

# SAGCNet: Spatial-Aware Graph Completion Network for Missing Slice Imputation in Population CMR Imaging

Junkai Liu<sup>1</sup>, Nay Aung<sup>4,5</sup>, Theodoros N. Arvanitis<sup>1</sup>, Stefan K. Piechnik<sup>2</sup>, Joao A C Lima<sup>3</sup>, Steffen E. Petersen<sup>4,5</sup>, and Le Zhang<sup>1,4(✉)</sup>

<sup>1</sup> School of Engineering, College of Engineering and Physical Sciences, University of Birmingham, Birmingham, UK

<sup>2</sup> Oxford Center for Clinical Magnetic Resonance Research (OCMR), Division of Cardiovascular Medicine, John Radcliffe Hospital, University of Oxford, Oxford, UK

<sup>3</sup> Division of Cardiology, Johns Hopkins University School of Medicine, Baltimore, Maryland, USA

<sup>4</sup> William Harvey Research Institute, NIHR Barts Biomedical Research Centre, Queen Mary University London, London, UK

<sup>5</sup> Barts Heart Centre, St Bartholomew's Hospital, Barts Health NHS Trust, West Smithfield, London, UK

jx11920@student.bham.ac.uk; l.zhang.16@bham.ac.uk

**Abstract.** Magnetic resonance imaging (MRI) provides detailed soft-tissue characteristics that assist in disease diagnosis and screening. However, the accuracy of clinical practice is often hindered by missing or unusable slices due to various factors. Volumetric MRI synthesis methods have been developed to address this issue by imputing missing slices from available ones. The inherent 3D nature of volumetric MRI data, such as cardiac magnetic resonance (CMR), poses significant challenges for missing slice imputation approaches, including (1) the difficulty of modeling local inter-slice correlations and dependencies of volumetric slices, and (2) the limited exploration of crucial 3D spatial information and global context. In this study, to mitigate these issues, we present **Spatial-Aware Graph Completion Network (SAGCNet)** to overcome the dependency on complete volumetric data, featuring two main innovations: (1) a volumetric slice graph completion module that incorporates the inter-slice relationships into a graph structure, and (2) a volumetric spatial adapter component that enables our model to effectively capture and utilize various forms of 3D spatial context. Extensive experiments on cardiac MRI datasets demonstrate that SAGCNet is capable of synthesizing absent CMR slices, outperforming competitive state-of-the-art MRI synthesis methods both quantitatively and qualitatively. Notably, our model maintains superior performance even with limited slice data. Code is available at <https://github.com/JK-Liu7/SAGCNet>.

**Keywords:** Medical image synthesis · Incomplete graph · Cardiac MRI


## 1 Introduction

Magnetic resonance imaging (MRI) provides significant insights into tissue and anatomical characteristics and is widely used in medical research and clinical diagnosis. Nonetheless, missing slice is a common issue for volumetric MRI data, particularly in cardiac magnetic resonance (CMR) volumes, within clinical applications and practice, caused by factors such as excessive scanning times, image deterioration, motion artifacts, and disparate acquisition techniques [19]. Consequently, the development of a unified and effective approach for imputing missing slices using available data is critically needed [23].

Missing data imputation is a general method for tackling the incomplete volumetric data problem, employing medical slice synthesis techniques to generate missing slices from available 3D CMR images [17] [20] [21]. Deep learning-based medical slice synthesis, using algorithm such as Convolutional neural network (CNN) [18] and transformer [9], has demonstrated notable progress and become an emerging research topic. Even though multi-modal MRI synthesis for missing modality imputation has made significant strides in recent years [24], these methods are not specifically devised for the missing slice imputation task, resulting in a lack of flexibility in handling arbitrary missing scenarios.

Below, we highlight the following two key challenges that need be addressed for missing slice imputation task: *(a) How to model inter-slice correlations and dependencies explicitly to capture local interactions between slices?* The inter-slice correlations focus on the interaction between adjacent slices along the through-plane in each volume, which is essential for learning discriminative and informative slice representations [4]. Additionally, the internal dependencies among different slices contribute to describing detailed anatomical structures and lesions, which should be carefully modeled to fully explore and extract high-level and hierarchical concepts [26]. Nevertheless, these inter-slice dependencies are often overlooked due to their complex interdependence characteristics, making them difficult to estimate and model. *(b) How to extract the crucial 3D spatial information in medical volumetric data to fully describe the global spatial context?* CMRs are inherently volumetric and three dimensional, making it challenging for existing methods to handle their depth-wise properties [12]. Moreover, directly applying simple 3D models to high-resolution volumes is infeasible due to the significant computational memory, and data acquisition [2]. Therefore, for 3D CMR synthesis, inherent isotropic 3D spatial information must be carefully and comprehensively considered to effectively learn 3D volumetric spatial patterns and capture structural and fine-grained details.

