



Published in final edited form as:

Shape Med Imaging (2023). 2023 October ; 14350: 287–300. doi:10.1007/978-3-031-46914-5_23.

SADIR: Shape-Aware Diffusion Models for 3D Image Reconstruction

Nivetha Jayakumar¹, Tonmoy Hossain², Miaomiao Zhang^{1,2}

¹Department of Electrical and Computer Engineering, School of Engineering and Applied Science, University of Virginia, VA, USA

²Department of Computer Science, School of Engineering and Applied Science, University of Virginia, VA, USA

Abstract

3D image reconstruction from a limited number of 2D images has been a long-standing challenge in computer vision and image analysis. While deep learning-based approaches have achieved impressive performance in this area, existing deep networks often fail to effectively utilize the shape structures of objects presented in images. As a result, the topology of reconstructed objects may not be well preserved, leading to the presence of artifacts such as discontinuities, holes, or mismatched connections between different parts. In this paper, we propose a shape-aware network based on diffusion models for 3D image reconstruction, named SADIR, to address these issues. In contrast to previous methods that primarily rely on spatial correlations of image intensities for 3D reconstruction, our model leverages shape priors learned from the training data to guide the reconstruction process. To achieve this, we develop a joint learning network that simultaneously learns a mean shape under deformation models. Each reconstructed image is then considered as a deformed variant of the mean shape. We validate our model, SADIR, on both brain and cardiac magnetic resonance images (MRIs). Experimental results show that our method outperforms the baselines with lower reconstruction error and better preservation of the shape structure of objects within the images.

1 Introduction

The reconstruction of 3D images from a limited number of 2D images is fundamental to various applications, including object recognition and tracking [12], robot navigation [44], and statistical shape analysis for disease detection [4,36]. However, inferring the complete 3D geometry and structure of objects from one or multiple 2D images has been a long-standing ill-posed problem [25]. A bountiful literature has been investigated to recover the data from a missing dimension [9,32,34,37]. Initial approaches to address this challenge focused on solving an inverse problem of projecting 3D information onto 2D images from geometric aspects [8]. These solutions typically require images captured from different viewing angles using precisely calibrated cameras or medical imaging machines [7,28]. In spite of producing a good quality of 3D reconstructions, such methods are often impractical or infeasible in many real-world scenarios.

Recent advancements have leveraged deep learning (DL) techniques to overcome the limitations posed in previous methods [5,15,27]. Extensive research has explored various network architectures for 3D image reconstruction, including UNets [30], transformers [14,22], and state-of-the-art generative diffusion models [37]. These works have significantly improved the reconstruction efficiency by learning intricate mappings between stacks of 2D images and their corresponding 3D volumes. While the DL-based approaches have achieved impressive results in reconstructing detailed 3D images, they often lack explicit consideration of shape information during the learning process. Consequently, important geometric structures of objects depicted in the images may not be well preserved. This may lead to the occurrence of artifacts, such as discontinuities, holes, or mismatched connections between different parts, that break the topology of the reconstructed objects.

Motivated by recent studies highlighting the significance of shape in enhancing image analysis tasks using deep networks [6,20,26,39,43], we introduce a novel shape-aware 3D image reconstruction network called SADIR. Our methodology builds upon the foundation of diffusion models while incorporating shape learning as a key component. In contrast to previous methods that mainly rely on spatial correlations of image intensities for 3D reconstruction, our SADIR explicitly incorporates the geometric shape information aiming to preserve the topology of reconstructed images. To achieve this goal, we develop a joint deep network that simultaneously learns a shape prior (also known as a mean shape) from a given set of full 3D volumes. In particular, an atlas building network based on deformation models [39] is employed to learn a mean shape representing the average information of training images. With the assumption that each reconstructed object is a deformed variant of the estimated mean shape, we then utilize the mean shape as a prior knowledge to guide the diffusion process of reconstructing a complete 3D image from a stack of sparse 2D slices. To evaluate the effectiveness of our proposed approach, we conduct experiments on both real brain and cardiac magnetic resonance images (MRIs). The experimental results show the superiority of SADIR over the baseline approaches, as evidenced by substantially reduced reconstruction errors. Moreover, our method successfully preserves the topology of the images during the shape-aware 3D image reconstruction process.

2 Background: Fréchet Mean via Atlas Building

In this section, we briefly review an unbiased atlas building algorithm [21], a widely used technique to estimate the Fréchet mean of group-wise images. With the underlying assumption that objects in many generic classes can be described as deformed versions of an ideal template, descriptors in this class arise naturally by matching the mean (also referred as atlas) to an input image [21,38,45,42,46]. The resulting transformation is then considered as a shape that reflects geometric changes.

Given a number of N images $\{\mathcal{Y}_1, \dots, \mathcal{Y}_N\}$, the problem of atlas building is to find a mean or template image \mathcal{S} and deformation fields ϕ_1, \dots, ϕ_N with derived initial velocity fields v_1, \dots, v_N that minimize the energy function

$$E(\mathcal{S}, \phi_n) = \sum_{n=1}^N \frac{1}{\sigma^2} \text{Dist}[\mathcal{S} \circ \phi_n(v_i), \mathcal{Y}_n] + \text{Reg}[\phi_n(v_i)], \quad (1)$$

where σ^2 is a noise variance and \circ denotes an interpolation operator that deforms image \mathcal{Y}_n with an estimated transformation ϕ_n . The $\text{Dist}[\cdot, \cdot]$ is a distance function that measures the dissimilarity between images, i.e., sum-of-squared differences [3], normalized cross correlation [2], and mutual information [40]. The $\text{Reg}[\cdot]$ is a regularizer that guarantees the smoothness of transformations.

