



Fully Automatic 3D Bi-Atria Segmentation from Late Gadolinium-Enhanced MRIs Using Double Convolutional Neural Networks

Zhaohan Xiong¹(✉), Aaqel Nalar¹, Kevin Jamart¹, Martin K. Stiles², Vadim V. Fedorov³, and Jichao Zhao¹

¹ Auckland Bioengineering Institute, University of Auckland, Auckland, New Zealand

zxio506@aucklanduni.ac.nz

² Waikato Clinical School, Faculty of Medical and Health Sciences, University of Auckland, Auckland, New Zealand

³ Department of Physiology and Cell Biology, The Ohio State University Wexner Medical Center, Columbus, USA

Abstract. Segmentation of the 3D human atria from late gadolinium-enhanced (LGE)-MRIs is crucial for understanding and analyzing the underlying atrial structures that sustain atrial fibrillation (AF), the most common cardiac arrhythmia. However, due to the lack of a large labeled dataset, current automated methods have only been developed for left atrium (LA) segmentation. Since AF is sustained across both the LA and right atrium (RA), an automatic bi-atria segmentation method is of high interest. We have therefore created a 3D LGE-MRI database from AF patients with both LA and RA labels to train a double, sequentially used convolutional neural network (CNN) for automatic LA and RA epicardium and endocardium segmentation. To mitigate issues regarding the severe class imbalance and the complex geometry of the atria, the first CNN accurately detects the region of interest (ROI) containing the atria and the second CNN performs targeted regional segmentation of the ROI. The CNN comprises of a U-Net backbone enhanced with residual blocks, pre-activation normalization, and a Dice loss to improve accuracy and convergence. The receptive field of the CNN was increased by using 5×5 kernels to capture large variations in the atrial geometry. Our algorithm segments and reconstructs the LA and RA within 2 s, achieving a Dice accuracy of 94% and a surface-to-surface distance error of approximately 1 pixel. To our knowledge, the proposed approach is the first of its kind, and is currently the most robust automatic bi-atria segmentation method, creating a solid benchmark for future studies.

Keywords: Atrial segmentation · Convolutional neural network · MRI

1 Introduction

Atrial fibrillation (AF) is the most common form of cardiac arrhythmia and is associated with substantial morbidity and mortality [1]. Current clinical treatments for AF perform poorly due to a lack of basic understanding of the underlying atrial anatomical

structure which directly sustains AF in the human atria [2]. In recent years, gadolinium-based contrast agents are utilized in a third of all MRI scans to improve the clarity of a patient's internal structures such as the atria by enhancing the visibility of disease-associated structures such as fibrosis/scarring, inflammation, tumors, and blood vessels [3]. Late gadolinium-enhanced MRI (LGE-MRI) is widely used to study fibrosis/scarring [4], and clinical studies on AF patients using LGE-MRIs have shown that the extent and distribution of atrial fibrosis can be used to reliably predict ablation success rates [5]. As a result, direct analysis of the atrial structure in patients with AF is vital to improving the understanding and patient-specific treatment of AF.

Segmentation of both the left atrial (LA) and right atrial (RA) chambers is a crucial task for aiding medical management for AF patients based on structural analysis of the segmented 3D geometry. Due to the rising popularity of convolutional neural networks (CNN) in the field of medical imaging, many algorithms have been developed utilizing CNNs, particularly for the segmentation of the LA directly from LGE-MRIs [6]. These methods have drastically improved on the previous traditional atlas-based or shape-based approaches [7] in terms of both performance and adaptability due to their fully data-driven nature. In 2018, numerous CNN methods were submitted to the STACOM 2018 Left Atrial Segmentation Challenge [8] aiming at optimizing LA segmentation performance from LGE-MRIs. Through the challenge, the U-Net was shown to be the most widely used and most easily adaptable architecture for the task [9]. In particular, pipelines with enhancements to the U-Net baseline such as the addition of residual connections, dilated convolutions, and custom loss functions achieved far superior segmentation accuracies.

Despite the extensive research in LA segmentation, no established study has been conducted for the fully automatic segmentation of the RA directly from LGE-MRIs. A 2017 benchmarking study investigated methods of segmenting the LA, RA, left ventricle (LV), and right ventricle (RV) from non-contrast MRIs [10]. However, segmentation from LGE-MRIs compared to non-contrast MRIs is much more challenging due to the attenuation caused by the contrast agent resulting in a lack of distinguishable features between the atrial tissue and background. Thus, there is still an urgent need of an intelligent algorithm capable of automatically segmenting both the LA and the RA simultaneously from LGE-MRIs.