To address the aforementioned challenges, we propose **Spatial-Aware Graph Completion Network (SAGCNet)** for missing slice imputation. SAGCNet leverages graph structure to explicitly model inter-slice relationships and incorporates adapters to learn 3D spatial-related knowledge and characteristics within volumes. SAGCNet is intricately designed to balance computational cost and performance, achieving lightweight and efficient CMR synthesis with **moderate computational cost** and **rapid convergence rate**, as depicted in Fig. 1.



**Fig. 1.** (a) Performance comparison of SAGCNet and baselines under missing rate of 0.1, 0.4 and 0.7. (b) PSNR vs. floating point operations (FLOPs). The area of each blob denotes the number of model parameters. (c) Convergence rate of SAGCNet and baselines under missing rate of 0.1. All experiments are conducted on the UK Biobank dataset. More details are presented in the experiments section.

Our main contributions are as follows: **(1)** A unified missing slice imputation framework, SAGCNet, is proposed to effectively synthesize CMR images for arbitrary missing scenarios. **(2)** We propose a volumetric slice graph completion (VSGC) module, employing graphs to capture inter-slice correlations. To the best of our knowledge, our work is the first to leverage graph-based modeling at slice level for medical image synthesis. **(3)** We introduce a simple yet effective volumetric spatial adapter (VSA) to preserve 3D volumetric spatial information, enabling our model to extract crucial volumetric insights and thus be spatial-aware. **(4)** Experimental results on three datasets demonstrate the quantitative and qualitative superiority of SAGCNet under various missing rates.

## 2 Methodology

### 2.1 Problem Formulation and Model Overview

Given a volume with randomly missing slices, our aim is to construct a unified and robust framework to handle arbitrary missing slice scenarios, i.e., various missing positions and numbers, which simulates the practical clinical scenario. Mathematically, consider the given incomplete volume  $\mathbf{V} \in \mathbb{R}^{N \times H \times W}$  with available slices  $\mathbf{V}_a \in \mathbb{R}^{M \times H \times W}$  and missing slices  $\mathbf{V}_c \in \mathbb{R}^{P \times H \times W}$ , where  $M$  and  $P$  denote the numbers of available and missing slices, respectively, and  $N = M + P$ . SAGCNet reasons about the missing slice position automatically and synthesizes the complete volume as output according to the input available slices at once. The overall architecture of SAGCNet is depicted in Fig. 2. SAGCNet utilizes modified UNETR [6] with VSA as visual encoder to extract hierarchical features. Moreover, VSGC blocks are incorporated into the backbone encoder to fully exploit inter-slice relationships from multi-views. Ultimately, following the U-shaped network, the CNN-based decoders with skip connections are employed to reconstruct and synthesize images. Next, we will elaborate the detailed information of VSA and VSGC modules.

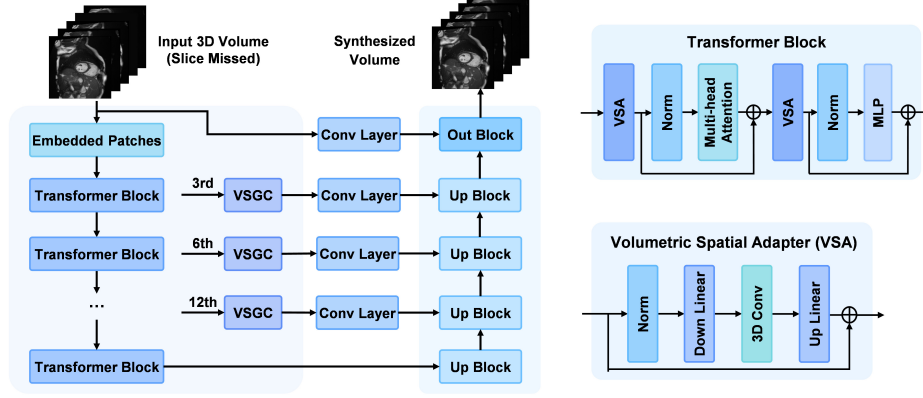


Fig. 2. The overview of our proposed SAGCNet.

