

Improved Heatmap-Based Landmark Detection

Huifeng Yao^{1(\boxtimes)}, Ziyu Guo¹, Yatao Zhang¹, and Xiaomeng Li²

¹ Shandong University, Shandong, China

 $^2\,$ Hong Kong University of Science and Technology, Hongkong, China

Abstract. Mitral valve repair is a very difficult operation, often requiring experienced surgeons. The doctor will insert a prosthetic ring to aid in the restoration of heart function. The location of the prosthesis' sutures is critical. Obtaining and studying them during the procedure is a valuable learning experience for new surgeons. This paper proposes a landmark detection network for detecting sutures in endoscopic pictures, which solves the problem of a variable number of suture points in the images. Because there are two datasets, one from the simulated domain and the other from real intraoperative data, this work uses cycleGAN to interconvert the images from the two domains to obtain a larger dataset and a better score on real intraoperative data. This paper performed the tests using a simulated dataset of 2708 photos and a real dataset of 2376 images. The mean sensitivity on the simulated dataset is about 75.64 \pm 4.48% and the precision is about $73.62 \pm 9.99\%$. The mean sensitivity on the real dataset is about $50.23 \pm 3.76\%$ and the precision is about 62.76 \pm 4.93%. The data is from the AdaptOR MICCAI Challenge 2021, which can be found at https://zenodo.org/record/4646979#.YO1zLUxCQ2x.

Keywords: Heatmap \cdot Landmark detection \cdot CycleGAN

1 Introduction

In mitral valve repair, the surgeon repairs part of the damaged mitral valve to allow the valve to fully close and stop leaking. The surgeon may tighten or reinforce the ring around a valve by implanting an artificial ring. The surgeon may place approximately 12 to 15 sutures on the mitral annulus [1]. We need to know how sutures are placed because analyzing the pattern and distances between them can help us improve the quality of this surgery. Furthermore, the position of the sutures may aid the medico in learning how to perform this surgery by reconstructing it in a 3D virtual environment.

Deep learning methods have been widely used in the field of medical images. This task belongs to the landmark detection task in computer vision. In general, people mainly use the heatmap-based [6] method, coordinate regression method,

© Springer Nature Switzerland AG 2021

Huifeng Yao and Ziyu Guo—contributed equally to this work and should be considered joint first-authors.

S. Engelhardt et al. (Eds.): DGM4MICCAI 2021/DALI 2021, LNCS 13003, pp. 125–133, 2021. https://doi.org/10.1007/978-3-030-88210-5_11

and patch-based method. Payer et al. [6] used the SpatialConfiguration-Net which combines the local appearance of landmarks with their spatial configuration. Because the coordinate regression method is too difficult to converge and the patch-based method is difficult to distinguish adjacent points, we choose the heatmap-based method.

Many state-of-the-art heatmap-based deep learning methods focus on detecting fixed key points which are not suitable for our task. Stern et al. [9] proposed a heatmap-based method to detecting a varying number of key points. Inspired by that, we present an improved heatmap-based method that can deal with a varying number of sutures and get better performance than that.

The data set is mainly split into two endoscopic sets. One is simulation data set and the other is real data set. Inspired by Engelhardt et al. [2], we also implement the image to image translation to get more real data. We use the cycleGAN [11] network to do this task.

The work proposed a network to detect a varying number of landmarks and used the cycleGAN to translate images from two different domains. And we are participating within the scope of the AdaptOR challenge.

2 Materials and Methods

2.1 Data Set

Our data set comes from the AdaptOR challenge [8]. The data set is mainly split into two endoscopic sets:

- (1) Sim-Domain is the image acquired during simulating mitral valve repair on a surgical simulator. More information on the simulator can be found in [3] and [4]. The simulator dataset used for training consists of 2708 frames, which were extracted from 10 surgeries. We divide it into 5 fold. To prevent data leakage, dataset splitting was always carried out on the level of the surgeries.
- (2) Intraop-Domain is the Intraoperative endoscopic data from real minimally invasive mitral valve repair. Since the intraoperative dataset consists of 2376 frames extracted from 4 simulated surgeries, we split it into 4 fold with each surgery comprising one fold.

The Label of this data set is stored in the format of a JSON file. In addition, the data splitting is shown in Table 1.

Domain	Split	Number of frames					
		f1	f2	f3	f4	f5	
Sim	Train	2246	2144	1960	2174	2308	
	Validation	462	564	748	534	400	
Intraop	Train	1582	1852	2004	1690	_	
	Validation	794	524	372	686	_	

Table 1. Data set.

