Editorial

I. Introduction

THE prevailing understanding in the field of machine learning and deep learning (ML/DL) is that, given a high-quality dataset, one can effectively learn data-related priors through supervised learning. However, in medical imaging, this assumption faces two critical challenges: 1) high-quality training data are often scarce and 2) data are highly heterogeneous, stemming from different imaging scanners, protocols, or populations at various institutions. This diversity makes it impractical to represent the data with a single, universal prior using traditional methods, leading to limited generalizability in medical imaging tasks.

The recent advent of score-based generative diffusion models has shown promise in addressing these challenges by learning complex data distribution priors in an unsupervised manner [1], [2]. The theoretical foundations of score-based generative models provide a robust and interpretable framework, suggesting their potential to overcome the limitations associated with traditional methods in medical imaging. As a result, diffusion models have gained significant attention within the medical imaging community [3], [4], [5], [6], [7], [8].

Following consultations with several experts in the field, we identified the growing need for a Special Issue to spotlight these pivotal advancements. This Special Issue of IEEE Transactions on Medical Imaging focuses on cutting-edge diffusion models and their applications in medical imaging. It attracted a large interest from the scientific community; i.e., 67 manuscript submissions, of which 20 articles were selected for publication. Each article underwent a regular peer-review process, typically involving two rounds of revisions, with evaluations from between three to six experts in the field. The following section provides an overview of the included articles and key insights from each manuscript in the Special Issue.

II. Topics of This Special Issue

Given the immense potential of diffusion models for building capable deep learning methods, there is no surprise that the articles in this Special Issue ranged broadly in terms of the imaging modality and task that they tackled. Topics of interest ranged from image reconstruction, artifact reduction, unsupervised quality assessment, segmentation, etc. The submitted works highlight how score-based diffusion models are shaping the future of medical imaging, providing new pathways to improve image quality, enhance computational efficiency, and solve longstanding clinical challenges. These contributions demonstrate how diffusion models can bridge the gap between research innovation and practical clinical application by addressing key issues such as data sparsity, radiation dose reduction, and diagnostic accuracy.

The pie chart in Fig. 1 summarizes how many papers focused on a particular topic of interest. One can notice a slight majority focus on medical image reconstruction, specifically addressing the challenges of sparse-view and limited-angle CT, as well as fast MRI reconstruction. The integration of advanced techniques such as multi-frequency priors, wavelet transforms, and data consistency terms significantly boost reconstruction accuracy and robustness, offering solutions to reduce radiation exposure while maintaining high image quality. In addition, several papers explored diverse topics, including metal artifact reduction, where diffusion models were employed to restore degraded CT regions, and unsupervised quality assessment, which introduced novel methods to evaluate medical images without relying on ground truth labels. Further contributions encompassed advancements in segmentation, counterfactual generation, data augmentation, radiotherapy planning, virtual staining in pathology, and multimodal 3-D image generation. The eminent diversity of the considered applications is a testament to the versatility of the diffusion modeling framework. Below, we summarize the key insights from each manuscript in the Special Issue.

III. Articles Included in the Special Issue

A. Medical Image Reconstruction Using Diffusion Models

These methods focus on improving the accuracy and stability of medical image reconstruction through the application of diffusion models. They are particularly effective in scenarios such as sparse-view and limited-angle CT reconstruction, and fast MRI reconstruction. By incorporating techniques such as multi-frequency priors, wavelet transforms, and data consistency terms, these models not only enhance reconstruction quality and efficiency but also reduce radiation dose and computational cost. These approaches improve diagnostic precision and expand clinical applicability by addressing the challenges of acquiring high-quality imaging data.

In [A1], Li et al. introduced the dual-domain collaborative diffusion sampling (DCDS) model to enhance sparse-view CT reconstruction. Unlike previous methods that focus exclusively on either sinogram or image domains, DCDS integrates both domains through a collaborative diffusion mechanism. This approach improves sinogram recovery and image generation synergistically. Extensive evaluations on simulations, phantoms, and clinical datasets showed that DCDS significantly outperforms state-of-the-art methods.

