techniques and machine learning models for lung cancer diagnosis. Expanding the anatomical scope, the Head-Neck-PET-CT [32] (2017) dataset offers specialized imaging data for head and neck cancers, while the FDG-PET-CT-Lesions [11] (2022) dataset includes comprehensive whole-body PET/CT scans with manual annotations performed using specialized software. These datasets are pivotal in advancing research on lesion detection and comprehensive cancer diagnostics across different body regions.

Despite the significant contributions of these datasets, supervised learning tasks such as I2I translation require precisely paired PET and CT images on a slice-by-slice basis, a criterion not typically met by existing datasets. Fur-

Table 1. Overview of our dataset and existing CT and PET datasets (sorted by their published year). **Auto-Paired** means that corresponding CT and PET slices are automatically paired slice-by-slice, a feature missing in existing datasets, which only provide unpaired image series in DICOM format.

Name of Dataset	Auto Paired	Body Parts	Total Slices	# of Studies
RIDER Lung PET-CT [23]	✗	Lung	266K	274
Head-Neck-PET-CT [32]	✗	Head & Neck	123K	504
Lung-PET-CT-Dx [19]	✗	Lung	251K	430
FDG-PET-CT-Lesions [11]	✗	Whole body	917K	1,014
Our Dataset	✓	Whole body	2M	3,454

thermore, generative tasks, including image generation and translation, generally require extensive datasets to develop high-quality models, a need that is not adequately fulfilled by current CT-PET datasets (refer to Table 1). In this context, our study contributes a large-scale paired CT-PET dataset, which is intended not only to advance the CT2PET translation problem but also to support research involving CT, PET, and medical imaging in general.

2.2. Image-to-Image Translation

Generative Models for Natural Images. Generative models have been crucial in advancing I2I translation tasks, especially for natural images. Pix2Pix [35] pioneers conditional GANs to translate images from one domain to another with controlled output generation. While this model demonstrates significant progress, it struggles to produce diverse outputs due to its one-to-one mapping approach. Subsequent models, CycleGAN [38] and DRIT++ [17] seek to address these limitations by enabling unpaired image translation and generating diverse samples. However, these models face challenges in training stability and mode collapse, restricting their performance in complex tasks. More recently, diffusion-based models have emerged as a promising alternative for image synthesis, offering increased stability and high-quality outputs. Palette [26] and SDEdit [21] demonstrate superior performance across various I2I tasks without requiring task-specific tuning, making them flexible solutions across domains. Subsequently, LDM [24] improves efficiency by conducting the diffusion process in the latent space of pre-trained models. Building on these advancements, BBDM [18] introduces a novel Brownian Bridge diffusion process for direct image domain translation, enhancing stability. However, despite these improvements, most of these models remain focused on natural images, with limited application to medical imaging.

Medical Image Translation Models. In the medical domain, several GAN-based models have been developed for translating medical images, including PET-CT, MRI, and CT-to-CT tasks. MedGAN [3] is one such framework that

Table 2. Comparison of the proposed CPDM with existing I2I translation methods. **S**, **R** and **F** stand for the **Sharpness**, **Realism**, and **Faithfulness** of the visual performance, estimated in three levels: High (H), Medium (M) and Low (L).

	Method	External Knowledge	Medical I2I	Visual Per.		
				S	R	F
GAN-based	Pix2Pix [35]			M	M	L
	MedGAN [3]		✓	M	M	M
	UP-GAN [38]		✓	M	M	M
	FCN-cGAN [5]		✓	M	M	M
Diffusion-based	Palette [26]			H	H	L
	DDIB [28]			H	H	L
	SDEdit [21]			H	H	L
	LDM [24]			H	H	L
	BBDM [18]			H	H	L
	CPDM (Ours)	✓	✓	H	H	H

applies GANs for medical image translation, particularly in PET-CT translation, demonstrating the utility of GANs in clinical applications. Similarly, UP-GAN [38] adopts an uncertainty-guided progressive learning approach for translating between medical imaging modalities. Although these models have made significant strides in medical image translation, they are often limited by their reliance on paired data or struggle with generating high-fidelity images useful for diagnostic purposes. In contrast, diffusion-based models have shown promise for overcoming the limitations of GANs in medical image translation. SynDiff [40] utilizes a conditional diffusion process to generate high-quality medical images. However, its reliance on cycle-consistent architecture for training still poses inherent limitations, particularly in generating diagnostically relevant outputs.

To address current challenges, we propose CPDM, a novel diffusion-based method designed for CT2PET translation tasks. Unlike existing methods, CPDM leverages a Brownian Bridge process to directly synthesize PET images from CT images by incorporating domain-specific knowledge to guide the diffusion process, enhancing the quality and clinical relevance of the results. We compare CPDM with existing I2I translation methods in Table 2.