Given an open and bounded d -dimensional domain $\Omega \subset \mathbb{R}^d$, we use $\text{Diff}(\Omega)$ to denote a space of diffeomorphisms (i.e., a one-to-one smooth and invertible smooth transformation) and its tangent space $V = T\text{Diff}(\Omega)$. A well-developed algorithm, large deformation diffeomorphic metric mapping (LDDMM) [3], provides a regularization that guarantees the smoothness of deformation fields and preserves the topological structures of objects for the atlas building framework (Eq. (1)). Such a regularization is formulated as an integral of the Sobolev norm of the time-dependent velocity field $v_n(t) \in V(t \in [0,1])$ in the tangent space, i.e.,

$$\text{Reg}[\phi_n(v_i)] = \int_0^1 (Lv_i, v_i) dt, \quad \text{with} \quad \frac{d\phi_n(t)}{dt} = -D\phi_n(t) \cdot v_n(t), \quad (2)$$

where $L: V \rightarrow V^*$ is a symmetric, positive-definite differential operator that maps a tangent vector $v_i \in V$ into its dual space as a momentum vector $m_i \in V^*$. We write $m_i = Lv_i$, or $v_i = Km_i$, with K being an inverse operator of L . The operator D denotes a Jacobian matrix and \cdot represents element-wise matrix multiplication. In this paper, we use a metric of the form $L = (-\alpha\Delta + \gamma\mathbf{I})^3$, in which Δ is the discrete Laplacian operator, α is a positive regularity parameter that controls the smoothness of transformation fields, γ is a weighting parameter, and \mathbf{I} denotes an identity matrix.

The minimum of Eq. (2) is uniquely determined by solving an Euler-Poincaré differential equation (EPDiff) [1,29] with a given initial condition of velocity fields, noted as v_0 . This is known as the *geodesic shooting* algorithm [35], which nicely proves that the deformation-based shape descriptor ϕ_n can be fully characterized by an initial velocity field $v_n(0)$. The mathematical formulation of the EPDiff equation is

$$\frac{\partial v_n(t)}{\partial t} = -K \left[(Dv_n(t))^T \cdot m_n(t) + Dm_n(t) \cdot v_n(t) + m_n(t) \cdot \text{div} v_n(t) \right], \quad (3)$$

where the operator D denotes a Jacobian matrix, div is the divergence, and \cdot represents element-wise matrix multiplication.

We are now able to equivalently minimize the atlas building energy function in Eq. (1) as

$$E(\mathcal{S}, \phi_n) = \sum_{n=1}^N \frac{1}{\sigma^2} \text{Dist}[\mathcal{S} \circ \phi_n(v_n(t)), \mathcal{Y}_n] + (Lv_n(0), v_n(0)), \text{ s.t. Eq. (2) \& (3)}. \quad (4)$$

For notation simplicity, we will drop the time index in the following sections.

3 Our Method: SADIR

In this section, we present SADIR, a novel reconstruction network that incorporates shape information in predicting 3D volumes from a limited number of input 2D images. We introduce a sub-module of the atlas building framework, which enables us to learn shape priors from a given set of full 3D images. It is worth mentioning that while the backbone of our proposed SADIR is a diffusion model [16], the methodology can be generalized to a variety of network architectures such as UNet [33], UNet++ [47], and Transformer [11].

3.1 Shape-Aware Diffusion Models Based on Atlas Building Network

Given a number of N training data $\{I_n, \mathcal{Y}_n\}_{n=1}^N$, where I_n is a stack of sparse 2D images with its associated full 3D volume \mathcal{Y}_n . Our model SADIR consists of two submodules:

- i. An atlas building network, parameterized by θ^a , that provides a mean image \mathcal{S} of $\{\mathcal{Y}_n\}$. In this paper, we employ the network architecture of Geo-SIC [39];
- ii. A reconstruction network, parameterized by θ^r , that considers each reconstructed image $\hat{\mathcal{Y}}_n$ as a deformed variant of the obtained atlas, i.e., $\hat{\mathcal{Y}}_n \triangleq \mathcal{S} \circ \phi_n(v_n(\theta^r))$. In contrast to current approaches learning the reconstruction process based on image intensities, our model is developed to learn the geometric shape variations represented by the predicted velocity field v_n .

Next, we introduce the details of our shape-aware diffusion models for reconstruction, which is a key component of SADIR. Similar to existing diffusion models [16,37], we develop a forward diffusion and a reverse diffusion process to predict the velocity fields associated with the pair of input training images and an atlas image. For the purpose of simplified math notations, we omit the index n for each subject in the following sections.

Forward diffusion process.—Let y^0 denote the original 3D image with full volumes and τ denote the time point of the diffusion process. We assume the data distribution of y^τ is a normal distribution with mean μ and variance β , i.e., $y^\tau \sim \mathcal{N}(\mu, \beta)$. The forward diffusion of $y^{\tau-1}$ to y^τ is then recursively given by

$$p(y^\tau | y^{\tau-1}) = \mathcal{N}\left(y^\tau; \sqrt{1 - \beta^\tau} y^{\tau-1}, \beta^\tau \mathbf{I}\right), \quad (5)$$

where \mathbf{I} denotes an identity matrix, and $\beta^\tau \in [0,1]$ denotes a known variance increased along the time steps with $\beta^1 < \beta^2 < \dots < \beta^\tau$. The forward diffusion process is repeated for a fixed, predefined number of time steps.

It is shown in [16] that repeated application of Eq. (5) to the original image y^0 and setting $\alpha^\tau = 1 - \beta^\tau$ and $\bar{\alpha}^\tau = \prod_{i=1}^{\tau} \alpha^i$ yields

$$p(y^\tau | y^0) = \mathcal{N}\left(y^\tau; \sqrt{\bar{\alpha}^\tau} y^0, (1 - \bar{\alpha}^\tau) \mathbf{I}\right).$$

Therefore, we can write y^τ in terms of y^0 as

$$y^\tau = \sqrt{\bar{\alpha}^\tau} y^0 + \sqrt{1 - \bar{\alpha}^\tau} \epsilon \quad \text{with } \epsilon \sim \mathcal{N}(0, \mathbf{I}).$$

Reverse diffusion process.—Given a concatenation of a sparse stack of 2D images I , an atlas image \mathcal{S} , and y^τ from the forward process, our diffusion model is designed to remove the added noise in the reverse process. Following the work of [41], we will now predict $y^{\tau-1}$ from the input y^τ . The joint probability distribution $p(y^{\tau-1} | y^\tau)$ is predicted by a trained neural network (e.g., UNet) in each reverse time step for all $\tau \in \{1, \dots, T\}$, where T is the maximal time step. With the network model parameters denoted by θ' , we can write the reverse process as

$$p_{\theta'}(y^{\tau-1} | y^\tau) = \mathcal{N}(y^{\tau-1}; \mu_{\theta'}(y^\tau, \tau), \Sigma_{\theta'}(y^\tau, \tau)).$$

Similarly, we can write $y^{\tau-1}$ backward in terms of y^τ as

$$y^{\tau-1} = \frac{1}{\sqrt{\alpha^\tau}} \left(y^\tau \frac{1 - \alpha^\tau}{\sqrt{1 - \bar{\alpha}^\tau}} \epsilon_{\theta'}(y^\tau, \tau) \right) + \sigma^\tau \mathbf{z},$$

where σ^τ is the variance scheme the model can learn, the component \mathbf{z} is a stochastic sampling process. The model is trained with input y^τ to subtract the noise scheme $\epsilon_{\theta'}(y^\tau, \tau)$ from y^τ to produce $y^{\tau-1}$.