## 2.2 Volumetric Spatial Adapter Module

To bridge the gap between 2D images and volumetric medical data, we devise a series of VSA modules integrated into each transformer block, empowering SAGCNet to learn spatial information inherent in 3D volumetric medical data [11]. Specifically, as illustrated in Fig. 2, each VSA can be represented as

$$VSA(\mathbf{H}) = \mathbf{H} + \sigma(\text{Conv 3D}(\text{Norm}(\mathbf{H}) \cdot \mathbf{W}_{\text{down}})) \mathbf{W}_{\text{up}} \quad (1)$$

where  $\mathbf{H}$  denotes the original feature representation,  $\mathbf{W}_{\text{down}}$  and  $\mathbf{W}_{\text{up}}$  represent the down- and up-projection layer respectively, Conv 3D indicates the 3D depth-wise convolutional layer. The 3D convolutional layer serves as a core module, aiming to extract valuable volumetric information. The down-projection layer reduces the original dimensionality, thereby lowering the number of parameters and making the VSA more flexible and lightweight. For each Transformer block, two VSAs are plugged before and after the multi-head self-attention module to produce better performance in practice.

## 2.3 Volumetric Slice Graph Completion

Leveraging the strong representation learning capability of Graph Neural Networks (GNNs) in complex relation modeling, graphs can effectively depict the spatial correlations of inter-slice properties through the message propagation pipeline [10, 13]. To this end, we propose to model volume data as graphs and establish a multi-view graph completion network to discover node interactions within each volume and extract inter-slice dependencies.

**Graph Construction.** After obtaining feature representations from the transformer blocks, we first employ a channel adapter to construct the node features, as shown in Fig. 3. Next, the k-Nearest Neighbors (kNN) algorithm is used to construct the incomplete volume slice graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ , where  $\mathcal{V}$  and  $\mathcal{E}$



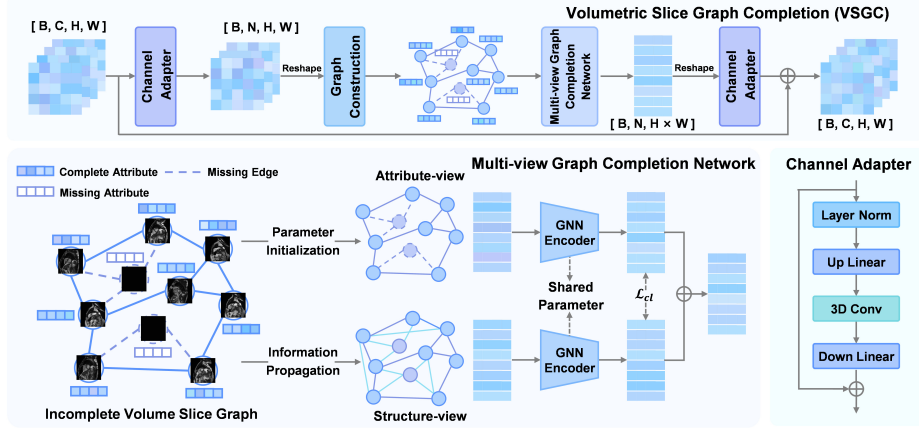


Fig. 3. An illustration of the VSGC module.

represent the node (slice) and edge (inter-slice correlation) set, respectively. Assume the slice number of a given volume, (i.e., node number) is  $n$ . Let  $\mathbf{X} \in \mathbb{R}^{n \times d}$  and  $\mathbf{A} \in \mathbb{R}^{n \times n}$  denote the original node attribute and adjacency matrix, respectively. Notably, nodes corresponding to missing slices are treated as nodes with missing attributes, and their associated edges are also considered incomplete.

**Multi-view Graph Completion Network.** Drawing inspiration from previous work [7], we devise a multi-view graph completion network to impute missing graph data at both attribute and structure levels. The attribute-view learning process, with parameterized completion, focus on extracting essential information from incomplete attributes while minimizing the impact of the incomplete graph structure on learning. Meanwhile, the structure-view learning process is designed to effectively model the incomplete graph structure without being hindered by missing node attributes.