3. Our Proposed PET/CT Dataset

Full Dataset. We present a comprehensive large-scale PET/CT dataset collected from hospitals (detailed statistics are provided in Table 3). The dataset consists of 2,028,628 paired CT-PET images from studies of 3,454 patients, designed to cover a wide range of anatomical regions. To ensure compliance with privacy and ethical guidelines, we remove all pathological labels rigorously. Each study includes approximately 250-500 paired CT and PET slices from the top of the head to the upper thigh region, above

Table 3. Detailed statistics of our large-scale CT-PET dataset.

General Information		CT Information		PET Information	
Total CT-PET slices	2,028,628	Resolution (per image)	512 x 512 x 1	Resolution (per image)	256 x 256 x 1
Total studies	3,454	Tube voltage (kVp)	120 / 140	Radioactive tracer	^{18}F -FDG
Total pairs per study	250-500	Slice thickness (mm)	3.75 / 5	Slice thickness (mm)	3.27
Body parts	Whole body	Slope coefficient	1.0	Attenuation correction	CT-based
PET/CT system	GE Discovery 710 / STE	Intercept coefficient	-1024.0	Uptake time (minutes)	60

the knees. The images are stored in DICOM format, encapsulating pixel data and relevant metadata, including patient age (in years), sex, body weight, injected radiotracer activity, and other acquisition parameters. Importantly, the PET images have undergone attenuation correction using the corresponding CT data. The images were acquired using the GE Discovery 710 PET/CT and GE Discovery STE PET/CT systems, with acquisition parameters defined by factors such as Kilovoltage peak (kVp), slope coefficient, and intercept coefficient. The dataset contains both diseased and non-diseased cases, providing a diverse representation of clinical scenarios.

Experimental Dataset. Due to computational constraints, we limit our experiments (in Section 5) to a subset of the complete dataset. Specifically, we derive a sub-dataset of 30,000 CT-PET images (15,000 pairs), corresponding to 598 studies, by randomly selecting 25-30 paired slices per study. The experimental dataset is organized with sequential indexing at the patient level, i.e., images with consecutive indices correspond to contiguous anatomical sections from the same individual. The dataset is subsequently stratified into training, validation, and test sets using an 80:10:10 split. That is, we use 80% of the studies allocated to the training set and the remaining 20% of data for the validation and test sets. Consequently, the test set contains entirely unseen data during model training. All image pairs have been standardized to a resolution of $256 \times 256 \times 1$ and normalized to a range of $[-1; 1]$, based on the maximum pixel intensities of $2^{11} - 1$ for CT images and $2^{15} - 1$ for PET images.

4. CPDM: Domain-Knowledge-Guided CT-to-PET Diffusion Model

4.1. Motivation and Architecture Overview

CT and PET images exhibit distinct characteristics compared to natural images. In these medical images, the majority of pixels are black, with only a small portion of pixels containing the critical non-black pixels that convey the most essential information. Specifically, in PET images, the intensely bright areas (High-SUV regions) are particularly significant, as they often signal regions with a high likelihood of abnormalities. Translating CT images into PET

images, thus, introduces unique challenges not encountered in natural image translation. First, because the bright regions are the most crucial, the translation process must prioritize these areas rather than distributing focus across the entire image. Second, the generated image must be sufficiently sharp to emphasize the critical bright regions effectively. Traditional generative models struggle to meet these specific requirements, as they tend to produce images with high variability, focusing more on fitting a general distribution than on achieving precise detail.

To address these challenges, we propose to use condition maps to guide the diffusion process, facilitating the generation of images that align with the inherent characteristics of PET images. For the first challenge, we introduce an attention map highlighting the regions likely to appear bright in the PET image. This attention map is derived from the CT image using a deep neural network trained through supervised learning. To tackle the second challenge, we based on the fact that different body parts absorb radiation at varying levels. We, thus, employ an attenuation map, a matrix that represents the absorption levels of radiation (light intensity) across different body parts. The attenuation map can be accurately calculated from the input CT image. Indeed, the attenuation map is traditionally used to correct PET images, making it an ideal tool for enhancing image clarity and focus. Additionally, inspired by [18], we integrate the Brownian Bridge process into our diffusion model for CT2PET translation. In contrast to existing diffusion methods, the proposed approach directly establishes a mapping between the input CT and output PET domains. This direct mapping eliminates the stochasticity typically associated with generative models.

Figure 2 illustrates the architecture of the proposed method, named CPDM. Starting with a CT image from the CT domain, we first extract its latent feature L_{CT} . Following this, the Brownian Bridge process maps L_{CT} to its corresponding latent representation $L_{CT \rightarrow PET}$ in the PET domain. During this process, domain knowledge in the form of Attention map (M_{σ}) and Attenuation map (M_{μ}) are employed. Subsequently, the synthesized PET image is generated by the decoder. We leverage a pre-trained VQGAN [10] for the Encoder/Decoder on our PET/CT dataset to enhance learning efficiency and model generalization.