In this paper, we propose and evaluate a robust approach for fully automatic segmentation of the atria, particularly the RA, from 3D LGE-MRIs. In order to effectively learn the complex geometry of the atrial chambers, we designed a double CNN pipeline for targeted segmentation and reconstruction of the LA and RA without human intervention. This exciting study is the first of its kind to present a method of segmenting both atrial chambers simultaneously, and is a very important step towards more effective and efficient patient specific diagnostics and treatment.

2 Methods

2.1 Data and Pre-processing

20 3D LGE-MRIs from patients with AF were provided by the University of Utah [4]. The *in vivo* patient images were acquired at a spatial resolution of 0.625 mm \times 0.625 mm \times 1.25 mm using either a 1.5T Avanto or 3.0T Verio clinical whole-body scanner. All 3D LGE-MRI scans contained 44 slices along the Z direction, each with an XY spatial size of 640 \times 640 pixels or 576 \times 576 pixels. The LA segmentations were provided by the University of Utah alongside the LGE-MRIs [4]. The RA segmentations were manually performed by our team based on the protocols used for LA segmentation to achieve consistency across both atrial chambers (Fig. 1c). Firstly, the RA endocardium was defined by manually tracing the RA blood pool in each slice of the LGE-MRI. The tricuspid valve connecting the RA and RV was defined by a 3D plane to create a smooth linear surface. The RA endocardium was then morphologically dilated and manually adjusted according to the RA geometry to obtain the boundary of the epicardium. Next, the septum, the region of tissue connecting the RA and LA, was manually traced such that the epicardial surfaces of the LA and RA joined together. Finally, the dilated tricuspid valve was manually removed from the RA epicardium. Overall, the three labels for the dataset were the background, the endocardium of the RA and LA, and the epicardium of the RA and LA.

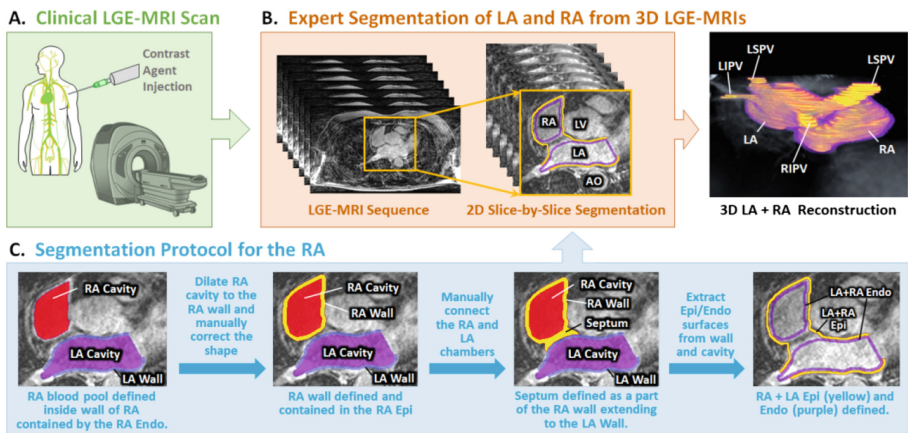


Fig. 1. Data acquisition and the protocol for labelling the left atrial (LA) and right atrial (RA) epicardium (Epi) and endocardium (Endo) from late gadolinium-enhanced magnetic resonance imaging (LGE-MRI). (A) Clinical MRI scanners were used to acquire LGE-MRIs. (B) The LGE-MRIs were manually segmented in a slice-by-slice manner by experts to obtain labels of the LA/RA epicardium and LA/RA endocardium. (C) RA manual annotation based on the LA annotation protocol provided by University of Utah [4]. AO, aorta; LV, left ventricle; PV, pulmonary vein; LS/LIPV, left superior/inferior PV; RS/RIPV, right superior/inferior PV.

The 20 3D LGE-MRI data were randomly split for performing 4-fold cross validation such that in each fold, 15 data was used for training and 5 data was used for validation. All data and labels were uniformly center cropped to a size of $576 \times 576 \times 44$ pixels, and mean and standard deviation normalization were performed. Contrast limited adaptive histogram equalization (CLAHE), a type of color intensity normalization, was also performed on the dataset during pre-processing.