The output of this reverse process is a predicted velocity field $v(\theta^r)$, which is then used to generate its associated transformation $\phi(v(\theta^r))$ to deform the atlas \mathcal{S} . Such a deformed atlas is the reconstructed image $\hat{\mathcal{Y}} = \mathcal{S} \circ \phi(v(\theta^r))$.

An overview of the proposed SADIR network architecture is shown in Fig. 1.

3.2 Network Loss and Optimization

The network loss function of our model, SADIR, is a joint loss of the atlas building network and the diffusion reconstruction network. We first define the atlas building loss as

$$\mathcal{L}(\theta^a) = \sum_{n=1}^N \frac{1}{\sigma^2} \|\mathcal{S}(\theta^a) \circ (\phi_n(v_n)) - \mathcal{Y}_n\|_2^2 + (Lv_n, v_n) + \text{reg}(\theta^a), \quad (6)$$

where $\text{reg}(\cdot)$ denotes a regularization on the network parameters.

We then define the loss function of the diffusion reconstruction network as a combination of sum-of-squared differences and Sørensen—Dice coefficient [10] loss (for distinct anatomical structure, e.g., brain ventricles or myocardium) between the predicted reconstruction and ground-truth in following

$$\mathcal{L}(\theta^r) = \sum_{n=1}^N \|\mathcal{S} \circ \phi_n(v_n(\theta^r)) - \mathcal{Y}_n\|_2^2 + \eta[1 - \text{Dice}(\mathcal{S} \circ \phi_n(v_n(\theta^r)), \mathcal{Y}_n)] + \text{reg}(\theta^r), \quad (7)$$

where η is the weighting parameter, and $\text{Dice}(\hat{\mathcal{Y}}, \mathcal{Y}_n) = 2(|\hat{\mathcal{Y}}| \cap |\mathcal{Y}_n|) / (|\hat{\mathcal{Y}}| + |\mathcal{Y}_n|)$, considering $\hat{\mathcal{Y}}_n \triangleq \mathcal{S} \circ \phi_n(v_n(\theta^r))$. Defining λ as a weighting parameter, we are now ready to write the joint loss of SADIR as

$$\mathcal{L} = \mathcal{L}(\theta^a) + \lambda \mathcal{L}(\theta^r).$$

Joint network learning with an alternative optimization.—We use an alternative optimization scheme [31] to minimize the total loss \mathcal{L} in Eq. (3.2). More specifically, we jointly optimize all network parameters by alternating between the training of the atlas building and diffusion reconstruction network, making it end-to-end learning. A summary of our joint training of SADIR is presented in Alg. 1.

Algorithm 1: Joint Training of SADIR.

Input : A group of N input images with full 3D volumes $\{\mathcal{Y}_n\}$ and a stack of sparse 2D images $\{I_n\}$.

Output: Generate mean shape or atlas \mathcal{S} , initial velocity fields v_n , and reconstructed images $\hat{\mathcal{Y}}_n$

```

1 for  $i = 1$  to  $p$  do
2   /* Train geometric shape learning network */
3   Minimize the atlas building loss in Eq. (6)
4   Output the atlas  $\mathcal{S}$ 
5   /* Train diffusion network */
6   Minimize the diffusion reconstruction loss in Eq. (7)
7   Output the initial velocity fields  $\{v_n\}$  and the reconstructed images  $\hat{\mathcal{Y}}_n$ 
8 end
9 Until convergence

```

4 Experimental Evaluation

We demonstrate the effectiveness of our proposed model, SADIR, for 3D image reconstruction from 2D slices on both brain and cardiac MRI scans.

3D Brain MRIs: For 3D real brain MRI scans, we include 214 public T1-weighted longitudinal brain scans from the latest released Open Access Series of Imaging Studies (OASIS-III) [23]. All subjects include both healthy and disease individuals,

aged from 42 to 95. All MRIs were pre-processed as $256 \times 256 \times 256, 1.25\text{mm}^3$ isotropic voxels, and underwent skull-stripped, intensity normalized, bias field corrected and pre-aligned with affine transformation. To further validate the performance of our proposed model on specific anatomical shapes, we select left and right brain ventricles available in the OASIS-III dataset [23].

3D Cardiac MRIs: For 3D real cardiac MRI, we include 215 publicly available 3D myocardium mesh data from MedShapeNet dataset [24]. We convert the mesh data to binary label maps using 3D slicer [13]. All the images were pre-processed as $222 \times 222 \times 222$ and pre-aligned with affine transformation.

4.1 Experimental Settings

We first validate our proposed model, SADIR, on reconstructing 3D brain ventricles, as well as brain MRIs from a sparse stack of eight 2D slices. We compare our model's performance with three state-of-the-art deep learning-based reconstruction models: 3D-UNet [9]; DDPM, a probabilistic diffusion model [16]; and DISPR, a diffusion model based shape reconstruction model with geometric topology considered [37]. Three evaluation metrics, including the Sørensen-Dice coefficient (DSC) [10], Jaccard Similarity [19], and RHD95 score [18], are used to validate the prediction accuracy of brain ventricles for all methods. For brain MR images, we show the error maps of reconstructed images for all the experiments.

To further validate the performance of SADIR on different datasets, we run tests on a relatively small dataset of cardiac MRIs to reconstruct 3D myocardium.