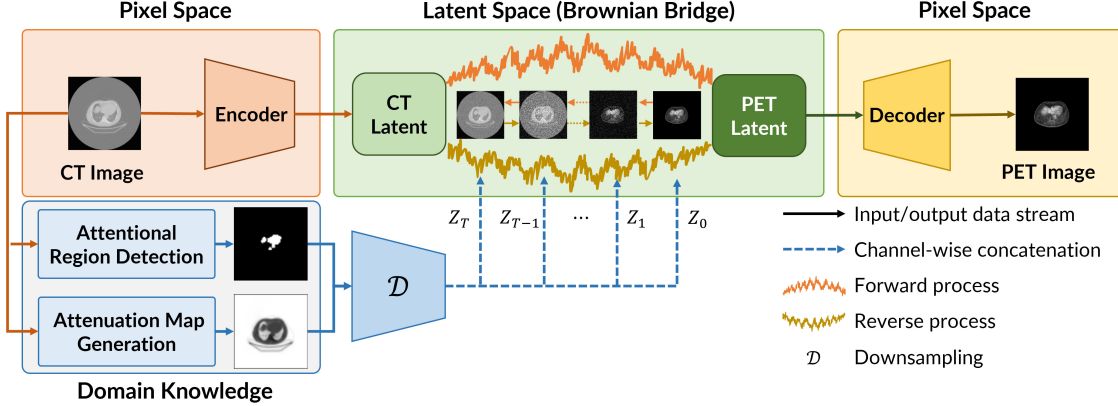


Figure 2. **Overview of CPDM.** The medical knowledge Attention map and the Attenuation map are used to guide each stage of the Brownian Bridge diffusion process.

4.2. Domain Knowledge-Guided Maps

Attention Map (M_σ). The Attention map identifies clinical regions of interest in CT images that show increased ^{18}F -Fluorodeoxyglucose (^{18}F -FDG) and exhibit higher SUVs uptake in PET images. To generate this map, we employ the U-Net [25] segmentation model, which is trained on CT images to distinguish these high-activity tissues, using Dice Loss [22] function to measure accuracy. PET images are employed to define the ground truth for the U-Net model by creating segmented masks that highlight areas of interest based on their brightness. In detail, we perform a qualitative assessment of PET images across different anatomical regions and utilize a straightforward thresholding approach to sample segmented masks, isolating pixels with intensity values that exceed a predefined threshold. The loss function is defined as follows:

$$\mathcal{L}_{M_\sigma} = 1 - \frac{2|M_\sigma \cap M_{gt}|}{|M_\sigma| + |M_{gt}|},$$

where M_{gt} represents the ground truth mask. The resulting Attention map M_σ then guides the Brownian Bridge diffusion process to focus on the highlighted regions when generating the PET images.

Attenuation Map (M_μ). In practical settings, PET/CT scanners typically utilize CT images to correct attenuation, thereby improving the quality of PET emission data. This crucial process involves converting CT data, acquired across a broad spectrum of photon energies (approximately 30 to 140 kVp), into attenuation coefficients for PET photon energies at 511 keV. Inspired by [1], we employ the tri-linear scaling methodology to calculate the Attenuation map. The Linear Attenuation Coefficient (LAC) at 511 keV, denoted as μ , is calculated using the equation $\mu = \alpha \times HU + \beta$, where α and β are constants derived from scanner calibration, which are listed in [1], and HU represents the Hounsfield Units calculated from CT data (i.e.

slope coefficient, intercept coefficient, and X-ray tube potential). The generated Attenuation map $M_\mu = e^{-\mu}$ is subsequently incorporated into the Brownian Bridge diffusion process, providing a guidance informed by medical expertise for model refinement.

4.3. Conditional Brownian Bridge Diffusion Process

Drawing inspiration from [18], we adopt the Brownian Bridge process to characterize the probability distribution throughout the diffusion. This process aims to facilitate a direct transformation from the source to the target domain. To enhance the training process’s efficiency, we use an encoder and decoder to transform PET and CT images into latent space vectors before the diffusion process. Denoting the original PET and CT images as \mathbf{X} and \mathbf{Y} , and the encoder and decoder as \mathcal{E} and \mathcal{D} , the corresponding latent vectors, \mathbf{x} and \mathbf{y} , are obtained as $\mathbf{x} := \mathcal{E}(\mathbf{X})$ and $\mathbf{y} := \mathcal{E}(\mathbf{Y})$. During the training phase, CPDM learns a mapping from \mathbf{y} to \mathbf{x} in the latent space through a Brownian Bridge process. At the end of the inference phase, the translated PET image $\tilde{\mathbf{X}}$ is generated by the decoder of the pre-trained VQGAN as $\tilde{\mathbf{X}} := \mathcal{D}(\tilde{\mathbf{x}})$, where $\tilde{\mathbf{x}}$ is the final latent state of the reverse diffusion process.

