3D Wavelet Latent Diffusion Model for Whole-Body MR-to-CT Modality Translation

Jiaxu Zheng†, Meiman He†, Xuhui Tang, Xiong Wang, Tuoyu Cao, Tianyi Zeng, Lichi Zhang, and Chenyu You

arXiv:2507.11557v1 [eess.IV] 14 Jul 2025

***Abstract*—Magnetic Resonance (MR) imaging plays an essential role in contemporary clinical diagnostics. It is increasingly integrated into advanced therapeutic workflows, such as hybrid Positron Emission Tomography/Magnetic Resonance (PET/MR) imaging and MR-only radiation therapy. These integrated approaches are critically dependent on accurate estimation of radiation attenuation, which is typically facilitated by synthesizing Computed Tomography (CT) images from MR scans to generate attenuation maps. However, existing MR-to-CT synthesis methods for whole-body imaging often suffer from poor spatial alignment between the generated CT and input MR images, and insufficient image quality for reliable use in downstream clinical tasks. In this paper, we present a novel 3D Wavelet Latent Diffusion Model (3D-WLDM) that addresses these limitations by performing modality translation in a learned latent space. By incorporating a Wavelet Residual Module into the encoder-decoder architecture, we enhance the capture and reconstruction of fine-scale features across image and latent spaces. To preserve anatomical integrity during the diffusion process, we disentangle structural and modality-specific characteristics and anchor the structural component to prevent warping. We also introduce a Dual Skip Connection Attention mechanism within the diffusion model, enabling the generation of high-resolution CT images with improved representation of bony structures and soft-tissue contrast. Quantitative assessments demonstrate that our method, 3D-WLDM, achieves superior results, with PSNR improvements of up to 3.98 dB (1.04 dB over the best baseline), SSIM improvements of up to 0.36 (0.02 over the best baseline), and an MAE reduction of up to 53.76 (7.76 lower than the best baseline). Qualitative evaluations and clinical utility assessments using an open-source organ segmentation model further reveal substantial gains in segmentation accuracy, highlighting the translational potential of our method for radiation planning, hybrid imaging, and broader biomedical applications requiring high-fidelity MR-to-CT synthesis. The source code for this work is publicly available on GitHub at https://github.com/hmm823/3D-Wavelet-Latent-Diffusion-Model.**

***Index Terms*—MR-to-CT synthesis, cross-modality translation, medical image generation, diffusion model, feature decouple.**

† Equally contributed to this work.

Corresponding authors: Tianyi Zeng and Lichi Zhang.

Jiaxu Zheng, Lichi Zhang are with the School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: jiaxu.zheng@sjtu.edu.cn; lichizhang@sjtu.edu.cn).

Meiman He, Xiong Wang are with the School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China (e-mail: hemm2022@shanghaitech.edu.cn; wangxiong@shanghaitech.edu.cn).

Xuhui Tang is with the School of Biomedical Engineering, ShanghaiTech University, Shanghai 201210, China (e-mail: tangxh2023@shanghaitech.edu.cn).

Tuoyu Cao is with the Shanghai United Imaging Healthcare Co., Ltd., Shanghai 201807, China (e-mail: tuoyu.cao@united-imaging.com).

Tianyi Zeng is with the Department of Radiology and Biomedical Imaging, Yale University, New Haven, CT, United States of America (e-mail: tianyi.zeng@yale.edu).

Chenyu You is with the Department of Applied Mathematics & Statistics and Department of Computer Science, Stony Brook University, Stony Brook, NY, United States of America (e-mail: chenyu.you@stonybrook.edu).

I. Introduction

MAGNETIC Resonance Imaging (MRI) provides superior soft tissue contrast compared to Computed Tomography (CT) and does not involve ionizing radiation, making it a preferred modality for tumor characterization, early disease detection, and strategic treatment planning [1], [2]. Leveraging these strengths, MRI has been increasingly integrated into advanced clinical applications, including PET/MR hybrid imaging and MR-only radiation therapy. However, both applications require accurate attenuation correction to compensate for radiation absorption in tissue, where CT remains the clinical gold standard due to its ability to directly provide electron density information. In contrast, MRI lacks such data, thereby presenting a significant limitation for PET attenuation correction and dose calculation in MR-only radiotherapy workflows [3], [4]. To address this gap, deep learning-based methods for MRI-to-CT synthesis have gained substantial attention [5], [6]. These approaches aim to generate synthetic CT images from MR data, enabling the estimation of attenuation maps suitable for PET reconstruction and radiotherapy planning. Additionally, the resulting synthetic CTs facilitate cross-modal registration by transforming heterogeneous alignment tasks into intra-modality problems, thereby improving the accuracy of longitudinal monitoring and lesion tracking [7].

Despite these advances, most existing MR-to-CT synthesis efforts focus on specific body regions such as the head or pelvis. Whole-body MR-to-CT synthesis remains underexplored, despite its clinical relevance in scenarios like total-body PET/MR imaging. Two major challenges have hindered progress in this domain: The first challenge lies in achieving accurate organ and tissue contrast in synthesized CT images. The nonlinear mapping between MRI and CT, rooted in their fundamentally different imaging physics, makes it difficult to achieve the precise reconstruction of fine-grained anatomical structures, such as bones. Effective translation of modality-specific features frequently requires integration of both local and global contextual information. However, traditional convolutional neural networks (CNNs), which rely on stacked local filters and pooling operations, are inherently limited in receptive field and contextual representation. Prior work has shown that larger spatial context improves synthesis

0000–0000/00$00.00 © 2021 IEEE

performance [8], and recent methods have sought to capture this advantage through frequency-domain enhancements and attention mechanisms [9]–[12]. However, these approaches are predominantly constrained to organ-specific or region-focused settings, and their scalability to full-body synthesis has yet to be demonstrated.

The second challenge concerns the anatomical misalignment between the MR and CT images, which often arises due to differences in patient posture, physiological motion, and variations in scanner bed geometry during separate acquisitions. These misalignments complicate the training of end-to-end modality translation models, which typically rely on the assumption of spatial correspondence. Complex anatomical structures and nonlinear distortions further exacerbate this issue. While CycleGAN-based methods incorporate additional constraints, such as gradients and local descriptors, to address anatomical misalignment [6], [13], they are typically limited to localized anatomical regions and often fail to maintain whole-body anatomical consistency. Furthermore, GAN-based models are vulnerable to mode collapse, which impairs their ability to generalize across diverse anatomical variations, resulting in limited synthesis fidelity and clinical applicability.

Diffusion-based generative models have recently emerged as a promising alternative for medical image synthesis, demonstrating superior performance over traditional GAN-based methods in terms of stability, fidelity, and diversity [10], [14]–[16]. The stochasticity of the forward diffusion process inherently reduces the risk of mode collapse, while latent space diffusion provides a natural mechanism for capturing long-range dependencies and contextual features [17]. These advantages motivate the application of diffusion models to the challenging task of whole-body MR-to-CT synthesis.

In this work, we propose a 3D Wavelet Latent Diffusion Model (3D-WLDM) for high-fidelity whole-body MR-to-CT synthesis. The proposed framework leverages the robust image generation capabilities of diffusion models to produce synthetic CT images that surpass existing GAN-based and diffusion-based methods in anatomical fidelity, contrast preservation, and structural consistency. To address the challenges inherent in this task, we introduce several key innovations. First, to enhance the accuracy of organ and tissue contrast in synthesized CT images, we integrate a *Wavelet Residual Module* into the encoder-decoder architecture. This module enhances the model's sensitivity to fine anatomical structures, thereby improving the quality of the encoded latent vectors and the reconstructed images derived from them. Second, to ensure structural consistency between input MR and synthetic CT images, we employ *Structure-Modality Disentanglement* in the latent space. This technique separates anatomical structures from modality-specific characteristics, preserving structural integrity throughout the iterative diffusion process. Additionally, we introduce a *Dual Skip Connection Attention* within the diffusion model to filter redundant information and suppress artifacts arising from MRI-specific textures. This mechanism selectively emphasizes structural features critical for accurate CT synthesis while minimizing interference from irrelevant modality details.

Our main contributions are summarized as follows:

- We propose **3D-WLDM**, a novel diffusion-based framework for whole-body MR-to-CT synthesis, which outperforms existing GAN-based and diffusion-based models in terms of anatomical fidelity and structural consistency.
- We develop *Wavelet Residual Module* to improve the quality of encoding and decoding between image and latent spaces, yielding CT images with more precise anatomical structure.
- We introduce *Structure-Modality Disentanglement* to anchor anatomical structure during latent space modalities translation, enhancing spatial consistency under cross-modal diffusion process.
- We incorporate *Dual Skip Connection Attention mechanism* to refines the fusion of structural and modality-specific features, resulting in greater synthesis stability and reduced artifact prevalence.