2.2 Convolutional Neural Network Architecture

Our pipeline consisted of two 2D CNNs used in a sequential manner (Fig. 2a). The first CNN performed coarse segmentation on a down-sampled version of the entire 3D LGE-MRI ($144 \times 144 \times 44$) in a slice-by-slice manner to construct an approximate segmentation of the atria. The center of mass of the atria was calculated from the coarse segmentation in each slice of the LGE-MRI, and a 240×240 patch was cropped around this point, leaving out the majority of background pixels which significantly decreased computational costs. 240×240 was chosen it could contain the entire atria which had a maximum size of 200×200 from measurements on the entire dataset. The second CNN then performed slice-by-slice regional segmentation on the ROIs cropped from the 3D LGE-MRI. Finally, the individual slice-by-slice segmentations were stacked together and zero-padded to a size of $576 \times 576 \times 44$ to obtain the final segmentation.

The same CNN was used for both stages of our pipeline and consisted of an enhanced U-Net architecture (Fig. 2b). The first half of the CNN was an encoder to learn dense features from the input through several convolutional layers of increasing depth. The convolutional layers contained 5×5 kernels and a stride of 1, and the number of feature maps increased from 16 to 256. At every 1–3 convolutional layers, residual connections were added to improve feature learning and 2×2 convolutions with a stride of 2 were used to progressively down sample the image by a factor of 2. The second half of the CNN was a decoder to reconstruct the image back to the original resolution for segmentation through several 5×5 convolutional layers of decreasing depth. The number of feature maps of the convolutions in this part of the network increased from 128 to 32. The images were progressively up-sampled by a factor of 2 with 2×2 deconvolutional, or transpose convolutional, layers with stride of 2. Residual connections were also added at every 1–3 convolutional layers. In order to directly preserve high-resolution features from the input, feature forwarding connections were also used to concatenate the outputs of the convolutional layers in the encoder part to those in the decoder path at 4 different points along the CNN. Batch normalization (BN) and parametric rectified linear units (PReLU) were used after every convolutional layer along the entire CNN for improving convergence, and 50% dropout was used at every layer for regularization. The final output layer of the CNN contained a 1×1 convolution with a stride of 1, 3 feature maps, and a softmax activation function to predict for the 3 classes in the data (LA/RA epicardium, LA/RA endocardium, background). All hyper-parameters were selected as a result of extensive experimentation under controlled settings to derive the optimal parameter combinations.

The CNN was trained with the 2D Dice loss function (Eq. 2) to prioritize the segmentation of the foreground pixels over the background. At each epoch, online augmentation was used to randomly augment each data with a probability of 50% and included translation, scaling, rotation and flipping. The adaptive moment estimation gradient descent optimizer was used with a learning rate of 0.0001 and the exponential decay rate of 0.9. After training, the CNN with the highest cross-validation accuracy was selected as the final model. The network was developed in TensorFlow and was trained on an Nvidia Titan V GPU with 5120 CUDA cores and 12 GB RAM. Training took 2 h and predictions on each 3D LGE-MRI took 2 s.

2.3 Evaluation

To measure the accuracy of the first CNN in detecting of the center of mass of the atria when extracting the ROI, the mean squared error (MSE) was defined as

$$MSE = \sqrt{(x - x')^2 + (y - y')^2} \quad (1)$$

for the ground truth co-ordinates (x, y) and the predicted coordinates (x', y') of the center of mass in each 2D slice of the 3D LGE-MRIs.

To measure the accuracy of the second CNN for segmentation, the Dice score, surface-to-surface distance (STSD), sensitivity, and specificity were used. The DICE score was calculated as


$$Dice = \frac{2N_{true\ positive}}{2N_{true\ positive} + N_{false\ positive} + N_{false\ negative}} \quad (2)$$

for the atrial epicardium and endocardium predictions and ground truths. The STSD measured the average distance error between the surfaces of the predicted LA volume, A , and the ground truth, B , and was calculated as

$$STSD(A, B) = \frac{1}{n_A + n_B} \left(\sum_{p=1}^{n_A} \sqrt{p^2 - B^2} + \sum_{p'=1}^{n_B} \sqrt{p'^2 - A^2} \right) \quad (3)$$

where n_A is the number of pixels in A , n_B is the number of pixels in B , and p and p' describes all points in A and B . The sensitivity and specificity were calculated and reflected the success for segmenting the foreground and the background respectively.