Forward Process. We utilize the formulation proposed in [18]. Starting from an initial state \mathbf{x}_0 (i.e., the latent vector of a PET image) and aiming for a destination state \mathbf{y} (i.e., the latent vector of the corresponding CT image), the intermediate state \mathbf{x}_t at timestep t are determined in discrete form as $\mathbf{x}_t = (1 - m_t)\mathbf{x}_0 + m_t\mathbf{y} + \sqrt{\delta_t}\epsilon_t$, where $m_t = t/T$, T represents the total number of steps in the diffusion process, δ_t indicates the variance of the Brownian Bridge, and ϵ depicts a Gaussian noise, i.e., $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The forward process is defined as:

$$q_{BB}(\mathbf{x}_t | \mathbf{x}_0, \mathbf{y}) = \mathcal{N}(\mathbf{x}_t; (1 - m_t)\mathbf{x}_0 + m_t\mathbf{y}, \delta_t \mathbf{I}). \quad (1)$$

Algorithm 1 CPDM Training Process

- 1: **repeat**
 - 2: Paired data: PET image $\mathbf{x}_0 \sim q(\mathbf{x}_0)$, CT image $\mathbf{y} \sim q(\mathbf{y})$; Attention map \mathbf{M}_σ ; Attenuation map \mathbf{M}_μ
 - 3: Timestep $t \sim \text{Uniform}(1, \dots, T)$
 - 4: Gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: Forward diffusion $\mathbf{x}_t = (1 - m_t) \mathbf{x}_0 + m_t \mathbf{y} + \sqrt{\delta_t} \epsilon$
 - 6: Take gradient descent:
 $\nabla_\theta \|m_t (\mathbf{y} - \mathbf{x}_0) + \sqrt{\delta_t} \epsilon - \epsilon_\theta (\text{concat}(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu), t)\|_1$
 - 7: **until** converged
-

To describe the transition probability between consecutive steps, we use the following formula:

$$q_{BB}(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{y}) = \mathcal{N} \left(\mathbf{x}_t; \frac{1 - m_t}{1 - m_{t-1}} \mathbf{x}_{t-1} + \left(m_t - \frac{1 - m_t}{1 - m_{t-1}} m_{t-1} \right) \mathbf{y}, \delta_{t|t-1} \mathbf{I} \right),$$

with $\delta_{t|t-1} = \delta_t - \delta_{t-1} \frac{(1 - m_t)^2}{(1 - m_{t-1})^2}$. (2)

Reverse Process. The reverse initiates by conditionally assigning $\mathbf{x}_T = \mathbf{y}$. Furthermore, to incorporate the domain knowledge, we employ the channel-wise concatenation to integrate the Attention map \mathbf{M}_σ and the Attenuation map \mathbf{M}_μ within the hidden state \mathbf{x}_t at every timestep t as:

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}, t), \tilde{\delta}_t \mathbf{I}),$$
$$\mu_\theta(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}, t) = c_{x_t} \mathbf{x}_t + c_{y_t} \mathbf{y} + c_{\epsilon_t} \epsilon_\theta(\text{concat}(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu), t), \quad (3)$$

where $\mu_\theta(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}, t)$ represents the predicted mean and $\tilde{\delta}_t$ represents the variance of the distribution at time step t , respectively; c_{x_t} , c_{y_t} and c_{ϵ_t} are non-trainable factors derived from m_t , m_{t-1} , δ_t and δ_{t-1} :

$$c_{x_t} = \frac{\delta_{t-1}}{\delta_t} \frac{1 - m_t}{1 - m_{t-1}} + \frac{\delta_{t|t-1}}{\delta_t} (1 - m_{t-1}),$$
$$c_{y_t} = m_{t-1} - m_t \frac{1 - m_t}{1 - m_{t-1}} \frac{\delta_{t-1}}{\delta_t}, \quad c_{\epsilon_t} = (1 - m_{t-1}) \frac{\delta_{t|t-1}}{\delta_t}.$$

Training Objective. The training process entails optimizing the disparity between the predicted distribution of the latent variables and the observed distribution in the forward diffusion process (see Algorithm 1). This is achieved by learning the mean value $\mu_\theta(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}, t)$ through a neural network parameterized by θ , employing maximum likelihood estimation. In practice, this is accomplished by minimizing the following Evidence Lower Bound (ELBO):

$$\text{ELBO} = -\mathbb{E}_q(D_{KL}(q_{BB}(\mathbf{x}_T | \mathbf{x}_0, \mathbf{y}) \| p(\mathbf{x}_T | \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y})))$$
$$+ \sum_{t=2}^T D_{KL}(q_{BB}(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0, \mathbf{y}) \| p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}))$$
$$- \log p_\theta(\mathbf{x}_0 | \mathbf{x}_1, \mathbf{M}_\sigma, \mathbf{M}_\mu, \mathbf{y}).$$