Specifically, for attribute-view imputation, we use learnable neural network parameters to initialize node attributes, reducing noise in the original attribute matrix through iterative refinement. This strategy generates the augmented graph  $\mathcal{G}^a = (\mathbf{X}^a, \mathbf{A})$ . Regarding structure-view completion, we utilize personalized PageRank [5] to propagate information and enhance its diversity, producing  $\mathcal{G}^s = (\mathbf{X}^s, \mathbf{A}^s)$ , which is crucial for learning comprehensive and discriminative node representations under conditions where partial edges are missing. We choose Graph attention network [16] (GAT) as the GNN encoder due to its expressive power, as follows:

$$\mathbf{Z}^a = f_{GAT}(\mathcal{G}^a), \mathbf{Z}^s = f_{GAT}(\mathcal{G}^s), \quad (2)$$

Afterwards, to bridge the semantic gap between different views resulting from graph incompleteness and to achieve consistent representations, the contrastive learning paradigm is applied to maximum the mutual information:

$$\mathcal{L}_{cl}^v = -\frac{1}{2n} \sum_{i=1}^n \left( \log \frac{\varphi(\mathbf{z}_i^a, \mathbf{z}_i^s)}{\sum_{j \neq i}^n \varphi(\mathbf{z}_i^a, \mathbf{z}_j^s)} + \log \frac{\varphi(\mathbf{z}_i^s, \mathbf{z}_i^a)}{\sum_{j \neq i}^n \varphi(\mathbf{z}_i^s, \mathbf{z}_j^a)} \right), \quad (3)$$

where  $\varphi(a, b) = e^{\text{sim}(a, b)/\tau}$ ,  $\text{sim}(\cdot, \cdot)$  is the cosine similarity function,  $\tau$  is the temperature parameter. Finally, the fused embeddings from the two views are passed through another channel adapter with a residue connection to transform dimensionality for subsequent operation. The VSGC components are inserted after the 3rd, 6th, and 12th transformer blocks, respectively.

**Loss Function.** The overall loss function of SAGCNet is defined as

$$\mathcal{L} = \lambda_r \mathcal{L}_{rec} + \lambda_s \mathcal{L}_{syn} + \lambda_{cl} \mathcal{L}_{cl} \quad (4)$$

where  $\mathcal{L}_{rec}$ ,  $\mathcal{L}_{syn}$ , and  $\mathcal{L}_{cl}$  denote the reconstruction loss, synthesis loss and contrastive loss, respectively, with  $\lambda_r$ ,  $\lambda_s$ , and  $\lambda_{cl}$  are trade-off parameters for balance.  $\mathcal{L}_{rec}$ ,  $\mathcal{L}_{syn}$  consist of L1 loss and perceptual loss.

### 3 Experiments

**Datasets.** We validate the performance of the proposed SAGCNet on three cardiac MRI datasets: UK Biobank (UKBB) [14], Multi-Ethnic Study of Atherosclerosis (MESA) [1] and Automatic Cardiac Diagnosis Challenge (ACDC) [22]. The UKBB dataset we utilized contains CMR scans from 600 patients, each with 50 temporal phases. The MESA dataset consists of 299 CMR volumes, with the number of slices per volume ranging from 6 to 14. The ACDC dataset is a publicly available resource containing CMR images from 100 patients, with 6 to 21 slices per volume. We divide all datasets into training and test sets in the ratio of 8:2. During preprocessing, each 2D slice image is resized to  $256 \times 256$ . Zero-padding is applied for all volumes in the through-plane direction to ensure the fixed input size. Min-max normalization is employed to scale the intensity range of all images to  $[-1, 1]$ . To enhance data diversity, random flipping and rotation are applied as data augmentation.

**Implementation Details.** The missing rate  $\eta$  is defined as  $\eta = P/N$ , which is fixed during training and inference stages. For training phase, the missing slices are randomly sampled based on  $\eta$  at each iteration to ensure that our model is robust to arbitrary missing scenario. For model hyperparameters, we set the number of kNN neighbors to 3 and the number of GAT layers to 2. The hyperparameter  $\tau$  is set empirically to 0.8. Additionally,  $\lambda_r$ ,  $\lambda_s$ , and  $\lambda_{cl}$  are set to 5, 20, and 0.001, respectively. All experiments are implemented on an Nvidia A100 GPU. We train SAGCNet for 2000 epochs using the Adam optimizer with an initial learning rate of  $1e-4$ . The cosine annealing scheduler is applied to decay the learning rate to  $5e-6$ , and the training batch size is set to 8.