Parameter setting: We set the mean and standard deviation of the forward diffusion process to be 0 and 0.1, respectively. The scheduling is `linear` for the noising process and is scaled to reach an isotropic Gaussian distribution irrespective of the value of T . For the atlas building network, we set the depth of the UNet architecture as 4. We set the number of time steps for Euler integration in EPDiff (Eq. (3)) as 10, and the noise variance $\sigma = 0.02$. For the shooting, we use a kernel map valued [0.5, 0, 1.0]. Besides, we set the parameter $\alpha = 3$ for the operator L . Similar to [37], we set the batch size as 1 for all experiments. We utilize the cosine annealing learning rate scheduler that starts with a learning rate of $\eta = 1e^{-3}$ for network training. We run all models on training and validation images using the Adam optimizer and save the networks with the best validation performance.

In the reverse process of the diffusion network, we set the depth of the 3D attention-UNet backbone as 6. We introduce the attention mechanism via spatial excitation channels [17], with ReLU (Rectified Linear Unit) activation. The UNet backbone has ELU activation (Exponential Linear Unit) in the hidden convolution layers and GeLU (Gaussian error Linear Unit) activation with tanh approximation. For each training experiment, we utilize Rivanna (high-performance computing servers of the University of Virginia) with NVIDIA A100 and V100 GPUs for ~18 hours (till convergence). For all the experimental datasets, we split all the training datasets into 70% training, 15% validation, and 15% testing. For both training and testing, we downsample all the image resolutions to $64 \times 64 \times 64$.

4.2 Experimental Results

Fig. 2 visualizes examples of ground truth and reconstructed 3D volumes of brain ventricles from all methods. It shows that SADIR outperforms all baselines in well preserving the structural information of the brain ventricles. In particular, models without considering the shape information of the images (i.e., 3D-UNet and DDPM) generate unrealistic shapes such as those with joint ventricles, holes in the volume, and deformed ventricle tails. While the other algorithm, DISPR, shows improved performance of enforcing topological consistency on the object surface, its predicted results of 3D volumes are inferior to SADIR.

Tab. 1 reports the average scores along with the standard deviation of the Dice similarity coefficient (DSC), Jaccard similarity, and Hausdorff distance computed between the brain ventricles reconstructed by all the models and the ground truth. Compared to all the baselines, SADIR achieves the best performance with a 1.6% – 5.6% increase in the average DSC with the lowest standard deviations across all metrics.

Fig. 3 visualizes the ground truth and reconstructed 3D brain MRIs as a result of evaluating DDPM and our method SADIR on the test data, along with their corresponding error maps. The error map is computed as absolute values of an element-wise subtraction between the ground truth and the reconstructed image. The images reconstructed by SADIR outperform the DDPM with a low absolute reconstruction error. Our method also preserves crucial anatomical features such as the shape of the ventricles, corpus callosum and gyri, which cannot be seen in the images reconstructed by the DDPM. This can be attributed to the lack of incorporating the shape information to guide the 3D MRI reconstruction. Moreover, our model has little to no noise in the background as compared to the DDPM.

Tab. 2 reports the average scores of DSC, Jaccard similarity, and Hausdorff distance evaluated between the reconstructed myocardium from all algorithms and the ground truth. Our method proves to be competent in reconstructing 3D volumes without discontinuities, artifacts, jagged edges or amplified structures, as can be seen in results from the other models. Compared to the baselines, SADIR achieves the best performance in terms of DSC, Jaccard similarity, and RHD95 with the lowest standard deviations across all metrics.

Fig. 4 visualizes a comparison of the reconstructed 3D myocardium between the ground truth and all models. It shows that our method consistently produces reconstructed volumes that preserve the original shape of the organ with less artifacts.

Fig. 5 shows examples of the superior, left, anterior and left-anterior views of the 3D ground truth and SADIR-reconstructed volumes of the myocardium for different subjects. We observe that the results predicted by SADIR have little to no difference from the ground truth, thereby efficiently preserving the anatomical structure of the myocardium.

5 Conclusion

This paper introduces a novel shape-aware image reconstruction framework based on diffusion model, named as SADIR. In contrast to previous approaches that mainly rely on the information of image intensities, our model SADIR incorporates shape features in the

deformation spaces to preserve the geometric structures of objects in the reconstruction process. To achieve this, we develop a joint deep network that simultaneously learns the underlying shape representations from the training images and utilize it as a prior knowledge to guide the reconstruction network. To the best of our knowledge, we are the first to consider deformable shape features into the diffusion model for the task of image reconstruction. Experimental results on both 3D brain and cardiac MRI show that our model efficiently produces 3D volumes from a limited number of 2D slices with substantially low reconstruction errors while better preserving the topological structures and shapes of the objects.

Acknowledgement.

This work was supported by NSF CAREER Grant 2239977 and NIH 1R21EB032597.

References

1. Arnold V. Sur la géométrie différentielle des groupes de lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. In *Annales de l'institut Fourier*, volume 16, pages 319–361, 1966.
2. Avants BB, Epstein CL, Grossman M, and Gee JC. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1):26–41, 2008. [PubMed: 17659998]
3. Beg MF, Miller MI, Trouné A, and Younes L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005.
4. Bruse JL, Zuluaga MA, Khushnood A, McLeod K, Ntsinjana HN, Hsia T-Y, Sermesant M, Pennec X, Taylor AM, and Schievano S. Detecting clinically meaningful shape clusters in medical image data: metrics analysis for hierarchical clustering applied to healthy and pathological aortic arches. *IEEE Transactions on Biomedical Engineering*, 64(10):2373–2383, 2017. [PubMed: 28221991]
5. Cetin I, Stephens M, Camara O, and Ballester MAG. Attri-vae: Attribute-based interpretable representations of medical images with variational autoencoders. *Computerized Medical Imaging and Graphics*, 104:102158, 2023. [PubMed: 36638626]
6. Chen C, Biffi C, Tarroni G, Petersen S, Bai W, and Rueckert D. Learning shape priors for robust cardiac mr segmentation from multi-view images. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 523–531. Springer, 2019.
7. Chen L, Bentley P, Mori K, Misawa K, Fujiwara M, and Rueckert D. Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis*, 58:101539, 2019. [PubMed: 31374449]
8. Chung H, Ryu D, McCann MT, Klasky ML, and Ye JC. Solving 3d inverse problems using pre-trained 2d diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22542–22551, 2023.
9. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, and Ronneberger O. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, pages 424–432. Springer, 2016.
10. Dice LR. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.
11. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, et al. An image is worth 16×16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