A. Proposed Pipeline for Automatic Bi-Atrial Segmentation



B. 2D CNN Architecture

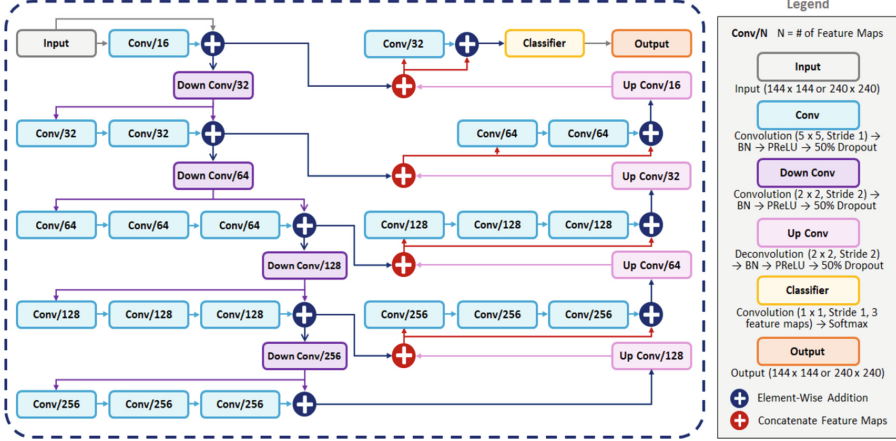


Fig. 2. The proposed double 2D convolutional neural network (CNN) pipeline for fully automatic segmentation of the left atrium (LA) and right atrium (RA) from late gadolinium-enhanced magnetic resonance imaging (LGE-MRI). (A) Overall workflow in which the first CNN detected the region of interest (ROI) containing the LA/RA and the second CNN performed regional segmentation of the ROI. (B) The architecture of the proposed 2D CNN. BN, batch normalization; PReLU, parametric rectified linear unit.

3 Results and Discussion

Table 1 summarizes the final evaluation metrics for the best performing CNN used in our proposed double CNN pipeline. Evaluation shows that the ROI containing the atria for each slice could be calculated within 5 mm, or 8 pixels, of the ground truth through the coarsely segmented atria provided by the first CNN. Individually for each co-ordinate, the CNN had an error of 2.3 mm, or approximately 3 pixels, in the x axis and an error of 3.8 mm, or approximately 6 pixels, in the y axis. The second CNN was evaluated on the entire final 3D segmentations produced for each test data. The LA and RA endocardium were segmented with a Dice score of 92.9% and a STSD of 0.63 mm. The LA and RA epicardium were segmented with a Dice score 94.0% and a STSD of 0.68 mm. Overall, the STSD showed that the predictions were approximately within 1 pixel of the ground on average for both the endocardium and epicardium.

3D visualization of the ground truth and the predictions for the test set in one cross-validation fold shows that the proposed CNN pipeline successfully captured the overall

Table 1. Overall evaluation results of the proposed double CNN pipeline. The first CNN was evaluated on the accuracy of the region of interest detected and the second CNN was evaluated on the accuracy of the final segmentation in 3D.

First CNN		Second CNN				
	Distance (mm)		Dice (%)	STSD (mm)	Sensitivity	Specificity
Error in (x, y)	(2.27 ± 2.78, 3.77 ± 4.67)	LA + RA Endocardium	92.9 ± 1.4	0.63 ± 0.11	92.9 ± 1.6	99.9 ± 0.01
MSE	4.79 ± 5.09	LA + RA Epicardium	94.0 ± 1.4	0.68 ± 0.15	93.7 ± 2.0	99.9 ± 0.01

MSE, mean squared error; STSD, surface-to-surface distance; LA/RA, left/right atrium.

geometry of the LA and RA endocardium with a high degree of precision (Fig. 3). 3D visualization of the STSD errors of the predictions shows the most erroneous regions to be the LA pulmonary veins which some parts containing an error of up to 10 mm from the ground truth. However, minimal error can be seen on the RA in general, showing that the RA may potentially be easier to segment considering there are no complex and small structures such as the pulmonary veins present. Visualizations of the LA and RA epicardium were not shown, however, had comparable accuracies as the endocardium as seen from the higher Dice score and similar STSD.

To further analyze the regions of the predictions containing the most errors, 2D slice-by-slice visualizations and comparisons were performed for each test LGE-MRI. Figure 4 illustrates the segmentation results for the LA and RA endocardium and epicardium by our proposed pipeline compared with the ground truth for selective slices at the same depth for a test 3D LGE-MRI. The results shown are representative of the errors seen in other test LGE-MRIs. The relative depth of each slice from the bottom of the LGE-MRI scan is provided in millimeters. At the bottom slices of the LGE-MRIs (12–18 mm), it can be seen that the network was not able to reproduce the linear plane used to define the tricuspid valve, but instead, produced a more rounded prediction at this region, potentially due to the fact that there is no anatomically visible border between the RA and the RV. The segmentation at the middle slices (18–31 mm) successfully captured the geometry of the RA and LA in detail, while also showing a clear gap between the epicardium and endocardium which denotes the atrial walls. At this region, the main source of error was also at the tricuspid and mitral valve as the CNNs produced segmentations which were smooth and rounded. The septum, on the other hand, was precisely captured despite the decreased contrast in this region, showing the network was able to successfully learn the shape of the epicardium in its entirety. At the top slices of the LGE-MRIs (31–38 mm), the pulmonary vein regions were the main sources of error, consistent with previous the 3D error visualizations. The predictions show that the pulmonary veins predicted by the CNN are smaller and thinner than that of the ground truth labels, further reflecting the difficulty of segmenting this structure. However, the epicardium segmentation is shown to be fairly consistent and effectively accounts for the gaps in its morphology caused by the pulmonary veins and valves across all slices of the LGE-MRI, which was also reflected in the higher Dice score.