**Baselines and Evaluation Metrics.** We compare the proposed SAGCNet with several state-of-the-art (SOTA) image synthesis methods as baselines, including Pix2Pix [8], MMGAN [15], ResViT [3], UNETR [6], and MTT-Net [25]. To quantitatively evaluate the performance of SAGCNet and baselines, three commonly-used evaluation metrics are adopted, including mean absolute error (MAE), peak signal-to-noise ratio (PSNR), and structural similarity (SSIM).

**Quantitative Results.** Table 1 presents the quantitative results of our proposed SAGCNet and baseline models on three datasets under the missing rate

$\eta=0.1$ , 0.4 and 0.7. The experimental results demonstrate that our SAGCNet consistently outperforms SOTA baselines across various missing rate settings, highlighting its superiority and effectiveness in synthesizing high-fidelity CMR images. Notably, SAGCNet shows greater performance gains under low missing rate settings (e.g.,  $\eta=0.1$ ), with improvements ranging from 1.02% to 4.37% in PSNR. Furthermore, even in high missing rate scenarios (e.g.,  $\eta=0.4$  and 0.7), SAGCNet also surpasses competitive baselines across three datasets, underscoring its robustness and adaptability to real-world slice missing patterns.


**Table 1.** Quantitative performance comparison of different methods on three datasets under three missing rate configurations. The best performance is in **bold**.

Dataset	Method	Missing Rate					
		0.1		0.4		0.7	
		PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$
UKBB	Pix2pix [8]	32.77	0.925	19.30	0.674	18.89	0.653
	MMGAN [15]	32.72	0.927	19.16	0.669	18.96	0.669
	ResViT [3]	34.74	0.948	19.29	0.677	18.90	0.655
	UNETR [6]	36.55	0.964	19.31	0.680	18.94	0.653
	MTT-Net [25]	35.24	0.951	19.16	0.669	18.82	0.650
	SAGCNet (Ours)	<b>38.15</b>	<b>0.973</b>	<b>19.39</b>	<b>0.688</b>	<b>19.03</b>	<b>0.673</b>
MESA	Pix2pix [8]	28.83	0.811	22.25	0.742	21.72	0.738
	MMGAN [15]	28.62	0.810	22.25	0.740	21.71	0.734
	ResViT [3]	29.03	0.815	21.80	0.738	21.81	0.739
	UNETR [6]	33.23	0.884	22.21	0.764	21.93	0.747
	MTT-Net [25]	28.96	0.808	22.29	0.746	21.87	0.740
	SAGCNet (Ours)	<b>33.57</b>	<b>0.888</b>	<b>22.51</b>	<b>0.770</b>	<b>21.95</b>	<b>0.762</b>
ACDC	Pix2Pix [8]	25.57	0.758	18.98	0.625	18.47	0.613
	MMGAN [15]	25.68	0.764	18.96	0.632	18.50	0.619
	ResViT [3]	25.09	0.751	18.91	0.623	18.38	0.610
	UNETR [6]	29.30	0.862	19.03	0.631	18.60	0.624
	MTT-Net [25]	25.16	0.756	19.02	0.620	18.55	0.611
	SAGCNet (Ours)	<b>30.19</b>	<b>0.889</b>	<b>19.12</b>	<b>0.637</b>	<b>18.63</b>	<b>0.627</b>

To further evaluate the performance, we also compare the quantitative results of SAGCNet with I2GAN [17], which is specifically designed for the single missing slice scenario (i.e.,  $P = 1$ ). As illustrated in Table 2, our SAGCNet yields better performance across three metrics. It can be noticed that I2GAN is limited to scenarios where a single slice is missing.

**Table 2.** Quantitative performance comparison of our SAGCNet and I2GAN [17] on two datasets under single missing slice configuration. The best performance is in **bold**.

Dataset	Method	MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
UKBB	I2GAN [17]	0.0257	26.64	0.873
	Ours	<b>0.0110</b>	<b>38.19</b>	<b>0.978</b>
ACDC	I2GAN [17]	0.0428	24.39	0.846
	Ours	<b>0.0304</b>	<b>30.86</b>	<b>0.895</b>



**Fig. 4.** Qualitative results of all methods on the UKBB dataset. Every two rows, from top to bottom, denote the experimental results and the error maps, respectively.