12. Duwek HC, Bitton A, and Tsur EE. 3d object tracking with neuromorphic event cameras via image reconstruction. In 2021 IEEE Biomedical Circuits and Systems Conference (BioCAS), pages 1–4. IEEE, 2021.
13. Fedorov A, Beichel R, Kalpathy-Cramer J, Finet J, Fillion-Robin J-C, Pujol S, Bauer C, Jennings D, Fennessy F, Sonka M, and et al. 3d slicer as an image computing platform for the quantitative imaging network, Nov 2012.
14. Feng C-M, Yan Y, Fu H, Chen L, and Xu Y. Task transformer network for joint mri reconstruction and super-resolution. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, pages 307–317. Springer, 2021.
15. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, and Bengio Y. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
16. Ho J, Jain A, and Abbeel P. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
17. Hu J, Shen L, Albanie S, Sun G, and Wu E. Squeeze-and-excitation networks, 2019.
18. Huttenlocher D, Klanderman G, and Rucklidge W. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993.
19. Jaccard P. Nouvelles recherches sur la distribution florale. *Bull. Soc. Vaud. Sci. Nat*, 44:223–270, 1908.
20. Jiang J and Veeraraghavan H. One shot pacs: Patient specific anatomic context and shape prior aware recurrent registration-segmentation of longitudinal thoracic cone beam cts. *IEEE Transactions on Medical Imaging*, 41(8):2021–2032, 2022.
21. Joshi S, Davis B, Jomier M, and Gerig G. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23:S151–S160, 2004. [PubMed: 15501084]
22. Korkmaz Y, Dar SU, Yurt M, Özbey M, and Cukur T. Unsupervised mri reconstruction via zero-shot learned adversarial transformers. *IEEE Transactions on Medical Imaging*, 41(7):1747–1763, 2022. [PubMed: 35085076]
23. LaMontagne PJ, Benzinger TL, Morris JC, Keefe S, Hornbeck R, Xiong C, Grant E, Hassenstab J, Moulder K, Vlassenko AG, et al. Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *MedRxiv*, 2019.
24. Li J. Medshapenet: A large-scale dataset of 3d medical shapes for computer vision, Mar 2023.
25. Lin DJ, Johnson PM, Knoll F, and Lui YW. Artificial intelligence for mr image reconstruction: an overview for clinicians. *Journal of Magnetic Resonance Imaging*, 53(4):1015–1028, 2021. [PubMed: 32048372]
26. Liu J, Aviles-Rivero AI, Ji H, and Schönlieb C-B. Rethinking medical image reconstruction via shape prior, going deeper and faster: Deep joint indirect registration and reconstruction. *Medical Image Analysis*, 68:101930, 2021. [PubMed: 33378731]
27. Maaløe L, Fraccaro M, Liévin V, and Winther O. Biva: A very deep hierarchy of latent variables for generative modeling. *Advances in neural information processing systems*, 32, 2019.
28. Maier-Hein L, Moutney P, Bartoli A, Elhawary H, Elson D, Groch A, Kolb A, Rodrigues M, Sorger J, Speidel S, et al. Optical techniques for 3d surface reconstruction in computer-assisted laparoscopic surgery. *Medical image analysis*, 17(8):974–996, 2013. [PubMed: 23837969]
29. Miller MI, Trouvé A, and Younes L. Geodesic shooting for computational anatomy. *Journal of mathematical imaging and vision*, 24(2):209–228, 2006. [PubMed: 20613972]
30. Nguyen T, Hua B-S, and Le N. 3d-ucaps: 3d capsules unet for volumetric image segmentation. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24, pages 548–558. Springer, 2021.
31. Nocedal J and Wright SJ. Numerical optimization. Springer, 1999.
32. Qin C, Schlemper J, Caballero J, Price AN, Hajnal JV, and Rueckert D. Convolutional recurrent neural networks for dynamic mr image reconstruction. *IEEE transactions on medical imaging*, 38(1):280–290, 2018. [PubMed: 30080145]
33. Ronneberger O, Fischer P, and Brox T. U-net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015:

- 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18, pages 234–241. Springer, 2015.
34. Schlemper J, Caballero J, Hajnal JV, Price AN, and Rueckert D. A deep cascade of convolutional neural networks for dynamic mr image reconstruction. *IEEE transactions on Medical Imaging*, 37(2):491–503, 2017. [PubMed: 29035212]
 35. Vialard F-X, Risser L, Rueckert D, and Cotter CJ. Diffeomorphic 3d image registration via geodesic shooting using an efficient adjoint calculation. *International Journal of Computer Vision*, 97(2):229–241, 2012.
 36. von Tycowicz C, Ambellan F, Mukhopadhyay A, and Zachow S. An efficient riemannian statistical shape model using differential coordinates: With application to the classification of data from the osteoarthritis initiative. *Medical image analysis*, 43:1–9, 2018. [PubMed: 28961450]
 37. Waibel DJE, Röell E, Rieck B, Giryas R, and Marr C. A diffusion model predicts 3d shapes from 2d microscopy images, 2023.
 38. Wang J and Zhang M. Bayesian atlas building with hierarchical priors for subject-specific regularization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 76–86. Springer, 2021.
 39. Wang J and Zhang M. Geo-sic: learning deformable geometric shapes in deep image classifiers. *Advances in Neural Information Processing Systems*, 35:27994–28007, 2022.
 40. Wells III WM, Viola P, Atsumi H, Nakajima S, and Kikinis R. Multi-modal volume registration by maximization of mutual information. *Medical image analysis*, 1(1):35–51, 1996. [PubMed: 9873920]
 41. Wolleb J, Bieder F, Sandkühler R, and Cattin PC. Diffusion models for medical anomaly detection. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VIII*, pages 35–45. Springer, 2022.
 42. Wu N, Wang J, Zhang M, Zhang G, Peng Y, and Shen C. Hybrid atlas building with deep registration priors. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022.
 43. Yang J, Wickramasinghe U, Ni B, and Fua P. Implicitatlas: learning deformable shape templates in medical imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15861–15871, 2022.
 44. Zelenskii A, Gapon N, Voronin V, Semenishchev E, Serebrenny V, and Cen Y. Robot navigation using modified slam procedure based on depth image reconstruction. In *Artificial Intelligence and Machine Learning in Defense Applications III*, volume 11870, pages 73–82. SPIE, 2021.
 45. Zhang M, Singh N, and Fletcher PT. Bayesian estimation of regularization and atlas building in diffeomorphic image registration. In *International conference on information processing in medical imaging*, pages 37–48. Springer, 2013.
 46. Zhang M, Wells WM, and Golland P. Low-dimensional statistics of anatomical variability via compact representation of image deformations. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part III 19*, pages 166–173. Springer, 2016.
 47. Zhou Z, Rahman Siddiquee MM, Tajbakhsh N, and Liang J. Unet++: A nested u-net architecture for medical image segmentation. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 3–11. Springer, 2018.