Algorithm 2 CPDM Sampling Process

- 1: Sample conditional input: CT image $\mathbf{x}_T = \mathbf{y} \sim q(\mathbf{y})$;
Attention map \mathbf{M}_σ ; Attenuation map \mathbf{M}_μ
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 - 4: $\mathbf{x}_{t-1} = c_{x_t} \mathbf{x}_t + c_{y_t} \mathbf{y} + c_{\epsilon_t} \epsilon_\theta(\text{concat}(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu), t) + \sqrt{\delta_t} \mathbf{z}$
 - 5: **end for**
 - 6: **return** \mathbf{x}_0
-

By substituting the values of $q_{BB}(\cdot)$ and $p_\theta(\cdot)$ from Equations 2 and 3, respectively, we derive the final training ELBO objective as follows:

$$\mathbb{E}_{\mathbf{x}_0, \mathbf{y}, \epsilon} \left[c_{\epsilon_t} \left\| m_t (\mathbf{y} - \mathbf{x}_0) + \sqrt{\delta_t} \epsilon - \epsilon_\theta(\text{concat}(\mathbf{x}_t, \mathbf{M}_\sigma, \mathbf{M}_\mu), t) \right\|_1 \right].$$

Sampling Process. We adopt the basic idea of DDIM (Denosing Diffusion Implicit Models) [27], where acceleration is achieved through a non-Markovian process while maintaining identical marginal distributions as Markovian inference processes (see Algorithm 2).

5. Performance Evaluation

In this section, we evaluate the performance of CPDM using our PET/CT dataset. We compare CPDM against SOTA image translation methods, including GAN-based models, i.e., Pix2Pix [14], MedGAN [3], and UP-GAN [31] and diffusion-based models, i.e., DDIM [28], Palette [26], SDEdit [21], LDM [24], and BBDM [18] (Section 5.2). We also provide the ablation studies in Section 5.3.

5.1. Experimental Settings

Evaluation Metrics. Following [3, 31], we utilize Learned Perceptual Image Patch Similarity (LPIPS), Mean Absolute Error (MAE), Structural Similarity Index Measure (SSIM), and Peak Signal-to-Noise Ratio (PSNR) metrics to evaluate the quality of generated PET images. These metrics ensure that the generated PET images are accurate and of high visual quality. LPIPS evaluates perceptual similarity, which is crucial for capturing medical details. MAE measures the average pixel-wise error, indicating accuracy. SSIM reflects visual quality by assessing changes in the image structure, whereas PSNR quantifies pixel-level fidelity.

Models & Hyperparameters. We implement the CPDM as detailed in Section 4. The diffusion timesteps of the Brownian Bridge are set at 1000 for training and reduced to 200 for inference, aiming to balance quality and speed. The Unet [25], utilizing a pre-trained ResNet50 [13] as its backbone, is employed for extracting the attention map. The hyperparameters of the baselines remain consistent with those specified in the original papers. We adjust the input and output image channels to 1 and a batch size to 16 to optimize hardware processing capabilities. We utilize Adam [15] op-

Table 4. **Comparison of CPDM against SOTA methods.** The best and second-based results are highlighted by **red** and **blue**, respectively. **Diff.** shows the relative performance gaps of CPDM compared to the nearest methods. \downarrow/\uparrow means lower/higher is better.

Method	LPIPS \downarrow	MAE \downarrow	SSIM \uparrow	PSNR \uparrow
Pix2Pix [14]	0.0491	336.47	0.9298	27.52
MedGAN [3]	0.1322	336.97	0.9329	29.27
UP-GAN [31]	0.1082	373.49	0.9337	28.38
DDIB [28]	0.5556	4898.58	0.0343	16.10
Palette [26]	0.3827	4032.23	0.3054	20.92
SDEdit [21]	0.1577	813.65	0.8187	22.85
LDM [24]	0.0492	318.71	0.9317	28.30
BBDM [18]	0.0548	353.66	0.9232	27.56
CPDM (w/o M_σ)	0.0471	305.46	0.9353	28.76
CPDM (w/o M_μ)	0.0523	324.99	0.9277	28.47
CPDM (Ours)	0.0466	284.61	0.9396	29.68
Diff.	-5.1%	-10.7%	+0.6%	+1.4%

optimizer with an initial learning rate of 10^{-4} , with the step decay for learning rate scheduling and a batch size of 16 across 200 epochs.