II. Related Works

A. MR-to-CT Modality Translation Methods

MR-to-CT image synthesis has garnered substantial research attention owing to its critical role in a range of clinical applications, including PET attenuation correction and radiotherapy planning. A variety of modality translation approaches have been proposed, with conditional Generative Adversarial Networks (cGANs) demonstrating particular success in synthesizing CT images from MR data across anatomical regions such as the pelvis, brain, and head [18]–[22]. However, MR-to-CT translation remains challenging in regions characterized by significant physiological deformation—most notably the abdomen—due to the pronounced spatial misalignment between MR and CT acquisitions [18], [23]. This issue becomes even more complex in whole-body synthesis tasks, where inter-organ motion and posture variability further exacerbate anatomical inconsistencies. To mitigate these challenges, CycleGAN-based frameworks have been widely adopted for unpaired cross-modality synthesis, often supplemented with additional spatial constraints such as gradient consistency, mutual information, or anatomical priors [6]. For whole-body applications specifically, certain studies have proposed the use of structural contour regularization to preserve global anatomical layout and reduce misalignment artifacts [13]. Despite these efforts, CycleGAN and its variants remain constrained by inherent limitations of adversarial training—particularly mode collapse, which can reduce synthesis diversity and generalizability. Additionally, the fundamental asymmetry in information content between MR and CT violates the cyclic consistency assumption of CycleGANs, further impairing their efficacy in capturing realistic CT features across large anatomical scales [24].

Cross-modality translation is inherently a high-dimensional, nonlinear, and ill-posed problem [25], which often exceeds the modeling capacity of conventional deep learning frameworks. To enhance synthesis performance, several innovative architectures have been introduced. Shi et al. [26] proposed a frequency decomposition module to separately process high- and low-frequency information, thereby improving anatomical detail preservation. Dovletov et al. [27], [28] introduced a

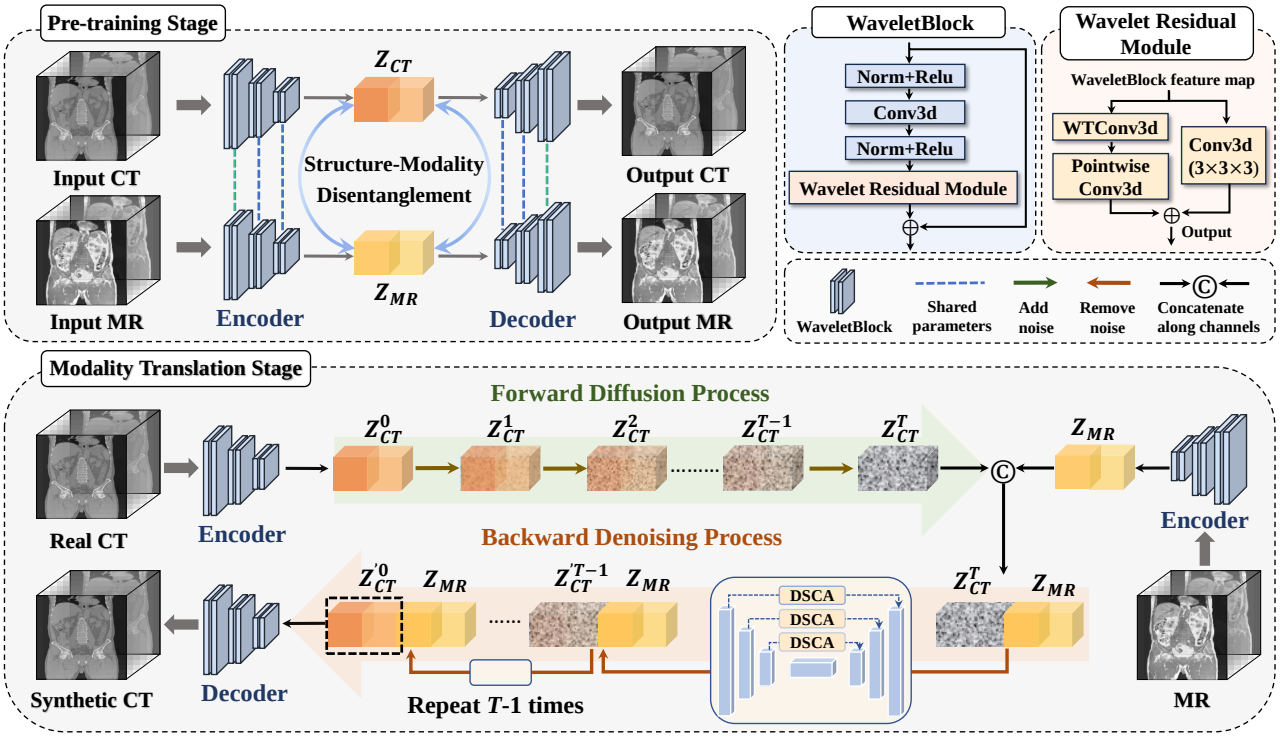

Fig. 1. Illustration of the proposed 3D-WLDM architecture. During the pre-training stage, we leverage Structure-Modality Disentanglement to align structural and modality information of the input CT and MR data in the latent space. Additionally, we enhance the response to frequency information by incorporating a Wavelet Residual Module into the WaveletBlocks. During the modality translation stage, the real CT is firstly encoded into the latent feature representation Z^0_{CT}, and Z^T_{CT} is subsequently derived through T timesteps of the forward diffusion process. In the backward denoising process, MR is encoded into the latent feature representation Z_{MR}, which are concatenated with Z^T_{CT} and fed into a U-net with the Dual Skip Connection Attention (DSCA) to predict the added noise. Through T time steps of iterative denoising, the CT feature representation Z'^0_{CT} is progressively recovered and decoded into synthetic CT.

Grad-CAM-guided framework that leverages pre-trained classifiers to focus the generator on diagnostically relevant regions. However, the reliance on convolutional neural networks (CNNs) in these methods limits their ability to capture long-range spatial dependencies and contextual relationships. To address this, hybrid CNN–Transformer models have been explored. For example, RTCGAN [29] integrates residual Transformers into the generator to extract hierarchical contextual features, while MTT-Net [30] incorporates image token-based Transformers within a 3D generator to model multi-scale spatial correlations. Despite their promising results, these GAN-based methods still face persistent challenges related to spatial misalignment, mode collapse, and limited contextual understanding—highlighting the need for alternative frameworks capable of more robust and anatomically consistent MR-to-CT synthesis.

B. Diffusion-Based Image Synthesis Methods

Diffusion models have recently demonstrated superior performance over GANs in image-to-image translation tasks by alleviating common issues such as mode collapse [31], [32]. For instance, StyleDiffusion [33] introduced a framework for content–style disentanglement that explicitly extracts content information while implicitly learning complementary style features. This allows for interpretable and controllable style transfer. To enable flexible conditional generation, ControlNet [34] proposed integrating lightweight adapter modules into pre-trained diffusion models, enabling efficient fine-tuning with external guidance. However, these methods have largely been applied to 2D image synthesis. When naively extended to 3D, they often suffer from inter-slice discontinuities due to the lack of spatial context across adjacent slices.

To overcome this limitation, Pan et al. proposed MC-DDPM [10], which integrates shifted-window transformers into a 3D diffusion framework to capture global anatomical structures for MR-to-CT translation. While effective, the model's sampling process is computationally intensive. In parallel, Friedrich et al. introduced WDM [35], a 3D wavelet-based diffusion model that applies denoising in the wavelet domain. By separating and treating frequency components independently, WDM enhances low-frequency global structures and suppresses high-frequency noise and artifacts. Tapp et al. [36] also explored diffusion-based MR-to-CT synthesis by embedding the denoising process within a 3D CNN architecture; however, the improvements in the synthesis fidelity remains to be modest.

On the other hand, training diffusion models in pixel space remains computationally demanding and leads to slow inference times. To mitigate these issues, Latent Diffusion Models (LDMs) [37] were introduced, operating in a lower-dimensional latent space to significantly reduce computational burden. Building on this, Cola-Diff [38] employed cross-attention mechanisms to incorporate diverse conditional signals for multimodal MRI synthesis within the LDMs frame-

work. Despite their efficiency, LDMs methods face challenges in aligning anatomical structures across modalities, particularly for MR-to-CT translation. Misalignment in the latent space and the presence of artifacts in the input MRI can significantly degrade the quality and reliability of the generated CT images.

III. Methods

The framework of the proposed 3D Wavelet Latent Diffusion Model (3D-WLDM) is illustrated in Fig. 1. The 3D-WLDM is designed to generate high-quality synthetic 3D CT images from MR images using a two-stage architecture that integrates a pre-training stage and a modality translation stage. In the pre-training stage, MR and CT images are encoded into latent space using shared-parameter encoders. To enhance the capture and reconstruction of fine-scale features, wavelet blocks with Wavelet Residual Module are incorporated into both the encoders and decoders. Within the latent space, a Structure-Modality Disentanglement mechanism separates structural information from modality-specific details across MR and CT, while aligning the structural information between the two modalities.

Subsequently, the modality translation stage employs a denoising diffusion process in the latent space to facilitate modality conversion from MR to CT, guided by a skip-connection architecture enhanced with Dual Skip Connection Attention (DSCA). This attention mechanism filters out discrepancies in contours and contrast between CT and MR that are propagated through skip connections, reducing interference during the modality conversion process. Finally, the transformed latent representation is decoded back into the image space via the decoder, yielding the synthesized CT image. Together, these components enable 3D-WLDM to effectively handle 3D medical imaging tasks with improved performance.

A. Wavelet Residual Module

Latent diffusion models typically employ Variational Autoencoders (VAEs) to map images into a compressed latent space, yet this process often compromises structural and contrast details, limiting efficacy in high-fidelity tasks like MR-to-CT modality conversion. To address this, we introduce the Wavelet Residual Module, which leverages wavelet decomposition to separately process high- and low-frequency feature components, thereby enhancing latent feature expressiveness and improving structural fidelity and modality conversion accuracy in the diffusion process.