2.2 Outline of the Proposed Method

We have a lot of simulated data, but we don't have enough real data. So the first step is the image to image translation. We use a cycleGAN to convert simulated data to real data in order to obtain more real data, which will help our model score higher on the real dataset. The second step is to generate the heatmap. Unlike other tasks about landmark detection, which use one channel for each landmark, we do not have fixed points in this task. So we generate all the points in one channel. And each of them is a 2D Gaussian kernel. We do some augmentation for both the original image and heatmap. Then the enhanced images would be the input of the U-net-based [7] network. The corresponding heatmap would be the label of the image. Then, we use the Otsu [5] to get the thresholding image. We also Use the open operation to remove the noise in the image and make the binarized area smoother. Finally, we use the cutting method to separate very close points and the centroid of each region is taken as the final result. All of these are shown in Fig. 1.



Fig. 1. Outline

2.3 Pre-processing

Image to Image GAN. In this task, our datasets come from two domains, one is the simulation domain and the other is the Intraop domain. The data set of the Intraop domain is smaller than the data set of the simulation domain. We decided to transform the simulation domain data into Intraop domain data to get a higher score on the Intraop domain. We introduced cycleGAN to solve this problem.

The cycleGAN has two mapping functions, as shown in Fig. 2, one is G and the other is F. G transforms the image of X domain into the image of Y domain, and F transforms the image of Y domain into the image of X domain. Two discriminators identify the real domain image and the generated image.



Fig. 2. CycleGAN

Applied to this task, the overall flow is shown in Fig. 3. This diagram only shows the process from the simulation domain to the Intraop domain and vice versa, which is not shown here.



Fig. 3. CycleGAN in this task

Heatmap. Unlike the traditional landmark detection method, we do not generate a heatmap for each point but generate all the points onto the same heatmap.

Because in other tasks, the number of feature points is fixed, while in our task, the number of feature points varies with the image, ranging roughly from 0 to 15.

Each of our points is a 2D Gaussian kernel, and a variable number of points make up this heatmap, which will be used as the model's label. The heatmap is shown in Fig. 4.



Fig. 4. Heatmap

Data Augmentation. During training the images are randomly augmented using Albumentations functions: horizontally and vertically with a probability of 50%, rotation of $\pm 40^{\circ}$, ColorJitter with a probability of 50%, RandomBrightnessContrast with a probability of 50%.

2.4 Point Detection

This work uses a U-Net-based architecture with a depth of 5. After each 3×3 convolution, batch normalization is applied. The first convolutional layer has 16 filter maps, while the bottleneck layer has 512 filter maps. We choose the Resnext [10] network as our encoder. We don't have an activation function after the final 1×1 -convolutional layer while training, but we apply the sigmoid function when we predict the heatmap. The loss function is dice loss.

The input images are RGB images with 3 channels. One channel output is the heatmap. The heatmap becomes the real output point after a series of subsequent operations.

2.5 Post-processing

Otsu. The maximum between-class variance method is a nonparametric and unsupervised method of automatic threshold selection for picture segmentation. According to the gray characteristics of the image, the image is divided into background and objects. Among them, the greater the variance between the background and objects shows that the difference between the two parts of the image is also greater. This method calculates the relationship between the average gray level between background pixels and foreground pixels and their proportion in the whole image, so as to obtain the global threshold when the image segmentation effect is the best, and finally segment the image according to this value.

Opening. We discovered that the network's predicted images were connected together in blocks that should have been separated after Otsu. The open operation is used to separate them. This also smoothes the edges of the segmented blocks and removes some of the noise.

Centre Mass and Cutting. After the opening process, we identify the centroid of each segmented block in the output image, and these are regarded the final predicted points, but we discovered that the form of some of these blocks compared to the circle generated by the point in the heatmap image is somewhat irregular. Therefore, we assess whether a segmentation block should be clipped depending on whether the area of each segmented block in an output image exceeds the average value of all its segmented blocks. Then, based on the height and width of the segmented blocks' bounding box, decide the cutting direction. The cutting point is the centroid of the segmented blocks that need to be sliced. Cutting is done in the x-axis direction if the bounding box's height is higher than its width. If the bounding box's height is less than its width, the cutting is done using Cut in the y-axis direction. We recalculate the centroid of the partitioned block after cutting as the output points and save them in JSON files.

The example of post-processing is shown in Fig. 5.



Fig. 5. Example of post-processing. (a) input, (b) predict, (c) Otsu, (d) opening, (e) Centre mass and Cutting

2.6 Evaluation

A point detection is considered successful if the centres of mass of ground truth and prediction are less than 6 pixels apart. On an image of size 512×288 , this radius roughly corresponds to the thickness of a suture when it enters the tissue. Every matched point from the produced mask is considered a true positive (TP). Predicted points that could not be matched to any ground truth point are defined as false positives (FP) and all ground truth points without a corresponding point in the produced mask are false negatives (FN). Precision and sensitivity are computed over all landmarks. And F1-score presents the harmonic mean of precision and sensitivity.

$$Precision = \frac{TP}{TP + FP} \tag{1}$$

$$Sensitivity = \frac{TP}{TP + FN} \tag{2}$$

$$F1 = \frac{2 * Precision * Sensitivity}{Precision + Sensitivity}$$
(3)

3 Results

There are some visual examples in Fig. 6(a) and Fig. 6(b).