5.2. Comparison with Existing Methods

Table 4 presents a thorough comparison of PET image generation models, where CPDM stands out by achieving the best results across all metrics: lowest LPIPS (0.0466), lowest MAE (284.61), highest SSIM (0.9396), and highest PSNR (29.68). It demonstrates that CPDM provides superior perceptual quality and structural fidelity. Other methods perform well on specific metrics but often underperform on others. For instance, MedGAN achieves the second-best PSNR (29.27) but ranks fourth-worst in LPIPS (0.1322), highlighting its trade-off between pixel-level noise reduction and perceptual quality. Similarly, UP-GAN performs well in SSIM but falls behind in LPIPS. Overall, CPDM offers a more balanced and robust performance across all metrics. The final row (Diff.) shows CPDM’s improvement over the second-best models, with a 5.1% reduction in LPIPS, a 10.7% decrease in MAE, a 0.6% increase in SSIM, and a 1.4% rise in PSNR.

To further investigate the differences between CPDM and other methods, we visualize several examples in Figure 3. PET images generated by CPDM exhibit sharper details than those produced by the other models. As error maps show, most of CPDM’s pixels appear dark blue, indicating minimal reconstruction error. In contrast, GAN-based models, such as MedGAN and UP-GAN, demonstrate relatively good faithfulness to the input data but often suffer from insufficient sharpness and lack of realism in the generated images. Conversely, diffusion-based models like LDM, which rely on conditional synthetic generation, produce images with high sharpness and detailed structures;

however, they exhibit lower faithfulness, struggling to preserve anatomical integrity accurately. BBDM achieves a more balanced performance, offering a reasonable trade-off between sharpness, faithfulness, and realism. However, it lacks the domain-specific knowledge necessary to match the anatomical accuracy and structural faithfulness that CPDM consistently delivers. It highlights CPDM’s ability to maintain a superior balance across crucial image quality attributes, outperforming other methods.

5.3. Ablation Studies

5.3.1 Impacts of Domain Knowledge Guidance

We implement two variants of CPDM, namely CPDM (w/o M_σ) and CPDM (w/o M_μ), to examine the impacts of the two proposed conditional maps.

Quantitative Results. The results in Table 4 show that both variants of CPDM perform worse than the original CPDM across all metrics, underscoring the effectiveness of utilizing both conditional maps. Specifically, omitting the Attention map M_σ leads to an increase of approximately 1.1% and 7.3% in LPIPS and MAE, along with decreases of 0.46% and 3.0% in SSIM and PSNR, respectively. In contrast, removing the Attenuation map M_μ results in significant increases of 12.2% and 14.2% in LPIPS and MAE, accompanied by decreases of 1.3% and 4.1% in SSIM and PSNR, respectively. The results demonstrate that the Attenuation map M_μ is more critical than the Attention map M_σ , as its absence causes more substantial degradation across all metrics.

Qualitative Results. To evaluate the effects of the Attention map M_σ and the Attenuation map M_μ on features generated by the diffusion process, we use Principal Component Analysis (PCA) [9] to visualize the intermediate spatial features throughout the denoising network ϵ_θ . Figure 4 shows that the BBDM model spreads features into the surrounding background rather than focusing on patient anatomy. In contrast, the proposed CPDM effectively integrates both two condition maps, achieving focused attention on regions of interest and meaningful anatomical structures. The CPDM without M_σ (second row) shows an enhanced focus on anatomical features during the diffusion process, particularly at intermediate timesteps ($t \approx T/2$), where it maintains consistency even in high noise conditions. Meanwhile, CPDM without M_μ (third row) demonstrates that the Attention map helps ϵ_θ focus on significant features related to regions of interest, maintaining this focus even at intermediate timesteps with minimal information.

5.3.2 Robustness of the Condition Maps

In this section, we evaluate the robustness of the proposed condition maps, M_σ , and M_μ , by integrating them




Figure 3. Visualization of PET images generated by the best-performing methods. Errors induced by PET produced by CPDM are significantly more minor than other methods.




Figure 4. Visualization of spatial features extracted from the denoising network ϵ_θ .

Table 5. Impact of the our conditional maps on existing methods.

Method		LPIPS ↓	MAE ↓	SSIM ↑	PSNR ↑
Pix2Pix	Original	0.0491	336.47	0.9298	27.52
	w/ (M_σ, M_μ)	0.0470	305.33	0.9345	28.49
	Diff.	-4.3%	-9.3%	+0.50%	+3.4%
MedGAN	Original	0.1322	336.97	0.9329	29.27
	w/ (M_σ, M_μ)	0.1158	294.05	0.9325	29.73
	Diff.	-12.4%	-12.7%	-0.04%	+1.5%
UP-GAN	Original	0.1082	373.49	0.9337	28.38
	w/ (M_σ, M_μ)	0.1001	301.67	0.9296	29.99
	Diff.	-7.5%	-19.2%	-0.44%	+5.7%
LDM	Original	0.0492	318.71	0.9317	28.30
	w/ (M_σ, M_μ)	0.0468	302.70	0.9348	28.92
	Diff.	-4.9%	-8.2%	+0.33%	+2.2%
CPDM (Ours)		0.0466	284.61	0.9396	29.68