As shown in Figs. 1 and 2, the Wavelet Residual Module is integrated into the WaveletBlock to enhance feature extraction by incorporating frequency-domain information. The module consists of three interconnected processing streams that manage information from the source domain, low-frequency domain, and high-frequency domain, respectively. The primary stream begins with the input feature map, which is processed through group-wise and point-wise convolutions and augmented by a residual connection to preserve hierarchical features. This stream serves as the backbone for integrating frequency-refined features from the other domains.

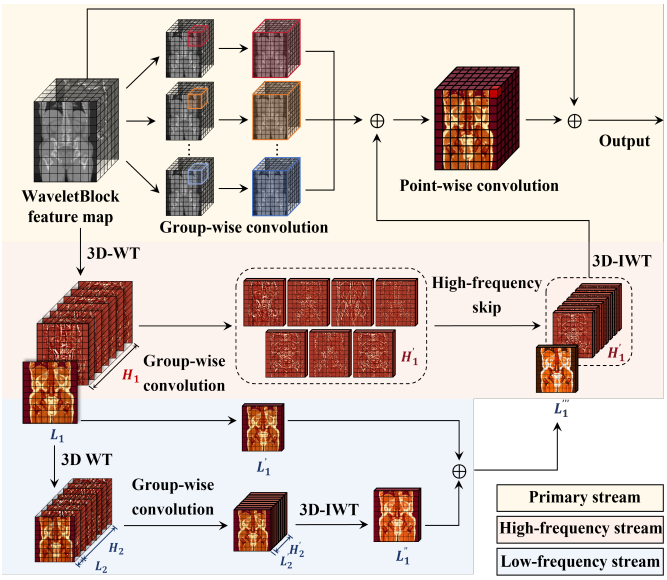

Fig. 2. Architecture of the Wavelet Residual Module in the WaveletBlock. The module processes the input feature map through three streams: a primary stream with group-wise and point-wise convolutions, a high-frequency stream generating H_1', and a low-frequency stream producing $L_1' + L_1''$. These two streams are fused via a 3D Inverse Wavelet Transform (3D-IWT) and integrated into the primary stream to enhance the extraction of structural and textural features.

In contrast to the primary stream, the second stream is dedicated to high-frequency processing, which starts with the initial 3D Wavelet Transform (3D-WT) that decomposes the input into a high-frequency component H_1 and a low-frequency component L_1. The H_1 component undergoes group-wise convolution to produce H_1', capturing detailed structural information, which is then held for subsequent fusion. The third stream, focused on low-frequency processing, further decomposes L_1 using a second 3D-WT into deeper wavelet frequency sub-bands L_2 (low-frequency) and H_2 (high-frequency). These sub-bands are processed through group-wise convolutions, and a 3D Inverse Wavelet Transform (3D-IWT) is applied to reconstruct L_1'', which is then fused with the low-frequency output L_1' to enhance low-frequency feature extraction.

Once the high-frequency stream (H_1') and the low-frequency stream ($L_1' + L_1''$) are obtained, they are subsequently combined and fed into another 3D-IWT to reconstruct the frequency-enhanced feature map. This reconstructed output is then integrated into the primary stream, which undergoes a point-wise convolution to produce the final output. This multi-stream design, leveraging wavelet-based decomposition, group-wise convolutions, and residual learning, increases the capability to encode and decode 3D medical images, contributing to the overall efficacy of the 3D-WLDM framework.

B. Latent Structure-Modality Disentanglement

Diffusion models rely on long Markov chains, where latent features risk losing structural information over extended iterations. To address this, we propose Structure-Modality Disentanglement, which decouples structural and modality-specific

information in latent vectors. This approach anchors structural features during diffusion, enhancing anatomical consistency.

The pipeline of Structure-Modality Disentanglement is shown in Fig. 3, which is proposed to enhance anatomical information in the latent space by decoupling structural information from modality-specific features within latent vectors Z. This decoupling process aims to improve anatomical consistency during the diffusion process, providing a foundation for subsequent information filtering.

The implementation of Latent Structure-Modality Disentanglement is based on the disentanglement loss function, $\mathcal{L}_D$, which comprises a structure loss $\mathcal{L}_{stru}$ and a modality loss $\mathcal{L}_{modal}$. As illustrated in Fig. 3, we pre-train encoders and decoders using paired MRI and CT volumes from the same patient. For input CT volumes $I_{CT} \in \mathbb{R}^{D\times H\times W}$ and MRI volumes $I_{MR} \in \mathbb{R}^{D\times H\times W}$, the encoder compresses these inputs along each dimension and outputs their respective distributions. From these distributions, we sample latent feature vectors $Z_{CT} \in \mathbb{R}^{c\times d\times h\times w}$ for CT and $Z_{MR} \in \mathbb{R}^{c\times d\times h\times w}$ for MRI.

Each latent vector Z (e.g., Z_{CT} or Z_{MR}) is explicitly divided into two parts: the structural component S and the modality component M, such that $Z = [S, M]$ along the channel dimension. Specifically, we designate the first half of the channels in Z_{CT} and Z_{MR} to encode their structural attributes as S_{CT} and S_{MR} respectively, while the second half encodes modality-specific attributes as M_{CT} and M_{MR}. This separation facilitates the disentanglement of anatomical structure (shared across modalities) from modality-specific characteristics (unique to CT or MRI).

To enforce this disentanglement, we define "paired" samples as Z_{CT} and Z_{MR} from the same patient, where the CT and MRI volumes are spatially aligned and thus share structural information. Conversely, "unpaired" samples from different patients, such as Z_{CT} and Z'_{CT}, share the same modality but differ in structure, facilitating the clustering of modality characteristics.

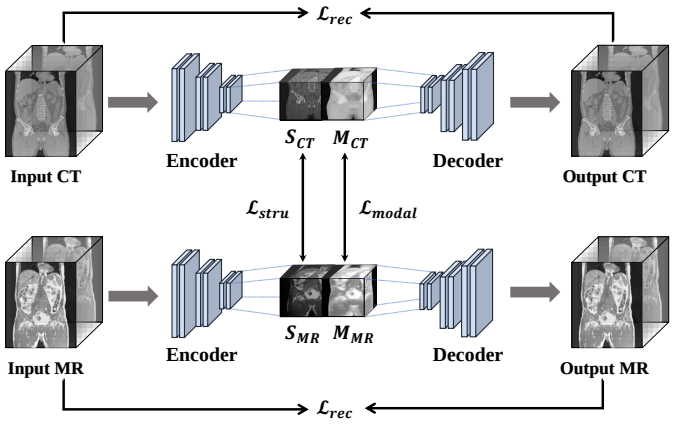

Fig. 3. Diagram of the Structure-Modality Disentanglement. The latent vectors Z_{CT} and Z_{MR} are split into the structural components S_{CT} and S_{MR} and the modality components M_{CT} and M_{MR}, respectively. The structure loss $\mathcal{L}_{stru}$ aligns S_{CT} and S_{MR} for structural similarity, while the modality loss $\mathcal{L}_{modal}$ enforces dissimilarity between M_{CT} and M_{MR}.

We impose constraints on S and M using the disentanglement loss $\mathcal{L}_D$. The structure loss $\mathcal{L}_{stru}$ applies a negative cosine similarity constraint within paired samples (e.g., S_{CT} and S_{MR}) to encourage structural alignment, and a positive cosine similarity constraint between unpaired samples (e.g., S_{CT} and S'_{CT}) to minimize structural similarity. Conversely, the modality loss $\mathcal{L}_{modal}$ enforces dissimilarity within paired samples (e.g., M_{CT} and M_{MR}) and similarity between unpaired samples (e.g., M_{CT} and M'_{CT}) to cluster modality-specific features. These losses are defined as:

$$\mathcal{L}_{stru} = -\cos\theta_s^{\text{pair}} + \cos\theta_s^{\text{unpair}}, \tag{1}$$

$$\mathcal{L}_{modal} = \cos\theta_m^{\text{pair}} - \cos\theta_m^{\text{unpair}}, \tag{2}$$

where the cosine similarities are:

$$\cos\theta_s^{\text{pair}} = \frac{S_{CT}\cdot S_{MR}}{\|S_{CT}\|\|S_{MR}\|} + \frac{S'_{CT}\cdot S'_{MR}}{\|S'_{CT}\|\|S'_{MR}\|}, \tag{3}$$

$$\cos\theta_s^{\text{unpair}} = \frac{S_{CT}\cdot S'_{CT}}{\|S_{CT}\|\|S'_{CT}\|} + \frac{S_{MR}\cdot S'_{MR}}{\|S_{MR}\|\|S'_{MR}\|}, \tag{4}$$

$$\cos\theta_m^{\text{pair}} = \frac{M_{CT}\cdot M_{MR}}{\|M_{CT}\|\|M_{MR}\|} + \frac{M'_{CT}\cdot M'_{MR}}{\|M'_{CT}\|\|M'_{MR}\|}, \tag{5}$$

$$\cos\theta_m^{\text{unpair}} = \frac{M_{CT}\cdot M'_{CT}}{\|M_{CT}\|\|M'_{CT}\|} + \frac{M_{MR}\cdot M'_{MR}}{\|M_{MR}\|\|M'_{MR}\|}. \tag{6}$$

The disentanglement loss $\mathcal{L}_D$ combines these terms as:

$$\mathcal{L}_D = \mathcal{L}_{stru} + \mathcal{L}_{modal}. \tag{7}$$

Here, S_{CT} and S_{MR} are structural components of paired CT and MRI data, and similarly for S'_{CT} and S'_{MR}. Likewise, M_{CT} and M_{MR} are modality components of paired data, with M'_{CT} and M'_{MR} following suit. For unpaired cases, S_{CT} and S'_{CT} represent structural components from different CT samples, with analogous definitions for MRI and modality components.