**Qualitative Results.** The qualitative results of representative examples from the UKBB dataset are shown in Fig. 4. It can be observed that baseline approaches tend to generate blurrier results that lack anatomical details and edges. In contrast, by modeling inter-slice dependencies and capturing global context information, our SAGCNet significantly reduces the red areas of the error maps of the synthesized images, particularly in anatomical structures and regions, indicating the capability of SAGCNet to produce results that are both visually realistic and closely aligned with the ground truth.

**Ablation Studies.** To verify the effectiveness of each component in SAGCNet, ablation studies are conducted by constructing variant models. As shown in Table 3, each module contributes differently to the model performance on the UKBB and ACDC dataset, highlighting the positive impact of key components of SAGCNet in the missing slice imputation task. Furthermore, the VSA and VSGC modules exhibit the most significant performance improvements, further demonstrating the necessity and rationale of capturing inter-slice correlations and extracting valuable 3D information for volumetric data.

## 4 Conclusion

In this paper, we introduce SAGCNet for unified missing slice synthesis in volumetric CMR data. SAGCNet flexibly handles arbitrary sets of missing slices, reducing excessive reliance on complete volumetric data in real-world scenarios. By integrating inter-slice dependencies, we propose volumetric slice graphs that

**Table 3.** Quantitative performance comparison of ablation variants.

Method	Dataset					
	MAE ↓	UKBB PSNR ↑	SSIM ↑	MAE ↓	ACDC PSNR ↑	SSIM ↑
w/o VSA	0.0130	37.05	0.967	0.0357	29.45	0.875
w/o VSGC	0.0117	37.94	0.971	0.0329	30.08	0.879
w/o attribute-view	0.0121	37.93	0.970	0.0331	29.98	0.875
w/o structure-view	0.0119	37.91	0.968	0.0334	30.01	0.869
w/o $\mathcal{L}_{cl}$	0.0116	38.03	0.971	0.0328	30.12	0.881
SAGCNet	<b>0.0113</b>	<b>38.15</b>	<b>0.973</b>	<b>0.0326</b>	<b>30.19</b>	<b>0.889</b>

employ the graph structure to model CMR slices and impute absent data in a graph completion manner. To learn spatial-aware representations, we introduce the volumetric spatial adapter, empowering our model to adaptively exploit 3D spatial knowledge. Our extensive experiments demonstrate the superior effectiveness and robustness of SAGCNet compared with SOTA methods in both sufficient and limited slice data scenarios.

## Disclosure of Interests

The authors have no competing interests in the paper.

## References

1. Chen, X., Wang, R., Zee, P., Lutsey, P.L., Javaheri, S., Alcántara, C., Jackson, C.L., Williams, M.A., Redline, S.: Racial/ethnic differences in sleep disturbances: The multi-ethnic study of atherosclerosis (mesa). *Sleep* **38**(6), 877–888 (2015)
2. Choo, K., Jun, Y., Yun, M., Hwang, S.J.: Slice-consistent 3d volumetric brain ct-to-mri translation with 2d brownian bridge diffusion model. In: MICCAI. pp. 657–667 (2024)
3. Dalmaz, O., Yurt, M., Çukur, T.: Resvit: Residual vision transformers for multi-modal medical image synthesis. *IEEE Trans. Med. Imag.* **41**(10), 2598–2614 (2022)
4. Fang, C., Wang, L., Zhang, D., Xu, J., Yuan, Y., Han, J.: Incremental cross-view mutual distillation for self-supervised medical ct synthesis. In: CVPR. pp. 20677–20686 (2022)
5. Gasteiger, J., Bojchevski, A., Günnemann, S.: Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997* (2018)
6. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: WACV. pp. 574–584 (2022)
7. Huo, C., Jin, D., Li, Y., He, D., Yang, Y.B., Wu, L.: T2-gnn: Graph neural networks for graphs with incomplete features and structure via teacher-student distillation. *AAAI* **37**(4), 4339–4346 (2023)
8. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
9. Liu, J., Pasumarthi, S., Duffy, B., Gong, E., Datta, K., Zaharchuk, G.: One model to synthesize them all: Multi-contrast multi-scale transformer for missing data imputation. *IEEE Trans. Med. Imag.* **42**(9), 2577–2591 (2023)