into various I2I models (i.e., Pix2Pix, MedGAN, UP-GAN, and LDM) and analyzing their impact on performance enhancement. Specifically, we developed a conditional variant for each model that incorporates the proposed maps to guide the generative process, embedding domain-specific knowledge to refine model outputs. Table 5 presents quantitative results demonstrating that integrating our conditional maps improves the performance of the original methods. The inclusion of M_σ and M_μ leads to significant enhancements in LPIPS and MAE, with UP-GAN achieving gains of 7.5% and 19.2%, respectively. The slight reduction in MAE can be attributed to the conditional maps, which provide additional guidance during generation and effectively reduce pixel-wise discrepancies by embedding domain-specific information. Despite these improvements, UP-GAN and MedGAN still exhibit considerably higher LPIPS than CPDM, indicating lower perceptual quality.

The conditional maps encourage these models to focus on faithfulness, resulting in smoother images with fewer pixel-level discrepancies but diminished sharpness and detail. This trade-off likely explains the elevated PSNR and slight decline in SSIM observed in UP-GAN and MedGAN, as the generated images prioritize smoothness at the expense of visual fidelity. In contrast, LDM, leveraging the gradual refinement process typical of diffusion models, shows more balanced improvements across all metrics by focusing on local and global consistency. Nonetheless, CPDM surpasses LDM by achieving a superior balance between sharpness, realism, and faithfulness, preserving finer details while delivering better overall performance.

6. Conclusion

This study proposed CPDM, a novel method for synthesizing PET images from CT images using a diffusion process. CPDM utilized a Brownian Bridge diffusion process to translate data directly from the CT domain to the PET domain, mitigating the inherent stochastic characteristics often present in generative models. Furthermore, CPDM integrated Attention and Attenuation maps, resulting in PET images with significantly improved quality and reduced error compared to existing methods. The quantitative and qualitative results demonstrated the superiority of CPDM over existing GAN-based and diffusion-based methods across various evaluation criteria.

Acknowledgment

This work was funded by Vingroup Joint Stock Company (Vingroup JSC), Vingroup, and supported by Vingroup Innovation Foundation (VINIF) under project code VINIF.2021.DA00128.

References

- [1] Monica Abella, Adam M Alessio, David A Mankoff, Lawrence R MacDonald, Juan Jose Vaquero, Manuel Desco, and Paul E Kinahan. Accuracy of CT-based attenuation correction in PET/CT bone imaging. *Physics in Medicine & Biology*, 57(9):2477, 2012. [5](#)
- [2] Suhaib Alameen, Nissren Tamam, Sami Awadain, Abdelmoneim Sulieman, Latifa Alkhalidi, and Amira Ben Hmed. Radiobiological risks in terms of effective dose and organ dose from 18f-fdg whole-body pet/ct procedures. *Saudi Journal of Biological Sciences*, 28:5947–5951, 2021. [1](#)
- [3] Karim Armanious, Chenming Jiang, Marc Fischer, Thomas Küstner, Tobias Hepp, Konstantin Nikolaou, Sergios Gatidis, and Bin Yang. Medgan: Medical image translation using GANs. *Computerized Medical Imaging and Graphics*, 79:101684, 2020. [2](#), [3](#), [6](#), [7](#)
- [4] Avi Ben-Cohen, Eyal Klang, Stephen P. Raskin, Michal Marianne Amitai, and Hayit Greenspan. *Virtual PET Images from CT Data Using Deep Convolutional Networks: Initial Results*, page 49–57. Springer International Publishing, 2017. [1](#)
- [5] Avi Ben-Cohen, Eyal Klang, Stephen P. Raskin, Shelly Soffer, Simona Ben-Haim, Eli Konen, Michal Marianne Amitai, and Hayit Greenspan. Cross-modality synthesis from ct to pet using fcn and gan networks for improved automated lesion detection, 2018. [3](#)
- [6] Simon R Cherry, Terry Jones, Joel S Karp, Jinyi Qi, William W Moses, and Ramsey D Badawi. Total-body PET: maximizing sensitivity to create new opportunities for clinical research and patient care. *Journal of Nuclear Medicine*, 59(1):3–12, 2018. [1](#)
- [7] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [2](#)
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [2](#)
- [9] George H Dunteman. *Principal components analysis*, volume 69. Sage, 1989. [7](#)
- [10] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. [4](#)
- [11] S. Gatidis and T. Kuestner. A whole-body fdg-pet/ct dataset with manually annotated tumor lesions (fdg-pet-ct-lesions), 2022. Dataset. [2](#), [3](#)
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014. [1](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [6](#)
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017. [1](#), [6](#), [7](#)
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *Computing Research Repository arXiv Preprints arXiv:1412.6980*, 2014. [6](#)
- [16] Lingke Kong, Chenyu Lian, Detian Huang, Yanle Hu, Qichao Zhou, et al. Breaking the dilemma of medical image-to-image translation. *Advances in Neural Information Processing Systems*, 34:1964–1978, 2021. [2](#)
- [17] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. Dri++: Diverse image-to-image translation via disentangled representations, 2019. [3](#)
- [18] Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. BBDM: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1952–1961, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [19] P. Li, S. Wang, T. Li, J. Lu, Y. HuangFu, and D. Wang. A large-scale ct and pet/ct dataset for lung cancer diagnosis (lung-pet-ct-dx), 2020. Data set. [2](#), [3](#)
- [20] Zhiwen Liang, Hui Wei, Gang Liu, Mengjie Cheng, Jiaqi Gao, Song Li, and Xin Tian. Leveraging GAN-based CBCT-to-CT translation models for enhanced image quality and accurate photon and proton dose calculation in adaptive radiotherapy. *Journal of Radiation Research and Applied Sciences*, 17(1):100809, 2024. [2](#)
- [21] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. [2](#), [3](#), [6](#), [7](#)
- [22] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Proceedings of the 4th International Conference on 3D Vision*, pages 565–571, 2016. [5](#)
- [23] P. Muzi, M. Wanner, and P. Kinahan. Data from rider lung pet-ct, 2015. Dataset. [2](#), [3](#)
- [24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [3](#), [6](#), [7](#)
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the 18th International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241, 2015. [5](#), [6](#)
- [26] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *Proceedings of the 2022 ACM Special Interest Group on Computer Graphics and Interactive Techniques Conference*, pages 1–10, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)