Finally, the encoders and decoders are pre-trained using the total loss function:

$$\mathcal{L}_{pre} = \mathcal{L}_{rec} + \alpha\cdot\mathcal{L}_{KL} + \beta\cdot\mathcal{L}_D + \gamma\cdot\mathcal{L}_{adv}, \tag{8}$$

where $\mathcal{L}_{rec}$ is the L2 reconstruction loss, $\mathcal{L}_{KL}$ aligns the latent distribution with a standard normal prior, and $\mathcal{L}_{adv}$ is an adversarial loss to enhance generated image quality. The hyperparameters α, β, and γ balance the contributions of each term.

C. Dual Skip Connection Attention

MRI offers superior soft tissue contrast compared to CT, but the asymmetric features between these modalities often introduce information redundancy in the CycleGAN frameworks, degrading synthesized CT image quality. To address this, we integrate the Dual Skip Connection Attention module (DSCA) into a skip-connection architecture to filter irrelevant information and enhance structural consistency, as shown in Fig. 4.

The DSCA modules within skip connections at different scales ($i = 1, 2, 3$), as depicted in the Fig. 4. Each DSCA module processes three inputs: the MR latent feature Z_{MR}, shallow features E_i from the i-th encoder layer, and deep features D_i from the i-th decoder layer. The DSCA employs a dual cross-attention mechanism, which comprises the Structure Emphasis

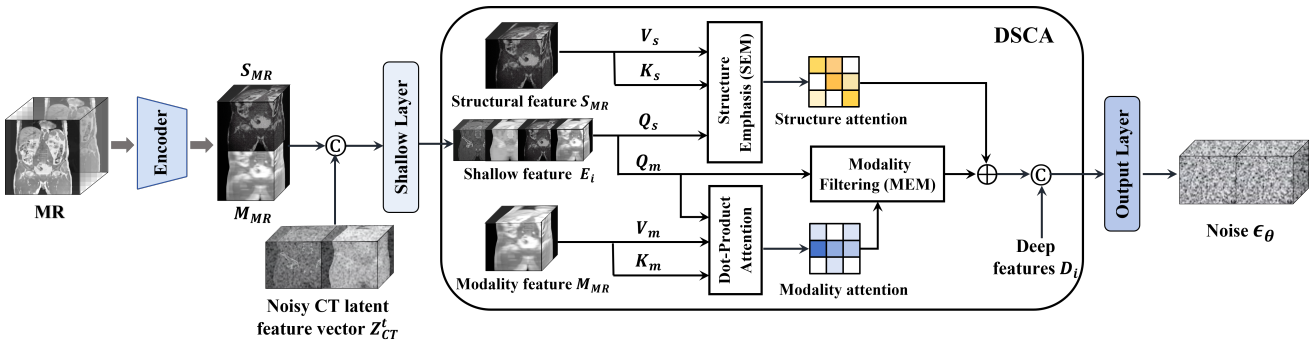

Fig. 4. Framework of Dual Skip Connection Attention (DSCA) modules in skip connections. Shallow feature map E_i is derived from the skip connection at the current resolution, while deep feature map D_i is upsampled from a lower-resolution level.

Module (SEM) and the Modality Filtering Module (MFM) to enhance anatomical accuracy and suppress modality-specific interference, as detailed in the right part of Fig. 4.

The SEM extracts structural information from S_{MR} to improve anatomical accuracy in synthesized CT images. Shallow features E_i are transformed into a query matrix Q_s, while S_{MR} generates key matrix K_s and value matrix V_s matrices. Structural attention A_{stru} is computed as:

$$A_{stru} = \text{Softmax}\left(\frac{\mathrm{Q}_s \times \mathrm{K}_s^T}{\sqrt{d_{K1}}}\right) \times \mathrm{V}_s, \tag{9}$$

where d_{K1} is the dimension of K_s. The output processed via captures consistent structural features.

The MFM excludes modality-specific interference from M_{MR}. Here, E_i forms the query Q_m, while M_{MR} produces key K_m and value V_m matrices. Modality attention A_{modal} is derived as:

$$A_{modal} = \mathrm{Q}_m - \text{Softmax}\left(\frac{\mathrm{Q}_m \times \mathrm{K}_m^T}{\sqrt{d_{K2}}}\right) \times \mathrm{V}_m, \tag{10}$$

where d_{K2} is the dimension of K_m. The output processed via filters out modality-specific noise.

The SEM and MFM outputs are combined with E_i via a residual connection, concatenated with D_i, and passed to the next decoder layer. The final DSCA output integrates structurally consistent and modality-filtered features for enhanced synthesis.

D. Condition-guided Training and Sampling

The 3D-WLDM framework employs a diffusion model to progressively transform the original CT latent feature Z_{CT}^0 into a noisy version, and subsequently recover it through a conditioned reverse process. This process is guided by the MR latent feature Z_{MR} to synthesize high-quality CT images.

In the forward process, Gaussian noise ϵ is incrementally added to Z_{CT}^0, generating the noisy CT latent feature Z_{CT}^t at time step t, expressed as:

$$Z_{CT}^t = \sqrt{\bar{\alpha}_t} \cdot Z_{CT}^0 + \sqrt{1-\bar{\alpha}_t} \cdot \epsilon, \tag{11}$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^{t} \alpha_i$. The variance $\beta_t \in (0, 1)$ serves as a hyperparameter controlling the noise level, typically scheduled over T time steps.

In the conditioned reverse process, the model learns to denoise Z_{CT}^t with guidance from Z_{MR}. Specifically, Z_{CT}^t and Z_{MR} are concatenated along the channel dimension and fed into a deeplearning network ϵ_θ, which is trained using the loss function:

$$\mathcal{L}_{WLDM} = \mathbb{E}_{\epsilon_t \sim \mathcal{N}(0,I)}\left[\left\|\epsilon_t - \epsilon_\theta\left(Z_{CT}^t, t, Z_{MR}\right)\right\|_2^2\right], \tag{12}$$

where ϵ_t is the true noise added at time step t, sampled from a standard normal distribution $\mathcal{N}(0, I)$, and $\epsilon_\theta\left(Z_{CT}^t, t, Z_{MR}\right)$ is the noise predicted by the network, conditioned on Z_{CT}^t, t, and Z_{MR}.

During the sampling phase, the trained network ϵ_θ is used to predict the noise at each step, enabling the iterative recovery of Z_{CT}^0. For inputs Z_{CT}^t and Z_{MR}, the latent feature vector Z_{CT}^{t-1} for the previous step is computed as:

$$Z_{CT}^{t-1} = \sqrt{\bar{\alpha}_{t-1}} \cdot Z'_{CT} + \sqrt{1-\bar{\alpha}_{t-1}} \cdot \epsilon_\theta\left(Z_{CT}^t, t, Z_{MR}\right), \tag{13}$$

where Z'_{CT} is the intermediate CT latent feature estimated in a single step, derived as:

$$Z'_{CT} = \frac{Z_{CT}^t - \sqrt{1-\bar{\alpha}_t} \cdot \epsilon_\theta\left(Z_{CT}^t, t, Z_{MR}\right)}{\sqrt{\bar{\alpha}_t}}. \tag{14}$$

This iterative denoising process gradually removes noise, guided by Z_{MR}, to reconstruct the clean CT latent feature.

We adopt a fast sampling strategy [39] that enables the selection of a subset of time steps for denoising, rather than processing all T steps sequentially. The recovered denoised CT latent feature Z_{CT}^0 is then passed to the decoder to generate the target CT modality data.