Digital Object Identifier 10.1109/TMI.2024.3464268

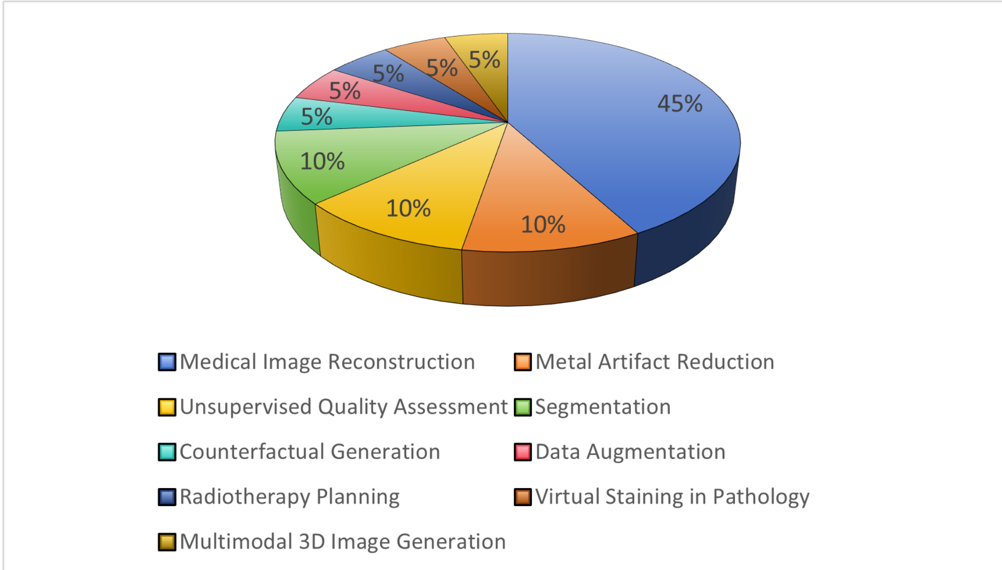

Fig. 1. Topics included in this Special Issue.

In [A2], Xu et al. introduced a stage-by-stage wavelet optimization refinement diffusion (SWORD) model for sparse-view CT reconstruction. This approach used wavelet transforms to enhance the robustness of diffusion models by addressing instability issues in the original sinogram domain. By integrating low- and high-frequency generative models within a unified mathematical framework and applying these models to wavelet-decomposed components, the SWORD model demonstrated excellent performance.

In [A3], Liu et al. proposed an unsupervised sparse-view spectral CT reconstruction and material decomposition algorithm using a multi-channel score-based generative model (SGM). This method addresses the challenge of inaccurate contrast agent quantification due to streaking artifacts in sparse-view scans. The approach involves training the SGM with multi-energy and tissue images to generate accurate multi-energy and tissue images from sparse-view projections. These generated images are then used in a material decomposition algorithm to determine the distribution and content of contrast agents. The method was validated in mouse scanning experiments, demonstrating its effectiveness in reducing radiation dose and improving quantification.

In [A4], Zhang et al. identified several limitations in applying score-based generative models (SGMs) to limited-angle CT (LACT) reconstruction, including the neglect of the directional distribution of artifacts and the varying properties of different frequency components. To address these challenges, they proposed a wavelet-inspired score-based model (WISM), which integrates wavelet transforms to model probability densities in both the image and wavelet domains. This method retains the directional properties of artifacts while simplifying density modeling by decomposing them into component-specific models. The unified sampling process, informed by observational data, produced high-quality LACT reconstructions, with experimental results showing superior performance compared to existing methods.

In [A5], Wang et al. introduced the time-reversion fast-sampling (TIFA) score-based model for LACT reconstruction. TIFA addresses the high computational cost of traditional score-based generative models (SGMs) by employing a rapid-sampling strategy that includes jump sampling, time-reversion with re-sampling, and compressed sampling. This approach improves reconstruction quality and efficiency by reducing the number of sampling steps. The experiments on various datasets demonstrated that TIFA outperforms existing methods with fewer steps, achieving high-quality reconstructions in as few as ten steps.

In [A6], Wu et al. proposed a multi-channel optimization generative model (MOGM) to achieve stable ultra-sparse-view CT reconstruction. The MOGM incorporated a novel data consistency term into the stochastic differential equation model, relying exclusively on original data to constrain the generative outcomes. Furthermore, an inference strategy was developed to trace back to the ground truth, which enhanced reconstruction stability. The rxperimental results demonstrated that MOGM consistently outperformed alternative methods, even when reconstructing from as few as ten and seven views.