Fig. 6. Example results of two domain. The green circles are true positives (TP). The red circles show false positives (FP). The yellow circles represent false negatives (FN).

The result of the simulation domain is shown in Table 2, and the result of the intraop domain is shown in Table 3.

The baseline results come from this paper [9]. We can not calculate the standard deviation of the baseline F1 score since the baseline does not give experimental data for F1 score.

As shown in Tables 2 and 3, while our precision is lower than the baseline, our sensitivity is much higher. As a result, when comparing F1 scores, our method outperforms the baseline on both the simulation and intraop domains. Because the images in the intraop domain have more interference factors and less data, the recognition effect of the two methods in the intraop domain is slightly inferior to that of the simulation domain.

Cross-validation result on Simu data							
Metric	Model	f1	f2	f3	f4	f5	$\mu \pm \sigma$
Precision	Baseline	-	-	-	-	-	81.50 ± 5.77
	Ours	84.37	54.79	76.84	74.18	77.89	73.62 ± 9.99
Sensitivity	Baseline	-	-	-	-	-	61.60 ± 6.11
	Ours	79.63	72.25	68.64	80.20	77.48	$\textbf{75.64} \pm \textbf{4.48}$
F1 score	Baseline	-	-	-	-	-	69.78
	Ours	81.94	62.33	72.51	77.07	77.69	$\textbf{74.31} \pm \textbf{6.69}$

Cross-validation result on Intra data							
Metric	Model	f1	f2	f3	f4	$\mu \pm \sigma$	
Precision	Baseline	-	-	-	-	66.68 ± 4.67	
	Ours	62.24	67.35	54.92	66.54	62.76 ± 4.93	
Sensitivity	Baseline	_	_	_	_	24.45 ± 5.06	
	Ours	51.81	54.44	44.22	50.45	$\textbf{50.23} \pm \textbf{3.76}$	
F1 score	Baseline	-	-	-	-	35.78	
	Ours	56.56	60.22	48.99	57.38	$\textbf{55.79} \pm \textbf{4.15}$	

Table 3. Intra result.

4 Conclusions

We present a novel method for predicting multiple key points in endoscopic images in this paper. Our method differs from traditional key point detection methods, which have a fixed number of prediction key points. Our method can detect multiple key points at the same time, significantly reducing detection time and model calculation. We also introduce cycleGAN, which can interconvert images from two domains to create a larger dataset. Our results outperform the baseline as well as other related methods after many repeated and rigorous experiments.

References

- 1. Carpentier, A., Adams, D.H., Filsoufi, F.: Carpentier's Reconstructive Valve Surgery E-Book. Elsevier Health Sciences, Amsterdam (2011)
- Engelhardt, S., De Simone, R., Full, P.M., Karck, M., Wolf, I.: Improving surgical training phantoms by hyperrealism: deep unpaired image-to-image translation from real surgeries. In: Bildverarbeitung für die Medizin 2019. I, pp. 282–282. Springer, Wiesbaden (2019). https://doi.org/10.1007/978-3-658-25326-4_62
- Engelhardt, S., Sauerzapf, S., Brčić, A., Karck, M., Wolf, I., De Simone, R.: Replicated mitral valve models from real patients offer training opportunities for minimally invasive mitral valve repair. Interact. Cardiovasc. Thorac. Surg. 29(1), 43–50 (2019)
- Engelhardt, S., Sauerzapf, S., Preim, B., Karck, M., Wolf, I., De Simone, R.: Flexible and comprehensive patient-specific mitral valve silicone models with chordae tendineae made from 3d-printable molds. Int. J. Comput. Assist. Radiol. Surg. 14(7), 1177–1186 (2019)
- Otsu, N.: A threshold selection method from gray-level histograms. IEEE Trans. Syst. Man Cybern. 9(1), 62–66 (1979)
- Payer, C., Štern, D., Bischof, H., Urschler, M.: Integrating spatial configuration into heatmap regression based CNNs for landmark localization. Med. Image Anal. 54, 207–219 (2019)
- 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

- Sharan, L., et al.: Mutually improved endoscopic image synthesis and landmark detection in unpaired image-to-image translation. arXiv preprint arXiv:2107.06941 (2021)
- Stern, A., et al.: Heatmap-based 2d landmark detection with a varying number of landmarks. Bildverarbeitung f
 ür die Medizin 2021. Informatik aktuell. Springer Vieweg, Wiesbaden (2021)
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1492–1500 (2017)
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)