IV. Experiments

A. Dataset and Implementation

This study was conducted with Institutional Review Board (IRB) approval and utilized a multi-centre dataset comprising 268 subjects recruited from collaborating hospitals. Each subject underwent whole-body PET/MR scanning using a UIH uPMR790 scanner to acquire magnetic resonance (MR) volumes based on the Water-Fat Imaging In-Phase (WFI-IP) sequence, and whole-body PET/CT scanning using a UIH uMI780 scanner to obtain corresponding computed tomography (CT) volumes. The native MR volumes had dimensions of $549 \times 384 \times 472$ with voxel spacing of $0.97\,\text{mm} \times 0.97\,\text{mm} \times$

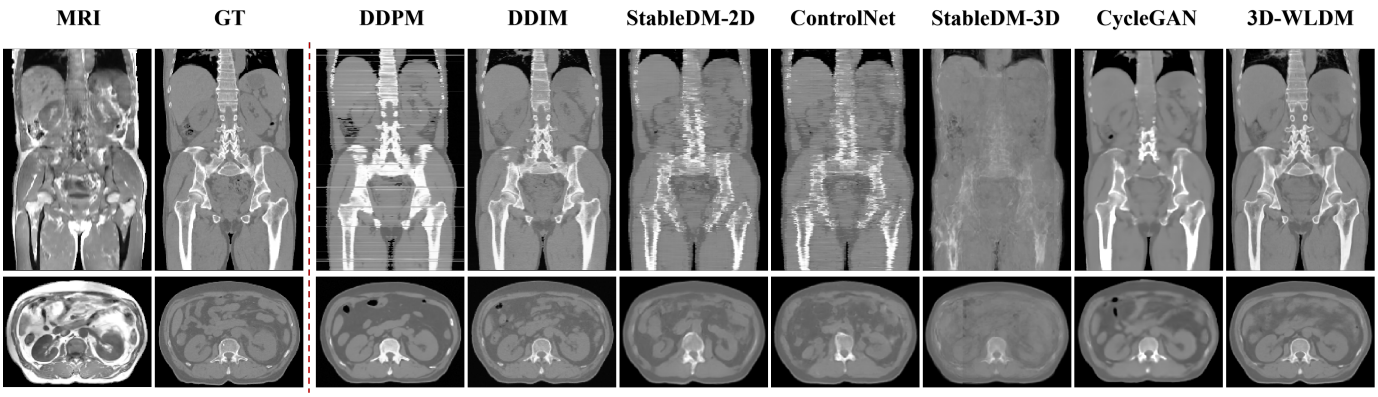

Fig. 5. Visualization of synthesized CT images in the coronal view (Row 1) and axial view (Row 2). The first column displays the input MRI images, the second column shows the Ground Truth (real CT images), and columns 3 to 9 present the synthesized CT results from various models.

2.40 mm; CT volumes were acquired at $512 \times 512 \times 900$ with voxel spacing of $0.98\,\text{mm} \times 0.98\,\text{mm} \times 1.00\,\text{mm}$.

To mitigate variability due to upper body positioning, anatomical regions above the shoulders were cropped. Black borders were removed using a rectangular envelope cropping algorithm, and all volumes were resampled to an isotropic voxel spacing of 2mm^3. Intensity normalization was performed for both modalities to range $[-1, 1]$, with CT images truncated to range $[-800, 800]$ HU to enhance contrast in soft tissues. Rigid and non-rigid intermodal registration was performed using Elastix [40], employing mutual information as the similarity metric. The final dataset was partitioned randomly into training and testing subsets in a 4:1 ratio.

During training, each volume was randomly cropped into subvolumes of size $128 \times 128 \times 128$. Data augmentation strategies included random flips, brightness mapping, and affine transformations, followed by normalization to $[-1, 1]$. Model training was performed using PyTorch with the Adam optimizer on an NVIDIA A40 GPU. The diffusion process comprised $T = 1000$ timesteps, with noise variance linearly scheduled from $\beta_1 = 0.0001$ to $\beta_T = 0.2$.

B. Comparisons of Multiple Baselines and Discussion

1) Qualitative Analysis: We performed MR-to-CT modality translation experiments to benchmark the proposed 3D Wavelet Latent Diffusion Model (3D-WLDM) against several state-of-the-art baselines, including DDPM [41], DDIM [39], StableDiffusion 2D (StableDM-2D) [37], ControlNet [34], StableDiffusion 3D (StableDM-3D) [37], and CycleGAN [42]. Representative qualitative comparisons are illustrated in Fig. 5.

In 2D transverse slice generation, DDPM occasionally produces failed outputs, notably horizontal line artifacts in coronal views. Our experiments suggest that DDPM is highly sensitive to the initial noise vector, resulting in substantial variability in output quality. These artifacts are attributed to the model's limited generative capacity. Increasing the number of diffusion steps reduces such interslice inconsistencies but further exacerbates inference inefficiency, underscoring the trade-off between quality and computational cost in DDPM. DDIM improves upon DDPM by enhancing sampling efficiency and reducing interlayer artifacts. Nevertheless, noise artifacts persist, particularly in transverse-plane reconstructions, suggesting suboptimal structural fidelity.

StableDM-2D and ControlNet generate CT images with visually realistic textures and intensity distributions that closely resemble ground truth. However, both suffer from anatomical misalignment. Structural discontinuities are evident, particularly in axial and coronal planes. These inconsistencies stem from limitations in the variational autoencoder (VAE) decoder, which, while enhancing perceptual quality, degrades structural integrity during latent encoding. The cross-attention mechanism in StableDM-2D and the condition fusion strategy in ControlNet both fall short in enforcing spatial coherence. StableDM-3D amplifies these limitations, yielding overly smooth and anatomically imprecise reconstructions, highlighting the current challenges faced by 3D diffusion frameworks in preserving high-fidelity structural information.

CycleGAN exhibits pronounced anatomical inconsistencies, especially in spinal structures, likely due to ambiguous MR signal characteristics. As a bidirectional generative framework, it tends to propagate non-physiological features introduced by low-contrast MR inputs. These artifacts degrade the fidelity of CT synthesis and compromise structural accuracy.

In contrast, our proposed 3D-WLDM demonstrates marked improvements in both visual fidelity and anatomical coherence. The integration of Wavelet Blocks enhances multiscale encoding and reconstruction, while the Structure-Modality Disentanglement module mitigates inconsistencies in latent space representation. Furthermore, the DSCA module significantly improves the generation of fine structural details during the diffusion process. Together, these components enable 3D-WLDM to produce CT reconstructions that are not only visually compelling but also structurally faithful to the underlying anatomy.

2) Quantitative Analysis: Table I summarizes the quantitative performance of all methods using four standard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), Mean Absolute Error (MAE), and Normalized Cross Correlation (NCC). The proposed 3D-WLDM achieves consistent improvements across all metrics, with PSNR gains of up to 3.98 dB (1.04 dB above the next best), SSIM improvements of 0.36 (0.02 above the next best), MAE reductions of 53.76 (7.76 below the next best), and NCC

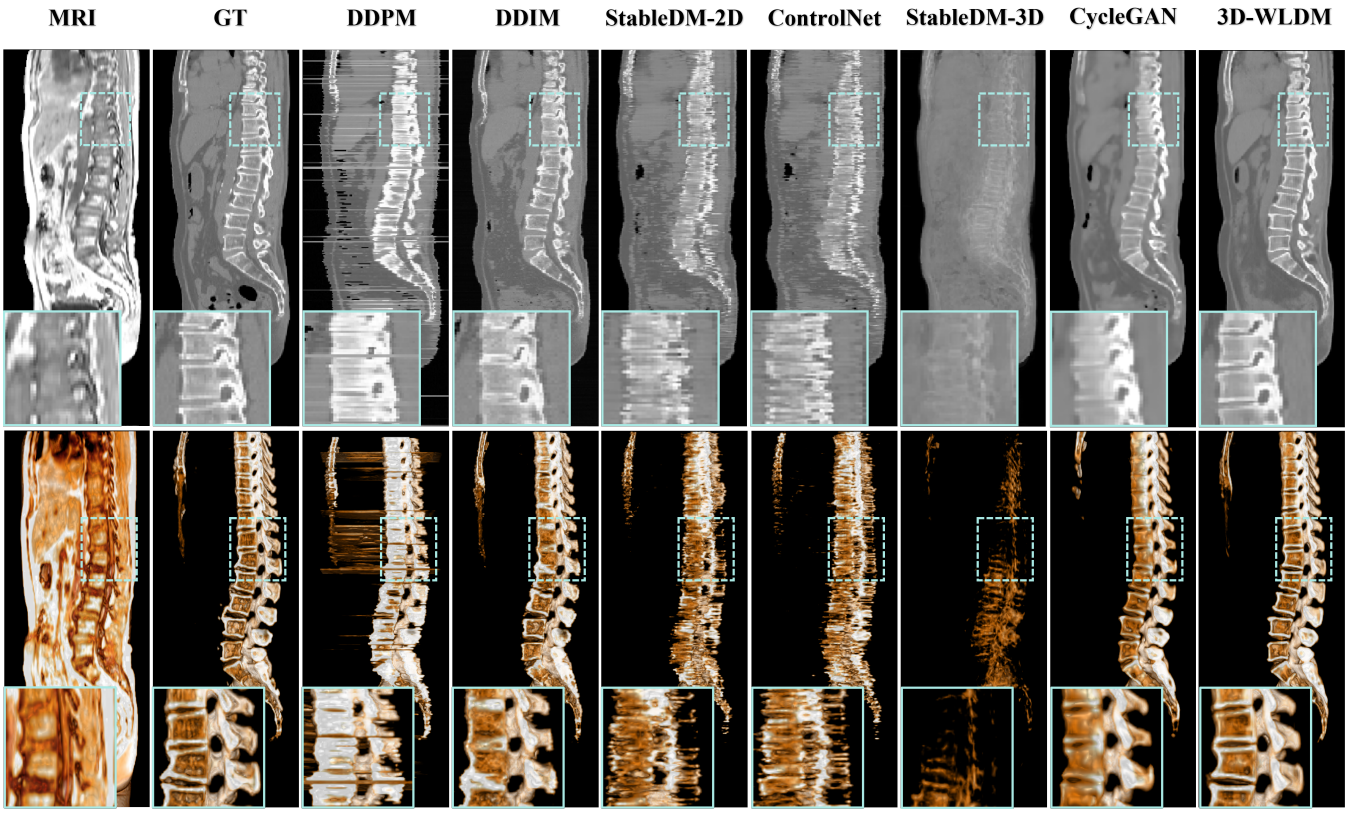

Fig. 6. Visualization of synthesized CT results on spine. The first row shows slices of the spine in the sagittal view, the second row presents the rendered 3D spine from the generated CT volumes.