In [A7], Chen et al. introduced PRECISION, a physics-constrained and noise-controlled diffusion model, to enhance the quality of material basis images and the quantitative accuracy of elemental composition in photon counting detectors in computed tomography (PCCT). This model addresses the issues caused by imperfect noise modeling and hand-crafted regularization in existing direct material basis image reconstruction approaches. PRECISION learns distribution-level regularization from ideal material basis images and samples the optimal images under physical constraints specific to PCCT systems and subject data, showing potential for improved image quality and elemental composition quantification in PCCT.

In [A8], Guan et al. introduced the correlated and multi-frequency diffusion model (CM-DM) to improve MRI reconstruction accuracy in highly under-sampled images. Unlike existing methods, CM-DM effectively combines high-frequency operators to form a multi-frequency prior, enhancing noise reduction and accelerating convergence in the diffusion process. The experimental results demonstrated that CM-DM achieved about 2 dB improvement in PSNR over state-of-the-art methods.

In [A9], Cui et al. introduced a physics-informed diffusion model for k-space interpolation. The approach utilizes interpolatable physical priors to interpolate high-frequency (HF) k-space data from low-frequency (LF) data. By connecting HF interpolation with the reverse heat diffusion process, a heat diffusion model was developed to generate missing HF data with enhanced accuracy. A k-space structural low-rank model was also integrated as a data fidelity term, further boosting the framework's performance. The experimental results demonstrated that this method significantly outperforms traditional image-domain diffusion models.

B. Metal Artifact Reduction in CT Imaging

These methods focus on reducing metal artifacts in CT imaging using diffusion models. Metal artifacts, caused by the presence of metal implants, can severely distort CT images, hindering accurate diagnosis. The diffusion models proposed in these studies aim to address this issue by either inpainting missing sinogram data or iteratively restoring degraded regions. By improving the handling of metal-induced distortions, these approaches enhance image clarity and diagnostic reliability.

In [A10], Karageorgos et al. proposed a denoising diffusion probabilistic model (DDPM) for inpainting missing sinogram data for improved metal artifact reduction (MAR). The model was trained unconditionally, without relying on information about metal objects, enhancing its generalizability across various types of metal implants. The DDPM-MAR technique was evaluated on clinical CT images with virtually introduced metal objects, demonstrating superior quality compared to other methods. This approach shows promise in improving the effectiveness of MAR, thereby enhancing the accuracy of CT.

In [A11], Liu et al. proposed an unsupervised MAR method using a diffusion model. Initially, a diffusion model was trained on CT images without metal artifacts. Subsequently, diffusion priors were iteratively introduced in both the sinogram and image domains to restore the degraded regions caused by metal artifacts. This approach outperformed existing unsupervised methods, including those based on diffusion models.

C. Unsupervised and Blind Image Quality Assessment

Unsupervised methods for medical imaging are essential for scenarios where labeled data is limited or unavailable. These studies apply diffusion models to either unsupervised model medical images or assess their quality without the need for explicit ground truth labels. These methods provide solutions for tasks such as dehazing in ultrasound and blind image quality assessment.

In [A12], Stevens et al. proposed a joint posterior sampling framework utilizing two diffusion models for clean ultrasound and haze distributions respectively. They introduced effective techniques for training diffusion models on radiofrequency ultrasound data. The proposed dehazing method proved effective in removing haze while preserving signals from weakly reflected tissue, as demonstrated through experiments on both in-vitro and in-vivo cardiac datasets.

In [A13], Shi et al. introduced a novel blind image quality assessment (BIQA) metric that emulates the human visual system's (HVS) active inference process. It utilized a DDPM to predict primary content and derived a dissimilarity map to assess the relationship between distorted images and their primary content. This multi-channel image, combining the distorted image and dissimilarity map, was evaluated using a transformer-based quality evaluator. The proposed method achieved competitive performance on a low-dose CT dataset.

D. Segmentation and Image Label Pair Generation

Segmentation and image-label generation are crucial for tasks like disease identification and treatment planning in medical imaging. However, data scarcity, class imbalance, and the complexity of accurately segmenting medical images often pose significant challenges. The methods in this category utilize diffusion models to generate high-quality image-label pairs and improve segmentation accuracy. These approaches enhance generalization and performance across different datasets, especially in challenging scenarios like small object segmentation or highly imbalanced datasets.