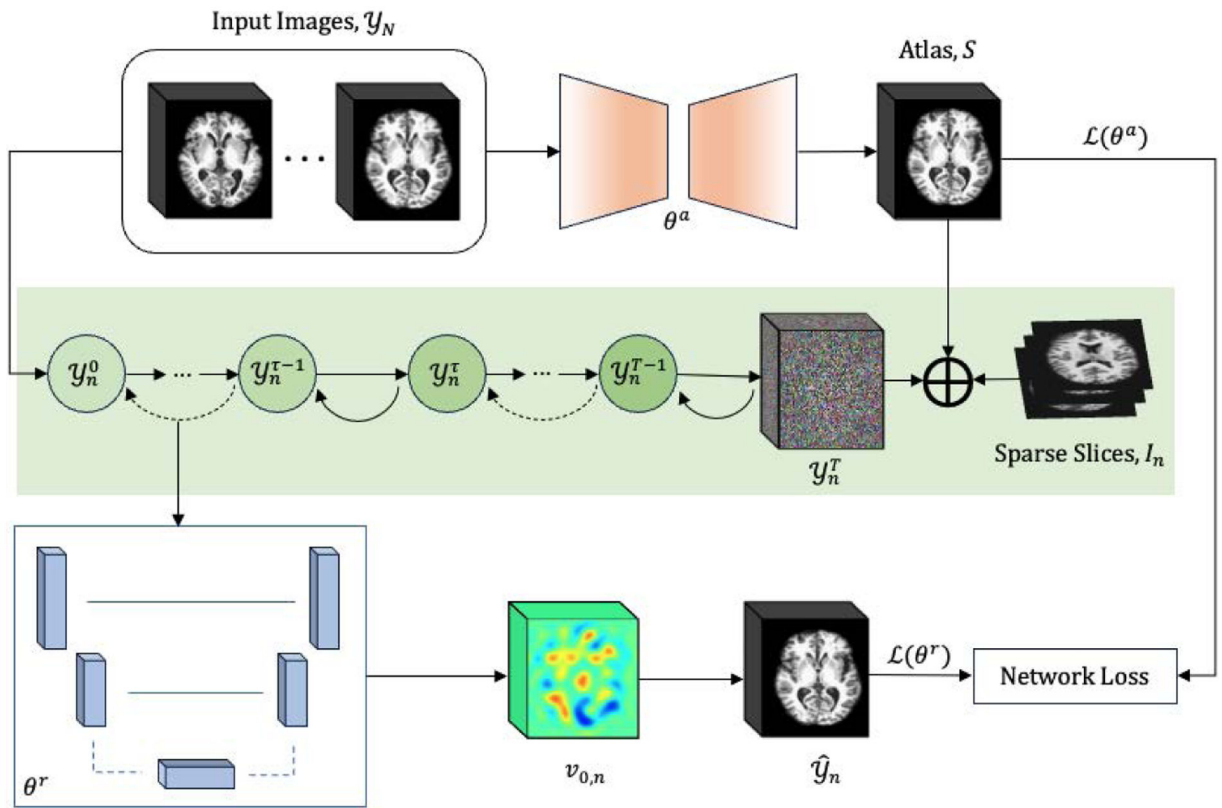


Fig. 1.
An overview of our proposed 3D reconstruction model SADIR.