TABLE I
QUANTITATIVE COMPARISONS OF DIFFERENT METHODS. BOLD VALUES INDICATE THE BEST PERFORMANCE. (↓ DENOTES LOWER IS BETTER; ↑ DENOTES HIGHER IS BETTER).

Method	Performance metrics			
	PSNR(↑)	SSIM(↑)	MAE (↓)	NCC(↑)
DDPM	21.35±3.75	0.60±0.10	101.04±45.52	0.93±0.11
DDIM	24.29±2.34	0.45±0.07	76.64±16.00	0.97±0.02
StableDM-2D	22.59±2.29	0.71±0.07	67.52±18.64	0.95±0.02
ControlNet	22.64±2.18	0.73±0.07	66.72±18.24	0.95±0.02
StableDM-3D	21.82±1.74	0.66±0.07	85.44±21.92	0.95±0.02
CycleGAN	23.44±2.31	0.79±0.06	55.04±17.28	0.96±0.02
3D-WLDM	**25.33±2.60**	**0.81±0.07**	**47.28±16.48**	**0.97±0.02**

increases of 0.04, highlighting its superiority in generating anatomically accurate and perceptually faithful CT images.

Quantitative results also reveal important discrepancies between visual quality and structural fidelity across baseline methods. For instance, DDIM generates visually sharper images in the coronal plane (Fig. 5) compared to StableDM-2D and ControlNet, yet achieves lower SSIM and higher MAE. This contrast is attributed to differences in generative pathways: latent vector–based decoding via variational autoencoders (VAEs) typically enhances perceptual quality and contrast but may fail to preserve accurate anatomical detail when compared to pixel-space diffusion outputs.

CycleGAN achieves the highest SSIM and lowest MAE among baseline methods, despite clear structural deficiencies, such as vertebral blurring and discontinuities (Fig. 5). This performance is primarily due to its strong contrast preservation and capacity for global structural consistency. However, its failure to reconstruct fine bone structures limits its clinical utility – particularly in applications such as PET/MR attenuation correction, where accurate bone delineation is critical due to its high electron density and poor visibility in MR imaging.

In contrast, 3D-WLDM delivers state-of-the-art performance across all quantitative and qualitative assessments. Its combination of multiscale wavelet encoding, structure-modality disentanglement, and enhanced structural attention enables improved anatomical fidelity, contrast consistency, and noise suppression. These properties underscore the model's potential for clinical translation, particularly in scenarios demanding high-resolution, structure-preserving CT synthesis.

3) Spine Reconstruction Analysis: The spine presents a significant challenge for MR-to-CT synthesis due to the inherently low soft-tissue contrast and blurred cortical boundaries in MR images. These limitations substantially affect the accuracy of vertebral bone reconstruction—an essential component for downstream clinical applications such as radiotherapy planning and PET/MR attenuation correction. To evaluate the fidelity of synthetic vertebral structures, sagittal-plane slices and rendered 3D spine visualizations are shown in Fig. 6, revealing notable differences among competing methods.

Consistent with observations from the axial-plane results in Fig. 5, both DDPM and DDIM fail to reconstruct anatomically accurate spinal structures. DDPM exhibits poor trabecular

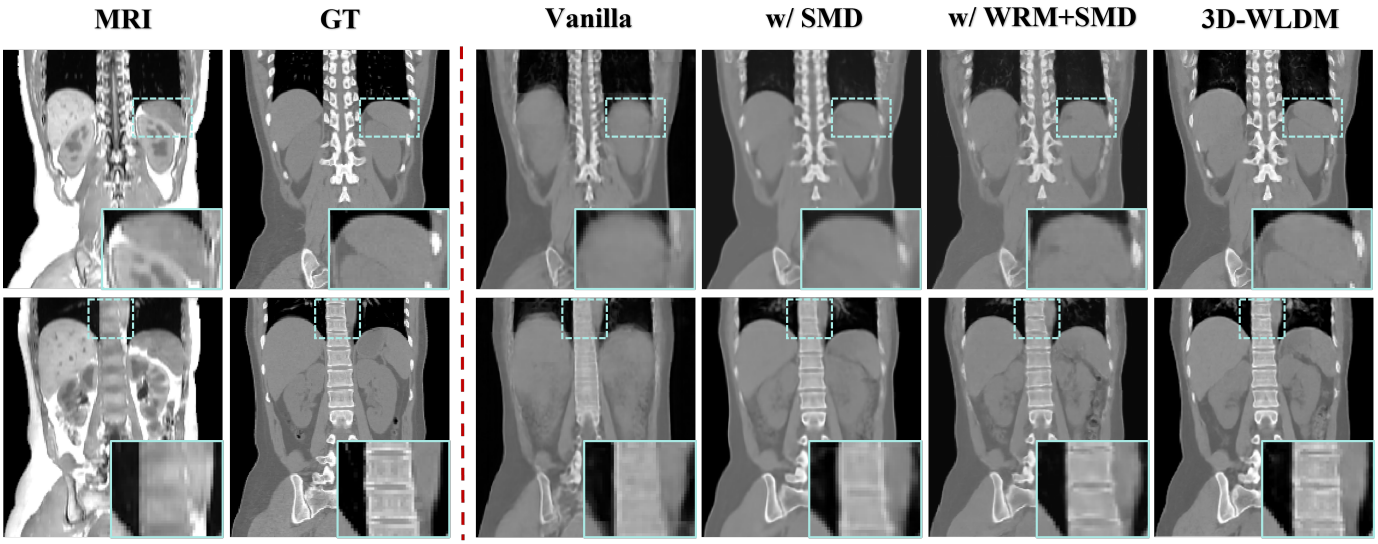

Fig. 7. Comparison of the results from ablation experiments: Vanilla (3D latent diffusion model), Wavelet Residual Module (WRM), Structure-Modality Disentanglement (SMD), and their combination with DSCA (3D-WLDM).

bone detail and fails to preserve internal vertebral architecture. Although DDIM shows improved contour alignment, it introduces contrast inversions at intervertebral boundaries, indicating erroneous intensity modeling. These findings suggest that pixel-space diffusion models offer better geometric alignment but are limited in capturing realistic image contrast and bone microstructure. StableDM-2D and ControlNet produce visually appealing outputs in the axial view, with enhanced contrast and texture fidelity. However, they fail to maintain structural continuity in the spine across sagittal slices, resulting in fragmented and anatomically implausible vertebral columns. This misalignment likely stems from inadequate structural encoding during latent-space diffusion, where spatial correlations between MR inputs and CT targets are poorly preserved. CycleGAN yields high-contrast outputs but fails to reconstruct the fine trabecular architecture of the spine. This is attributed to the network's difficulty in modeling cortical bone features, which appear hyperintense and sharply contoured in CT but indistinct in MR. Consequently, large portions of the spine are absent or inaccurately depicted, particularly in coronal slices (Fig. 5), reducing the method's utility in bone-critical applications.

In contrast, the proposed 3D-WLDM model achieves markedly improved vertebral bone reconstruction. It preserves global structural alignment with MR inputs while capturing fine-grained details of both cortical and trabecular bone. These results underscore the effectiveness of the model's wavelet-based multiscale representation and structure-modality disentanglement in overcoming the intrinsic limitations of MR imaging for bone synthesis.

C. Ablations

1) Synthesized Image Analysis: To examine the contribution of each model component within 3D-WLDM framework, we performed a series of ablation experiments, as detailed in Table II. The Vanilla model serves as the baseline, employing a standard 3D latent diffusion framework in which the MR latent representation is concatenated with Gaussian noise to form a dual-channel input, without any explicit structural or modality-specific enhancements.

Introducing SMD into the baseline enhances the model's ability to preserve anatomical structure during the generative process. By decoupling structural information from modality-specific features during pretraining, the model learns to anchor geometry consistently across the diffusion trajectory. While this improves overall spatial alignment, some fine-grained details remain suboptimal.

Incorporating WRM further improves image reconstruction by enhancing the representation of high-frequency features. Through wavelet-domain processing, WRM enriches the encoding and decoding between the image and latent spaces, yielding more precise recovery of complex anatomical structures, particularly in osseous regions.

The full model, combining WRM, SMD, and DSCA achieves the most robust performance across all evaluated metrics. DSCA reinforces the interaction between structural and modality-specific signals in the latent space, enabling the network to generate synthetic CT volumes with enhanced stability contrast, reduced artifact. Relative to the baseline, the full model improves PSNR by 2.31 dB, SSIM by 0.23, and reduces MAE by 28.08, underscoring the cumulative benefit of the proposed modules.

These quantitative findings are further supported by the qualitative results in Fig.7. The Vanilla model generates synthetic CT images with perceptual quality comparable to StableDM-3D (Fig.5), yet fails to capture modality-specific transformations with sufficient structural fidelity, particularly in anatomically complex 3D regions where spatial encoding is inherently more challenging.