10. Liu, Y., Ding, K., Wang, J., Lee, V., Liu, H., Pan, S.: Learning strong graph neural networks with weak information. In: KDD. p. 1559–1571. KDD '23 (2023)
11. Pan, J., Lin, Z., Zhu, X., Shao, J., Li, H.: St-adapter: Parameter-efficient image-to-video transfer learning. In: Proc. Adv. Neural Inform. Process. Syst. vol. 35, pp. 26462–26477 (2022)
12. Peng, C., Lin, W.A., Liao, H., Chellappa, R., Zhou, S.K.: Saint: Spatially aware interpolation network for medical slice synthesis. In: CVPR (2020)
13. Peng, X., Cheng, J., Tang, X., Zhang, B., Tu, W.: Multi-view graph imputation network. *Inf. Fusion* **102**, 102024 (2024)
14. Petersen, S.E., Matthews, P.M., Francis, J.M., Robson, M.D., Zemrak, F., Boubertakh, R., Young, A.A., Hudson, S., Weale, P., Garratt, S., Collins, R., Piechnik, S., Neubauer, S.: Uk biobank’s cardiovascular magnetic resonance protocol. *J. Cardio. Magn. Reson.* **18**(1), 8 (2016)
15. Sharma, A., Hamarneh, G.: Missing mri pulse sequence synthesis using multi-modal generative adversarial network. *IEEE Trans. Med. Imag.* **39**(4), 1170–1183 (2020)
16. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
17. Xia, Y., Zhang, L., Ravikumar, N., Attar, R., Piechnik, S.K., Neubauer, S., Petersen, S.E., Frangi, A.F.: Recovering from missing data in population imaging – cardiac mr image imputation via conditional generative adversarial nets. *Med. Image Analys.* **67**, 101812 (2021)
18. Yurt, M., Dar, S.U., Erdem, A., Erdem, E., Oguz, K.K., Çukur, T.: mustgan: multi-stream generative adversarial networks for mr image synthesis. *Med. Image Analys.* **70**, 101944 (2021)
19. Zhang, L., Bronik, K., Piechnik, S.K., Lima, J.A., Neubauer, S., Petersen, S.E., Frangi, A.F.: Automatic plane pose estimation for cardiac left ventricle coverage estimation via deep adversarial regression network. *IEEE Transactions on Artificial Intelligence* (2024)
20. Zhang, L., Pereañez, M., Bowles, C., Piechnik, S., Neubauer, S., Petersen, S., Frangi, A.: Missing slice imputation in population cmr imaging via conditional generative adversarial nets. In: MICCAI. pp. 651–659 (2019)
21. Zhang, L., Pereañez, M., Bowles, C., Piechnik, S.K., Neubauer, S., Petersen, S.E., Frangi, A.F.: Unsupervised standard plane synthesis in population cine mri via cycle-consistent adversarial networks. In: MICCAI. pp. 660–668 (2019)
22. Zhang, L., Wu, F., Bronik, K., Papiez, B.W.: Diffuseg: Domain-driven diffusion for medical image segmentation. *IEEE J. Biomed. Health Inform.* (2025)
23. Zhang, T., Tan, T., Han, L., Wang, X., Gao, Y., van Dijk, J., Portaluri, A., Gonzalez-Huete, A., D’Angelo, A., Lu, C., Teuwen, J., Beets-Tan, R., Sun, Y., Mann, R.: Important-net: Integrated mri multi-parametric increment fusion generator with attention network for synthesizing absent data. *Inf. Fusion* **108**, 102381 (2024)
24. Zhang, Y., Peng, C., Wang, Q., Song, D., Li, K., Kevin Zhou, S.: Unified multi-modal image synthesis for missing modality imputation. *IEEE Trans. Med. Imag.* **44**(1), 4–18 (2025)
25. Zhong, L., Chen, Z., Shu, H., Zheng, K., Li, Y., Chen, W., Wu, Y., Ma, J., Feng, Q., Yang, W.: Multi-scale tokens-aware transformer network for multi-region and multi-sequence mr-to-ct synthesis in a single model. *IEEE Trans. Med. Imag.* **43**(2), 794–806 (2024)
26. Zhou, H.Y., Guo, J., Zhang, Y., Han, X., Yu, L., Wang, L., Yu, Y.: nnformer: Volumetric medical image segmentation via a 3d transformer. *IEEE Trans. Image Process.* **32**, 4036–4045 (2023)