- [27] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *Computing Research Repository arXiv Preprints arXiv:2010.02502*, 2020. [6](#)
- [28] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. *Computing Research Repository arXiv Preprints arXiv:2203.08382*, 2022. [1](#), [3](#), [6](#), [7](#)
- [29] Hao Tang, Dan Xu, Nicu Sebe, Yanzhi Wang, Jason J Corso, and Yan Yan. Multi-channel attention selection GAN with cascaded semantic guidance for cross-view image translation. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2417–2426, 2019. [2](#)
- [30] Dmitrii Torbunov, Yi Huang, Haiwang Yu, Jin Huang, Shinjae Yoo, Meifeng Lin, Brett Viren, and Yihui Ren. Uvcgan: Unet vision transformer cycle-consistent GAN for unpaired image-to-image translation. In *Proceedings of the 2023 IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 702–712, 2023. [2](#)
- [31] Uddeshya Upadhyay, Yanbei Chen, Tobias Hepp, Sergios Gatidis, and Zeynep Akata. Uncertainty-guided progressive GANs for medical image translation. In *Proceedings of the 24th International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 614–624, 2021. [6](#), [7](#)
- [32] Martin Vallières, Emily Kay-Rivest, Léo Jean Perrin, Xavier Liem, Christophe Furstoss, Nader Khaouam, Phuc Félix Nguyen-Tan, Chang-Shu Wang, and Khalil Sultanem. Data from head-neck-pet-ct, 2017. Dataset. [2](#), [3](#)
- [33] Srinivasan Vijayakumar, Johnny Yang, Mary R Nittala, Alexander E Velazquez, Brandon L Huddleston, Nickhil A Rugnath, Neha Adari, Abhay K Yajurvedi, Abhinav Komanduri, Claus Chunli Yang, N Duggar William, P Berlin William, Duszak Richard, and Vijayakumar Vani. Changing role of PET/CT in cancer care with a focus on radiotherapy. *Cureus*, 14(12), 2022. [1](#)
- [34] Clinton J Wang, Natalia S Rost, and Polina Golland. Spatial-intensity transforms for medical image-to-image translation. *IEEE Transactions on Medical Imaging*, 2023. [2](#)
- [35] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans, 2018. [1](#), [3](#)
- [36] Lise Wei and Issam El Naqa. Artificial intelligence for response evaluation with PET/CT. In *Seminars in Nuclear Medicine*, volume 51, pages 157–169, 2021. [1](#)
- [37] Qianye Yang, Nannan Li, Zixu Zhao, Xingyu Fan, Eric I-Chao Chang, and Yan Xu. Mri cross-modality image-to-image translation. *Scientific reports*, 10(1):3753, 2020. [2](#)
- [38] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the 2017 IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. [2](#), [3](#)
- [39] Yingying Zhu, Youbao Tang, Yuxing Tang, Daniel C Elton, Sungwon Lee, Perry J Pickhardt, and Ronald M Summers. Cross-domain medical image translation by shared latent gaussian mixture model. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23*, pages 379–389. Springer, 2020. [2](#)
- [40] Muzaffer Özbey, Onat Dalmaz, Salman UH Dar, Hasan A Bedel, Şaban Öztürk, Alper Güngör, and Tolga Çukur. Un-supervised medical image translation with adversarial diffusion models, 2023. [3](#)