In [A14], Chen et al. proposed HiDiff, a hybrid diffusion framework for medical image segmentation that integrates discriminative and generative models. HiDiff combines a conventional discriminative segmentor with a novel binary Bernoulli diffusion model (BBDM) as a refiner. The segmentor provides initial segmentation masks, which the BBDM refines by modeling the underlying data distribution. Trained in an alternating manner, HiDiff showed superior performance on various segmentation tasks, including small object segmentation and generalization to new datasets, compared to state-of-the-art methods.

In [A15], Huang et al. proposed a framework for generating diverse and balanced image-label pairs for retinal layer segmentation in optical coherence tomography (OCT) images. This framework initially generates varied layer masks and subsequently creates corresponding OCT images using two customized diffusion probabilistic models. To address data imbalance, the approach incorporates pathological-related conditions and utilizes a structure modeling technique to transfer knowledge from less pathological to highly pathological samples. Extensive experiments on two public datasets demonstrated that this method produces OCT images with superior quality and diversity, enhancing performance in downstream segmentation tasks.

E. Counterfactual Generation and Data Augmentation

In medical imaging, data scarcity and class imbalance are common challenges that can hinder the training of robust models. The methods in this category employ diffusion models to generate counterfactual images and augment existing datasets, improving model performance in downstream tasks such as classification, localization, and disease detection. These techniques are particularly valuable in scenarios where obtaining real-world data is costly or time-consuming.

In [A16], Wang et al. proposed a score-based counterfactual generation (SCG) framework to address data scarcity and imbalance by creating counterfactual images from latent space. It incorporated a learnable FuzzyBlock into the classifier to manage uncertainties from external physical factors. The SCG framework showed significant improvements in both classification and lesion localization tasks, achieving an average enhancement of 3%–5% compared to state-of-the-art methods.

In [A17], Deshpande et al. systematically evaluated the capacity of DDPMs to learn spatial context relevant to medical imaging. The study used stochastic context models (SCMs) to generate training data and quantitatively evaluated DDPMs' performance in reproducing the spatial context. The results showed that DDPMs effectively generate contextually accurate images, offering significant advantages over generative adversarial networks (GANs) for data augmentation tasks. This assessment provides new insights into the potential of DDPMs for medical imaging applications.

F. Innovative Diffusion Models for Specialized Applications

These methods introduce cutting-edge diffusion models tailored for specific medical imaging applications, such as radiotherapy planning, virtual staining in pathology, and multimodal 3-D image generation.

In [A18], Zhang et al. proposed a distance-aware diffusion model (DoseDiff) for accurate dose distribution prediction in radiotherapy treatment planning. This model utilized signed distance maps (SDMs) obtained from target and organ-at-risk masks, incorporating these along with CT images into a sequence of denoising steps. The proposed multi-encoder and multi-scale fusion network (MMFNet) enhanced feature-level information fusion between CT images and SDMs using multi-scale and transformer-based modules. Evaluations on in-house and public datasets showed that DoseDiff outperformed existing methods in both quantitative performance and visual quality.

In [A19], He et al. developed PST-Diff, a method for generating virtual immunohistochemistry (IHC) images from hematoxylin and eosin (HE) images using diffusion models. PST-Diff addresses the challenges of high costs and information loss associated with IHC by allowing simultaneous viewing of multiple staining results from the same tissue slide. The method incorporates an asymmetric attention mechanism (AAM) and a latent transfer (LT) module to ensure pathological consistency and reduce bias. In addition, a conditional frequency guidance (CFG) module maintains structural consistency. PST-Diff demonstrated impressive effectiveness and generalizability in generating stable and functionally accurate IHC images, with promise for clinical virtual staining and pathological image analysis.

In [A20], Xu et al. proposed a method for generating high-resolution 3-D lung CT images guided by textual information and anatomical components. Their approach relies on a hierarchical scheme using a modified UNet architecture to first create low-resolution images from text and then refine these into high-resolution volumetric data. Additional anatomical guidance is provided through segmentation masks for vascular, airway, and lobular structures. As evaluated by radiologists, the proposed method demonstrated superior performance over GAN and other diffusion-based models in preserving anatomical details and texture.

IV. Conclusion

This Special Issue was curated to provide a unique snapshot of the latest advances in AI-based medical imaging, with a particular emphasis on the growing influence and potential of diffusion models. By assembling a collection of cutting-edge algorithms and systems built upon modern generative AI frameworks, we aim for this issue to be a valuable resource for the community. We also hope it will inspire future studies and innovations in imaging methods and tools within this domain.