Introducing SMD leads to more accurate anatomical alignment by explicitly preserving structure throughout the diffusion process. Organ and tissue boundaries become more coherent, though finer anatomical details remain underrepresented. The addition of WRM markedly improves the synthesis of

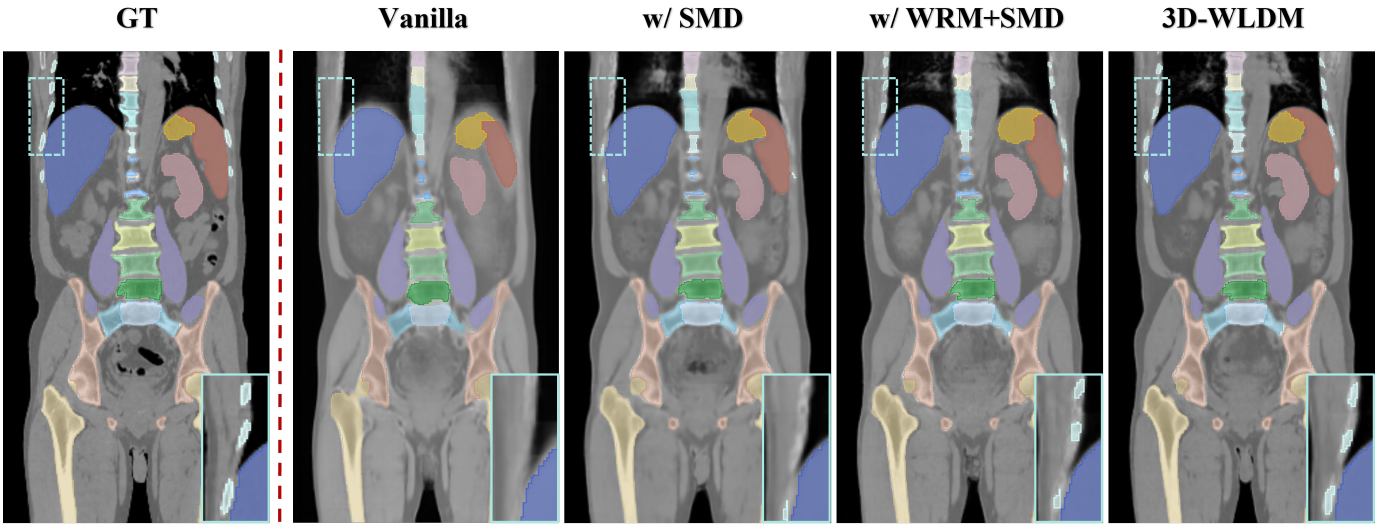

Fig. 8. The segmentation results of important anatomical structures. GT represents the segmentation results of the real CT, while others represent the segmentation results of the synthesized CT in the ablation experiments.

TABLE II
ABLATION ON PERFORMANCE OF WRM, SMD, AND THEIR COMBINATION WITH DSCA (3D-WLDM). BOLD VALUES INDICATE THE BEST PERFORMANCE ACROSS CONFIGURATIONS. (↓ DENOTES LOWER IS BETTER; ↑ DENOTES HIGHER IS BETTER).

Method	Performance metrics			
	PSNR(↑)	SSIM(↑)	MAE(↓)	NCC(↑)
Vanilla	23.02±2.45	0.58±0.09	75.36±23.84	0.96±0.03
w/ WRM	24.65±2.42	0.68±0.09	59.60±18.80	0.97±0.02
w/ WRM+SMD	25.02±2.49	0.78±0.10	52.24±19.28	0.97±0.02
3D-WLDM	**25.33±2.60**	**0.81±0.07**	**47.28±16.48**	**0.98±0.02**

TABLE III
ABLATION ON SEGMENTATION PERFORMANCE OF KEY STRUCTURES USING WRM, SMD, AND THEIR COMBINATION WITH DSCA (3D-WLDM). BOLD VALUES INDICATE THE BEST PERFORMANCE ACROSS CONFIGURATIONS. (↑DENOTES HIGHER IS BETTER).

Method	DSC(%)(↑)				
	Spleen	Kidney	Liver	Stomach	Vertebrae
Vanilla	64.2±17.9	63.3±27.6	85.6±06.0	55.7±14.2	69.8±25.1
w/ WRM	75.9±13.2	70.8±26.3	88.2±05.9	62.1±13.3	77.5±19.0
w/ WRM+SMD	76.2±10.9	74.9±24.4	89.4±04.7	61.2±14.0	78.8±17.5
3D-WLDM	**81.4±09.6**	**75.9±21.8**	**89.7±05.4**	**71.1±10.5**	**80.1±17.0**

high-frequency features, especially in regions with intricate bone morphology such as the vertebrae and ribs, by leveraging wavelet-domain encoding to reinforce multiscale structural representation.

The full model further benefits from the integration of DSCA, which refines the fusion of structural and modality cues and enhances both generation stability and contrast accuracy. As shown in Fig. 7, this results in improved delineation of abdominal organs, such as the stomach and kidneys, with sharper boundaries and more realistic texture.

Collectively, these findings substantiate our central hypothesis: that explicit disentanglement of structure and modality, combined with multiscale encoding and attention-guided feature fusion, is critical for achieving accurate, high-quality 3D medical image synthesis in cross-modality tasks.

2) Semantic segmentation Analysis: To assess the downstream utility of the synthesized CT images produced by different ablation variants, we conducted a segmentation study using the publicly available TotalSegmentator tool [43], which automatically delineates major anatomical structures. Segmentation performance was quantified using the Dice Similarity Coefficient (DSC), with results for representative organs and skeletal structures summarized in Table III. Ours model demonstrated substantial improvements in segmentation accuracy compared to the Vanilla, with DSC gains of 17.2% for the spleen, 12.6% for the kidneys, 4.1% for the liver, 15.4% for the stomach, and 10.3% for vertebrae. These improvements reflect enhanced anatomical fidelity and structural consistency in the synthesized CT volumes, underscoring the clinical relevance of the proposed framework.

To further illustrate the key contributions of each model component, Fig. 8 presents qualitative comparisons of segmented anatomical structures across model variants. The introduction of the SMD module leads to markedly improved spatial alignment of structures such as the femur, reflecting better retention of global anatomical geometry. The addition of WDM and the DSCA module respectively yields enhanced delineation of fine anatomical details, particularly in complex bony regions such as the ribs and vertebrae. Together, these findings demonstrate that improved synthesis quality directly translates to higher segmentation performance, highlighting the potential of 3D-WLDM to support clinically relevant tasks that depend on structurally accurate CT representations.

V. CONCLUSION

We present 3D-WLDM, a novel three-dimensional wavelet latent diffusion model designed for whole-body MR-to-CT modality translation. To address the limitations of conventional latent diffusion in cross-modality synthesis, our 3D-WLDM

introduces a wavelet-based residual mechanism that enhances the quality of encoding and decoding between image and latent spaces, enabling sharper structural detail and improved tissue contrast. A structure–modality disentanglement strategy further strengthens anatomical consistency by maintaining a persistent structural representation throughout the generative process. Additionally, a dual skip connection attention mechanism refines the fusion of structural and modality-specific features, resulting in greater synthesis stability and reduced artifact prevalence. Through extensive quantitative and qualitative evaluations, the model demonstrates clear advantages over the state-of-the-art GAN-based and diffusion-based methods, particularly in anatomical fidelity, contrast preservation, and structural consistency.

Beyond performance gains, our 3D-WLDM offers practical benefits for clinical imaging workflows. It enables accurate synthetic CT generation from MR input alone, transforming the inherently challenging MR-to-CT registration task into a mono-modality alignment problem, and supports a range of downstream applications, including PET/MR attenuation correction, MR-only radiation therapy planning, and MR-based organ segmentation using established CT algorithms. The method also provides CT-like visualization capabilities for MR-guided interventions without exposing patients to ionizing radiation. Despite its strengths, the model's reliance on accurate MR-CT alignment during training may limit robustness in regions affected by anatomical variation or motion. Future work will focus on incorporating semantic- and self-supervised learning strategies to mitigate this sensitivity. Moreover, extending the model's adaptability to diverse MR sequences and improving computational efficiency and uncertainty quantification will be essential steps toward clinical translation.

REFERENCES

[1] W. B. Overcast, K. M. Davis, C. Y. Ho, G. D. Hutchins, M. A. Green, B. D. Graner, and M. C. Veronesi, "Advanced imaging techniques for neuro-oncologic tumor diagnosis, with an emphasis on pet-mri imaging of malignant brain tumors," *Current Oncology Reports*, vol. 23, pp. 1–15, 2021.

[2] R. Ranjbarzadeh, A. Caputo, E. B. Tirkolaee, S. J. Ghoushchi, and M. Bendechache, "Brain tumor segmentation of mri images: A comprehensive review on the application of artificial intelligence tools," *Computers in Biology and Medicine*, vol. 152, p. 106405, 2023.

[3] C. Catana, "Attenuation correction for human pet/mri studies," *Physics in Medicine & Biology*, vol. 65, no. 23, p. 23TR02, 2020.

[4] A. M. Owrangi, P. B. Greer, and C. K. Glide-Hurst, "Mri-only treatment planning: benefits and challenges," *Physics in Medicine & Biology*, vol. 63, no. 5, p. 05TR01, 2018.

[5] S. K. Kang, H. J. An, H. Jin, J.-i. Kim, E. K. Chie, J. M. Park, and J. S. Lee, "Synthetic ct generation from weakly paired mr images using cycle-consistent gan for mr-guided radiotherapy," *Biomedical Engineering Letters*, vol. 11, no. 3, pp. 263–271, 2021.