Fig. 3. 3D visualizations of the left atrial (LA) and right atrial (RA) chambers for 3 representative test late gadolinium-enhanced magnetic resonance imaging (LGE-MRI). (A) Ground truth provided. (B) Predictions segmented by the convolutional neural network (CNN). (C) Surface-to-surface error of the predictions from the ground truth in millimeters (mm).




Fig. 4. The left atrial (LA) and right atrial (RA) endocardium (orange) and epicardium (blue) results from the proposed convolutional neural network (CNN) pipeline compared to the ground truth for representative slices on the same 3D late gadolinium-enhanced magnetic resonance imaging (LGE-MRI) for a test patient. (A) LGE-MRI scans. (B) Ground truths. (C) Predictions by the CNN. AO, aorta; LV, left ventricle; RV, right ventricle; LPV, left pulmonary veins; RPV, right pulmonary veins; TV, tricuspid valve; MV, mitral valve. (Color figure online)

4 Conclusion

We have developed and evaluated a double convolutional neural network for robust automatic bi-atria segmentation from LGE-MRIs. Our algorithm enables the reconstruction of both the LA and RA chambers in 3D with a Dice accuracy of 94% and a surface-to-surface distance error of approximately 1 pixel from the ground truth. Our study is the first automated method to segment both atrial chambers, particularly the RA, creating a solid benchmark for future studies. The exciting findings from this study may lead to the development of a more accurate and efficient atrial reconstruction and analysis approach, which can potentially be used for improved clinical diagnosis, patient stratification, and clinical guidance during treatment for AF patients.

Acknowledgements. This work was funded by the Health Research Council of New Zealand (16/385) and the National Institutes of Health (HL135109). We would like to acknowledge The NIH/NIGMS Center for Integrative Biomedical Computing (CIBC) at The University of Utah for providing the LGE-MRIs and LA labels. We also thank Nvidia for donating a Titan V GPU to accelerate our research.

References

1. Narayan, S.M., et al.: Ablation of focal impulses and rotational sources: what can be learned from differing procedural outcomes? *Curr. Cardiovasc. Risk Rep.* **11**(9), 27 (2017)
2. Zhao, J., et al.: Three dimensional integrated functional, structural, and computational mapping to define the structural “fingerprints” of heart specific atrial fibrillation drivers in human heart Ex Vivo. *J. Am. Hear. Assoc.* **6**(8), e005922 (2017)
3. Hansen, B.J., Zhao, J., Fedorov, V.V.: Fibrosis and atrial fibrillation: computerized and optical mapping: a view into the human atria at submillimeter resolution. *JACC: Clin. Electrophysiol.* **3**(6), 531–546 (2017)
4. McGann, C., et al.: Atrial fibrillation ablation outcome is predicted by left atrial remodeling on MRI. *Circ. Arrhythmia Electrophysiol.* **7**(1), 23–30 (2014)
5. Oakes, R., et al.: Detection and quantification of left atrial structural remodeling with delayed-enhancement magnetic resonance imaging in patients with atrial fibrillation. *Circulation* **119**, 1758, 1767 (2009)
6. Xiong, Z., Fedorov, V.V., Fu, X., Cheng, E., Macleod, R., Zhao, J.: Fully automatic left atrium segmentation from late gadolinium enhanced magnetic resonance imaging using a dual fully convolutional neural network. *IEEE Trans. Med. Imaging* **38**(2), 515–524 (2019)
7. Tobon-Gomez, C., et al.: Benchmark for algorithms segmenting the left atrium from 3D CT and MRI datasets. *IEEE Trans. Med. Imaging* **34**(7), 1460–1473 (2015)
8. Pop, M.: Statistical Atlases and Computational Models of the Heart. Atrial Segmentation and LV Quantification Challenges. Springer, Heidelberg (2019)
9. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
10. Zhuang, X., et al.: Evaluation of algorithms for multi-modality whole heart segmentation: an open-access grand challenge. arXiv preprint [arXiv:1902.07880](https://arxiv.org/abs/1902.07880) (2019)