We believe that future of score-based generative models, especially physics-inspired diffusion models, in medical imaging, is promising, with the potential to transform clinical practice. Over the coming years, we anticipate notable advancements in image quality, predictive accuracy, interpretability, and computational efficiency, alongside broader clinical adoption following large-scale validations. The integration of these advanced models into clinical workflows could enable faster, higher quality imaging, even in challenging scenarios such as low-dose or rapid imaging protocols. Ultimately, these technological advances could play a pivotal role in improving patient outcomes and setting new standards for disease detection, monitoring, and personalized healthcare.

Before closing, we would like to extend our heartfelt thanks to all the authors who submitted their valuable work to this Special Issue, and to the reviewers who generously contributed their time and expertise to provide insightful feedback on the submissions. We also express our deepest gratitude to the Editor-in-Chief of IEEE Transactions on Medical Imaging, Prof. Leslie Ying, and the Managing Editor, Prof. Rutao Yao, for giving us the opportunity to organize this Special Issue and for their unwavering guidance and support throughout the process.

Dong Liang, *Guest Editor*
Institute of Biomedical and Health Engineering
Shenzhen Institutes of Advanced Technology
Chinese Academy of Sciences
Shenzhen 518055, China
e-mail: dong.liang@siat.ac.cn

Daniel Rueckert, *Guest Editor*
Klinikum rechts der Isar Technical University of Munich
80333 Munich, Germany
Department of Computing Imperial College London
SW7 2AZ London, U.K.
e-mail: daniel.rueckert@tum.de

Ge Wang, *Guest Editor*
Biomedical Imaging Center
Rensselaer Polytechnic Institute
Troy, NY 12180 USA
e-mail: ge-wang@ieee.org

Tolga Cukur, *Guest Editor*
Department of Electrical and Electronics Engineering
Bilkent University
06800 Ankara, Türkiye
e-mail: cukur@ee.bilkent.edu.tr

Hengyong Yu, *Guest Editor*
Department of Electrical and Computer Engineering
University of Massachusetts Lowell
Lowell, MA 01854 USA
e-mail: Hengyong_Yu@uml.edu

Appendix: Related Articles

[A1] Z. Li et al., "Dual-domain collaborative diffusion sampling for multi-source stationary computed tomography reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3398–3411, Oct. 2024, doi: 10.1109/TMI.2024.3420411.

[A2] K. Xu, S. Lu, B. Huang, W. Wu, and Q. Liu, "Stage-by-stage wavelet optimization refinement diffusion model for sparse-view CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3412–3424, Oct. 2024, doi: 10.1109/TMI.2024.3355455.

[A3] Y. Liu, X. Zhou, C. Wei, and Q. Xu, "Sparse-view spectral CT reconstruction and material decomposition based on multi-channel SGM," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3425–3435, Oct. 2024, doi: 10.1109/TMI.2024.3413085.

[A4] J. Zhang, H. Mao, X. Wang, Y. Guo, and W. Wu, "Wavelet-inspired multi-channel score-based model for limited-angle CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3436–3448, Oct. 2024, doi: 10.1109/TMI.2024.3367167.

[A5] Y. Wang, Z. Li, and W. Wu, "Time-reversion fast-sampling score-based model for limited-angle CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3449–3460, Oct. 2024, doi: 10.1109/TMI.2024.3418838.

[A6] W. Wu, J. Pan, Y. Wang, S. Wang, and J. Zhang, "Multi-channel optimization generative model for stable ultra-sparse-view CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3461–3475, Oct. 2024, doi: 10.1109/TMI.2024.3376414.

[A7] R. Chen, Z. Zhang, G. Quan, Y. Du, Y. Chen, and Y. Li, "PRECISION: A physics-constrained and noise-controlled diffusion model for photon counting computed tomography," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3476–3489, Oct. 2024, doi: 10.1109/TMI.2024.3440651.

[A8] Y. Guan, C. Yu, Z. Cui, H. Zhou, and Q. Liu, "Correlated and multi-frequency diffusion modeling for highly under-sampled MRI reconstruction," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3490–3502, Oct. 2024, doi: 10.1109/TMI.2024.3381610.