[6] H. Yang, J. Sun, A. Carass, C. Zhao, J. Lee, J. L. Prince, and Z. Xu, "Unsupervised mr-to-ct synthesis using structure-constrained cyclegan," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 4249–4261, 2020.

[7] R. Han, C. K. Jones, J. Lee, P. Wu, P. Vagdargi, A. Uneri, P. A. Helm, M. Luciano, W. S. Anderson, and J. H. Siewerdsen, "Deformable mr-ct image registration using an unsupervised, dual-channel network for neurosurgical guidance," *Medical Image Analysis*, vol. 75, p. 102292, 2022.

[8] Y. Li, S. Xu, Y. Lu, and Z. Qi, "Ct synthesis from mri with an improved multi-scale learning network," *Frontiers in Physics*, vol. 11, p. 1088899, 2023.

[9] Y. Liu, Y. Lei, T. Wang, J. Zhou, L. Lin, T. Liu, P. Patel, W. J. Curran, L. Ren, and X. Yang, "Liver synthetic ct generation based on a dense-cyclegan for mri-only treatment planning," in *Medical Imaging 2020: Image Processing*, vol. 11313. SPIE, 2020, pp. 659–664.

[10] S. Pan, E. Abouei, J. Wynne, C.-W. Chang, T. Wang, R. L. Qiu, Y. Li, J. Peng, J. Roper, P. Patel *et al.*, "Synthetic ct generation from mri using 3d transformer-based denoising diffusion model," *Medical Physics*, vol. 51, no. 4, pp. 2538–2548, 2024.

[11] Z. Wang, Z. Zhang, W. Qi, F. Yang, and J. Xu, "Freqgan: Infrared and visible image fusion via unified frequency adversarial learning," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[12] Z. He, W. Ran, S. Liu, K. Li, J. Lu, C. Xie, Y. Liu, and H. Lu, "Low-light image enhancement with multi-scale attention and frequency-domain optimization," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, no. 4, pp. 2861–2875, 2023.

[13] Y. Ge, Z. Xue, T. Cao, and S. Liao, "Unpaired whole-body mr to ct synthesis with correlation coefficient constrained adversarial learning," in *Medical Imaging 2019: Image Processing*, vol. 10949. SPIE, 2019, pp. 28–35.

[14] Z. Wang, Y. Yang, Y. Chen, T. Yuan, M. Sermesant, H. Delingette, and O. Wu, "Mutual information guided diffusion for zero-shot cross-modality medical image translation," *IEEE Transactions on Medical Imaging*, vol. 43, no. 8, pp. 2825–2838, 2024.

[15] X. Meng, K. Sun, J. Xu, X. He, and D. Shen, "Multi-modal modality-masked diffusion network for brain mri synthesis with random modality missing," *IEEE Transactions on Medical Imaging*, vol. 43, no. 7, pp. 2587–2598, 2024.

[16] T. Xie, C. Cao, Z.-x. Cui, Y. Guo, C. Wu, X. Wang, Q. Li, Z. Hu, T. Sun, Z. Sang *et al.*, "Synthesizing pet images from high-field and ultra-high-field mr images using joint diffusion attention model," *Medical Physics*, vol. 51, no. 8, pp. 5250–5269, 2024.

[17] Z. Li, Y. Zhou, H. Wei, C. Ge, and J. Jiang, "Towards extreme image compression with latent feature guidance and diffusion prior," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[18] Y. Fu, Y. Lei, J. Zhou, T. Wang, S. Y. David, J. J. Beitler, W. J. Curran, T. Liu, and X. Yang, "Synthetic ct-aided mri-ct image registration for head and neck radiotherapy," in *Medical Imaging 2020: Biomedical Applications in Molecular, Structural, and Functional Imaging*, vol. 11317. SPIE, 2020, pp. 572–578.

[19] O. Dalmaz, M. Yurt, and T. Çukur, "Resvit: residual vision transformers for multimodal medical image synthesis," *IEEE Transactions on Medical Imaging*, vol. 41, no. 10, pp. 2598–2614, 2022.

[20] M. C. Florkow, F. Guerreiro, F. Zijlstra, E. Seravalli, G. O. Janssens, J. H. Maduro, A. C. Knopf, R. M. Castelein, M. van Stralen, B. W. Raaymakers *et al.*, "Deep learning-enabled mri-only photon and proton therapy treatment planning for paediatric abdominal tumours," *Radiotherapy and Oncology*, vol. 153, pp. 220–227, 2020.

[21] K. Kläser, T. Varsavsky, P. Markiewicz, T. Vercauteren, A. Hammers, D. Atkinson, K. Thielemans, B. Hutton, M. J. Cardoso, and S. Ourselin, "Imitation learning for improved 3d pet/mr attenuation correction," *Medical Image Analysis*, vol. 71, p. 102079, 2021.

[22] H. Sun, Q. Xi, R. Fan, J. Sun, K. Xie, X. Ni, and J. Yang, "Synthesis of pseudo-ct images from pelvic mri images based on an md-cyclegan model for radiotherapy," *Physics in Medicine & Biology*, vol. 67, no. 3, p. 035006, 2022.

[23] D. Cusumano, J. Lenkowicz, C. Votta, L. Boldrini, L. Placidi, F. Catucci, N. Dinapoli, M. V. Antonelli, A. Romano, V. De Luca *et al.*, "A deep learning approach to generate synthetic ct in low field mr-guided adaptive radiotherapy for abdominal and pelvic cases," *Radiotherapy and Oncology*, vol. 153, pp. 205–212, 2020.

[24] K. N. Brou Boni, J. Klein, A. Gulyban, N. Reynaert, and D. Pasquier, "Improving generalization in mr-to-ct synthesis in radiotherapy by using an augmented cycle generative adversarial network with unpaired data," *Medical physics*, vol. 48, no. 6, pp. 3003–3010, 2021.

[25] D. Nie, R. Trullo, J. Lian, L. Wang, C. Petitjean, S. Ruan, Q. Wang, and D. Shen, "Medical image synthesis with deep convolutional adversarial networks," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 12, pp. 2720–2730, 2018.

[26] Z. Shi, P. Mettes, G. Zheng, and C. Snoek, "Frequency-supervised mr-to-ct image synthesis," in *Deep Generative Models, and Data Augmentation, Labelling, and Imperfections*. Springer, 2021, pp. 3–13.

[27] G. Dovletov, D. D. Pham, S. Lörcks, J. Pauli, M. Gratz, and H. H. Quick, "Grad-cam guided u-net for mri-based pseudo-ct synthesis," in *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022, pp. 2071–2075.

[28] G. Dovletov, S. Lörcks, J. Pauli, M. Gratz, and H. H. Quick, "Double grad-cam guidance for improved mri-based pseudo-ct synthesis," in *BVM Workshop*. Springer, 2023, pp. 45–50.

[29] B. Zhao, T. Cheng, X. Zhang, J. Wang, H. Zhu, R. Zhao, D. Li, Z. Zhang, and G. Yu, "Ct synthesis from mr in the pelvic area using residual transformer conditional gan," *Computerized Medical Imaging and Graphics*, vol. 103, p. 102150, 2023.

[30] L. Zhong, Z. Chen, H. Shu, K. Zheng, Y. Li, W. Chen, Y. Wu, J. Ma, Q. Feng, and W. Yang, "Multi-scale tokens-aware transformer network for multi-region and multi-sequence mr-to-ct synthesis in a single model," *IEEE Transactions on Medical Imaging*, vol. 43, no. 2, pp. 794–806, 2023.

[31] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-to-image diffusion models," in *ACM SIGGRAPH 2022 Conference Proceedings*, 2022, pp. 1–10.

[32] X. Chen, J. Tan, T. Wang, K. Zhang, W. Luo, and X. Cao, "Toward real-world blind face restoration with generative diffusion prior," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, no. 9, pp. 8494–8508, 2024.

[33] Z. Wang, L. Zhao, and W. Xing, "Stylediffusion: Controllable disentangled style transfer via diffusion models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7677–7689.

[34] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3836–3847.

[35] P. Friedrich, J. Wolleb, F. Bieder, A. Durrer, and P. C. Cattin, "Wdm: 3d wavelet diffusion models for high-resolution medical image synthesis," in *MICCAI Workshop on Deep Generative Models*. Springer, 2024, pp. 11–21.

[36] A. Tapp, A. Parida, C. Zhao, V. Lam, N. Lepore, S. M. Anwar, and M. G. Linguraru, "Mr to ct synthesis using 3d latent diffusion," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2024, pp. 1–5.

[37] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 684–10 695.

[38] L. Jiang, Y. Mao, X. Wang, X. Chen, and C. Li, "Cola-diff: Conditional latent diffusion model for multi-modal mri synthesis," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023, pp. 398–408.

[39] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *International Conference on Learning Representations*, 2021. [Online]. Available: https://openreview.net/forum?id=St1giarCHLP

[40] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, "Elastix: a toolbox for intensity-based medical image registration," *IEEE Transactions on Medical Imaging*, vol. 29, no. 1, pp. 196–205, 2009.

[41] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.

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

[43] J. Wasserthal, H.-C. Breit, M. T. Meyer, M. Pradella, D. Hinck, A. W. Sauter, T. Heye, D. T. Boll, J. Cyriac, S. Yang *et al.*, "Totalsegmentator: robust segmentation of 104 anatomic structures in ct images," *Radiology: Artificial Intelligence*, vol. 5, no. 5, 2023.