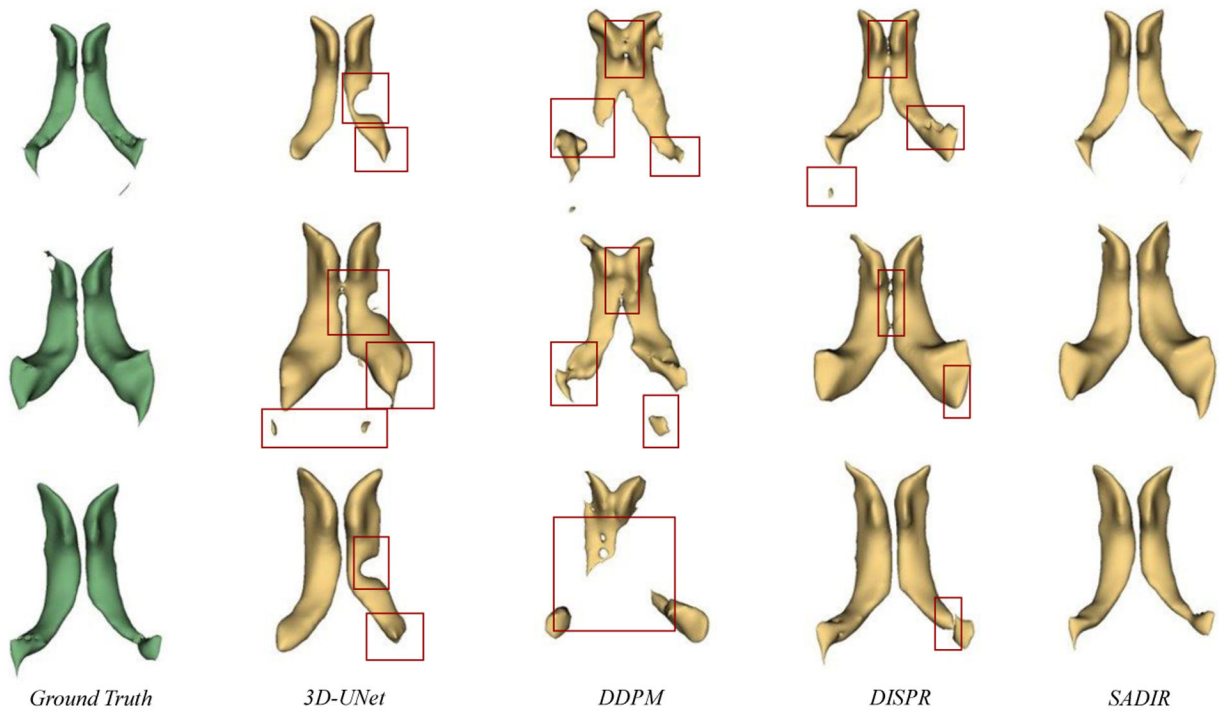


Fig. 2. Top to bottom: examples of reconstructed 3D brain ventricles from sparse 2D slices; Left to right: a comparison of brain ventricles of all reconstruction models with ground truth.

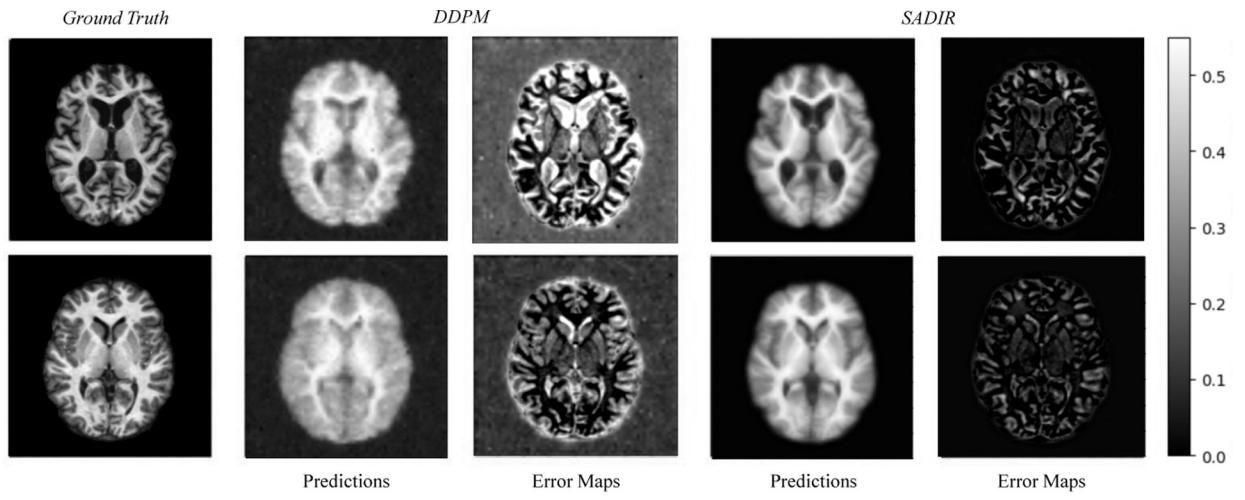


Fig. 3. Left to right: a comparison of ground truth, DDPM, and SADIR along with the error map.

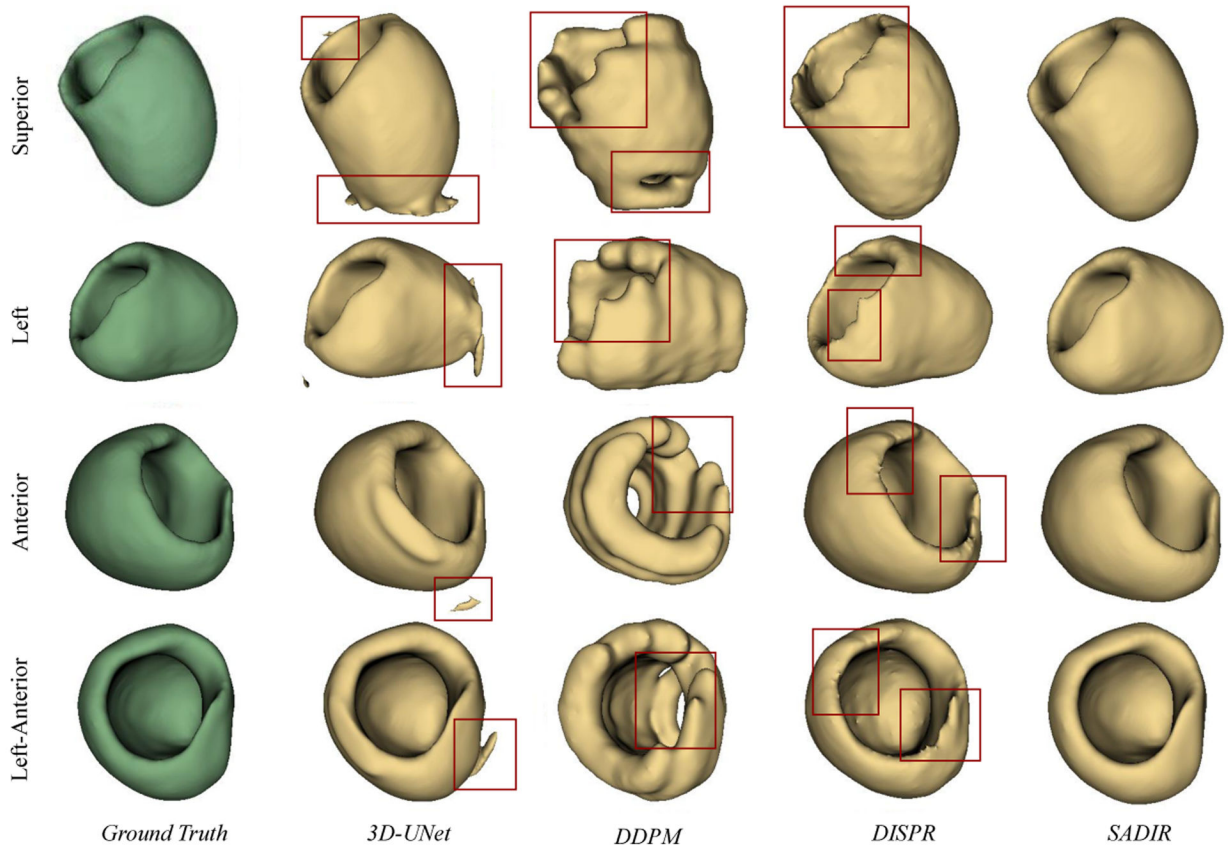


Fig. 4. A comparison of reconstructed 3D myocardium between ground truth, 3D-UNet, DDPM, DISPR, and SADIR over four different views.