[A9] Z. Cui et al., "Physics-informed DeepMRI: K-space interpolation meets heat diffusion," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3503–3520, Oct. 2024, doi: 10.1109/TMI.2024.3462988.

[A10] G. M. Karageorgos et al., "A denoising diffusion probabilistic model for metal artifact reduction in CT," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3521–3532, Oct. 2024, doi: 10.1109/TMI.2024.3416398.

[A11] X. Liu, Y. Xie, S. Diao, S. Tan, and X. Liang, "Unsupervised CT metal artifact reduction by plugging diffusion priors in dual domains," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3533–3545, Oct. 2024, doi: 10.1109/TMI.2024.3351201.

[A12] T. S. W. Stevens, F. C. Meral, J. Yu, I. Z. Apostolakis, J.-L. Robert, and R. J. G. Van Sloun, "Dehazing ultrasound using diffusion models," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3546–3558, Oct. 2024, doi: 10.1109/TMI.2024.3363460.

[A13] Y. Shi, W. Xia, G. Wang, and X. Mou, "Blind CT image quality assessment using DDPM-derived content and transformer-based evaluator," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3559–3569, Oct. 2024, doi: 10.1109/TMI.2024.3418652.

[A14] T. Chen, C. Wang, Z. Chen, Y. Lei, and H. Shan, "HiDiff: Hybrid diffusion framework for medical image segmentation," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3570–3583, Oct. 2024, doi: 10.1109/TMI.2024.3424471.

[A15] K. Huang et al., "Diverse data generation for retinal layer segmentation with potential structure modelling," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3584–3595, Oct. 2024, doi: 10.1109/TMI.2024.3384484.

[A16] K. Wang, Z. Chen, M. Zhu, Z. Li, J. Weng, and T. Gu, "Score-based counterfactual generation for interpretable medical image classification and lesion localization," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3596–3607, Oct. 2024, doi: 10.1109/TMI.2024.3375357.

[A17] R. Deshpande, M. Özbey, H. Li, M. A. Anastasio, and F. J. Brooks, "Assessing the capacity of a denoising diffusion probabilistic model to reproduce spatial context," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3608–3620, Oct. 2024, doi: 10.1109/TMI.2024.3414931.

[A18] Y. Zhang, C. Li, L. Zhong, Z. Chen, W. Yang, and X. Wang, "DoseDiff: Distance-aware diffusion model for dose prediction in radiotherapy," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3621–3633, Oct. 2024, doi: 10.1109/TMI.2024.3383423.

[A19] Y. He et al., "PST-diff: Achieving high-consistency stain transfer by diffusion models with pathological and structural constraints," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3634–3647, Oct. 2024, doi: 10.1109/TMI.2024.3430825.

[A20] Y. Xu et al., "MedSyn: Text-guided anatomy-aware synthesis of high-fidelity 3D CT images," *IEEE Trans. Med. Imag.*, vol. 43, no. 10, pp. 3648–3660, Oct. 2024, doi: 10.1109/TMI.2024.3415032.

References

[1] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 6840–6851.

[2] Y. Song et al., "Score-based generative modeling through stochastic differential equations," in *Proc. ICLR*, 2021, pp. 1–36.

[3] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated MRI," *Med. Image Anal.*, vol. 80, Aug. 2022, Art. no. 102479.

[4] R. Graf et al., "Denoising diffusion-based MRI to CT image translation enables automated spinal segmentation," *Eur. Radiol. Experim.*, vol. 7, no. 1, p. 70, Nov. 2023.

[5] M. Selim et al., "DiffusionCT: Latent diffusion model for CT image standardization," in *Proc. AMIA Annu. Symp.*, 2023, p. 624.

[6] M. Özbey et al., "Unsupervised medical image translation with adversarial diffusion models," *IEEE Trans. Med. Imag.*, vol. 42, no. 12, pp. 3524–3539, Dec. 2023.

[7] C. Cao et al., "High-frequency space diffusion model for accelerated MRI," *IEEE Trans. Med. Imag.*, vol. 43, no. 5, pp. 1853–1865, May 2024.

[8] Z. Li, Y. Wang, J. Zhang, W. Wu, and H. Yu, "Two-and-a-half order score-based model for solving 3D ill-posed inverse problems," *Comput. Biol. Med.*, vol. 168, Jan. 2024, Art. no. 107819.