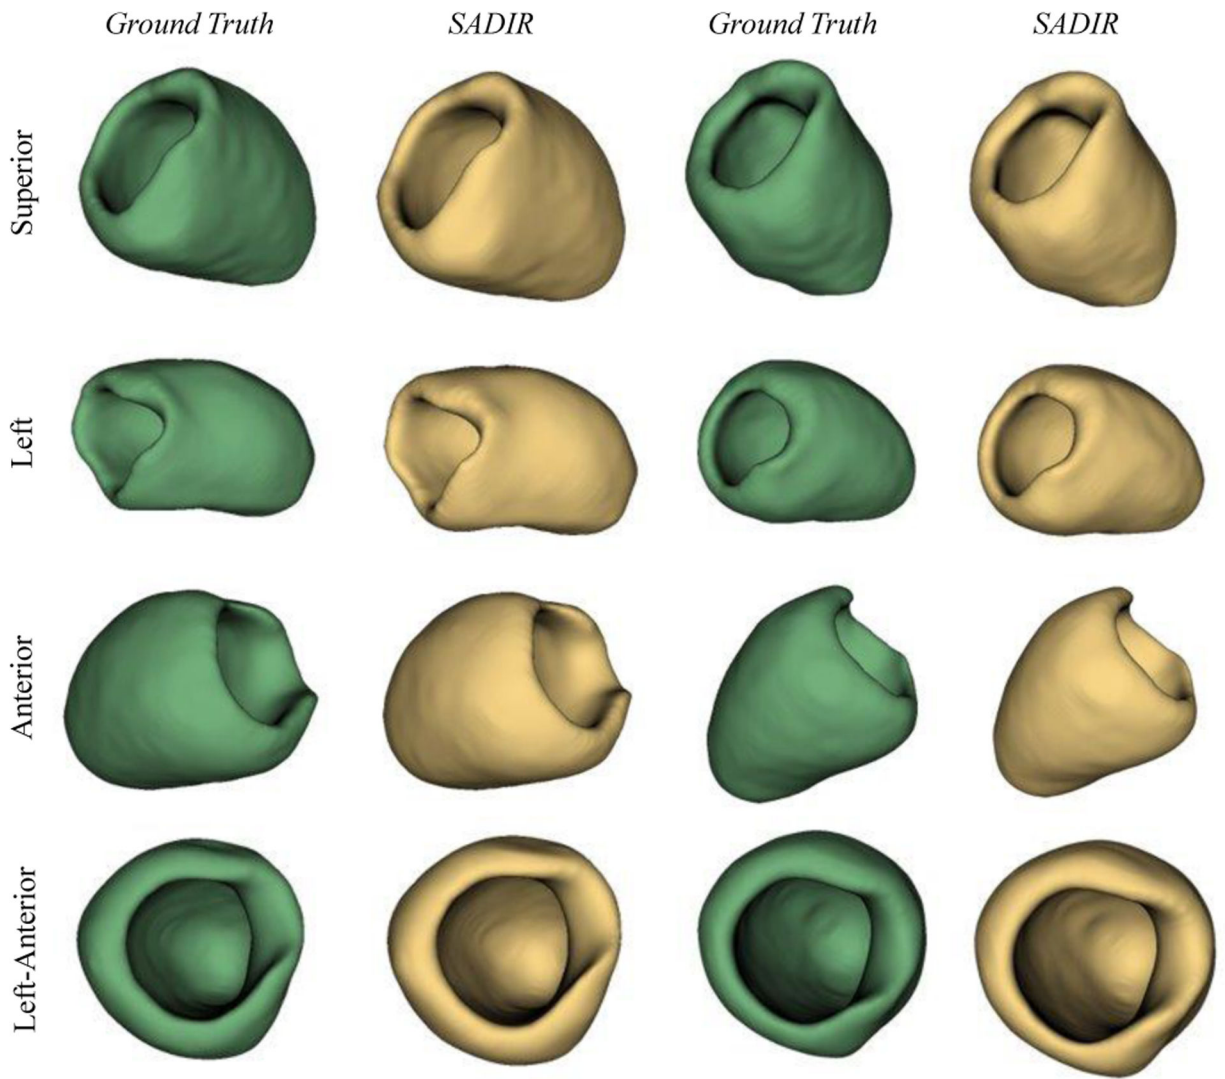


Fig. 5.
3D myocardium reconstructed from sparse 2D slices by SADIR over four different views.

Table 1.

A comparison of 3D brain ventricle reconstruction for all methods.

Model	DSC \uparrow	Jaccard similarity \uparrow	RHD95 \downarrow
3D-Unet	0.878 ± 0.0128	0.804 ± 0.0204	4.366 ± 1.908
DDPM	0.731 ± 0.0292	0.652 ± 0.0365	8.827 ± 9.212
DISPR	0.918 ± 0.0097	0.861 ± 0.0158	1.041 ± 0.130
SADIR	0.934 ± 0.013	0.900 ± 0.021	1.414 ± 0.190

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2.

A comparison of 3D myocardium reconstruction for all methods.

Model	DSC \uparrow	Jaccard similarity \uparrow	RHD95 \downarrow
3D-Unet	0.870 ± 0.0158	0.771 ± 0.024	0.840 ± 0.202
DDPM	0.823 ± 0.014	0.668 ± 0.019	1.027 ± 0.093
DISPR	0.950 ± 0.017	0.906 ± 0.031	0.347 ± 0.032
SADIR	0.978 ± 0.016	0.957 ± 0.031	0.341 ± 0.023

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript