

Heatmap-based 2D Landmark Detection with a Varying Number of Landmarks

Antonia Stern¹, Lalith Sharan¹, Gabriele Romano², Sven Koehler¹,
Matthias Karck², Raffaele De Simone², Ivo Wolf³, Sandy Engelhardt¹

¹Group Artificial Intelligence in Cardiovascular Medicine (AICM), Department of
Internal Medicine III, Heidelberg University Hospital, Heidelberg

²Department of Cardiac Surgery, Heidelberg University Hospital, Heidelberg

³Faculty of Computer Science, University of Applied Sciences, Mannheim
`sandy.engelhardt@med.uni-heidelberg.de`

Abstract. Mitral valve repair is a surgery to restore the function of the mitral valve. To achieve this, a prosthetic ring is sewed onto the mitral annulus. Analyzing the sutures, which are punctured through the annulus for ring implantation, can be useful in surgical skill assessment, for quantitative surgery and for positioning a virtual prosthetic ring model in the scene via augmented reality. This work presents a neural network approach which detects the sutures in endoscopic images of mitral valve repair and therefore solves a landmark detection problem with varying amount of landmarks, as opposed to most other existing deep learning-based landmark detection approaches. The neural network is trained separately on two data collections from different domains with the same architecture and hyperparameter settings. The datasets consist of more than 1,300 stereo frame pairs each, with a total over 60,000 annotated landmarks. The proposed heatmap-based neural network achieves a mean positive predictive value (PPV) of $66.68 \pm 4.67\%$ and a mean true positive rate (TPR) of $24.45 \pm 5.06\%$ on the intraoperative test dataset and a mean PPV of $81.50 \pm 5.77\%$ and a mean TPR of $61.60 \pm 6.11\%$ on a dataset recorded during surgical simulation. The best detection results are achieved when the camera is positioned above the mitral valve with good illumination. A detection from a sideward view is also possible if the mitral valve is well perceptible.

1 Introduction

Mitral valve reconstruction is a complicated surgery, where the surgeon implants a prosthetic ring to the tissue by placing approx. 12 to 15 sutures on the mitral annulus [1]. Analysing how sutures are placed (e.g., their pattern and distances) can help in improving the quality and consistency of the procedure. This could be useful in surgery itself and in preoperative simulation for trainees [2]. Besides that, the position of the sutures may be utilized to reconstruct the 3D position of the mitral annulus in a stereo-setting and a virtual ring model can be superimposed onto the endoscopic video stream prior to ring implantation [3,4].

Several deep learning-based methods exist for landmark detection in computer-vision and medical applications. In general, they can be classified into heatmap approaches [5,6], coordinate regression and patch-based [7] approaches. Gilbert et al. [5] used a U-Net-like architecture to detect anatomical landmarks by predicting one heatmap per landmark. The input label heatmaps contain 2D Gaussians centred at the annotated point location. In contrast to that, coordinate regression approaches directly output the coordinates instead of heatmaps, which is less computationally expensive as no upsampling is necessary. Usually, these approaches consist of a sequence of convolutional layers along with a fully connected layer or a 1×1 -convolution. Another common approach is to detect landmarks in global and local patches with the help of displacement vectors [7].

State-of-the-art literature on deep learning-based landmark detection focuses mainly on detecting a predefined number of keypoints. In medical applications, landmarks are often used to define a fixed number of anatomical points. These approaches output one heatmap for each landmark or use a vector representation of fixed length with the length corresponding to the number of keypoints. However, these methods cannot be applied to a setting with an arbitrary number of keypoints. Furthermore, in our scenario, unlike in most other works, the landmarks can lie in close proximity to each other, which makes the application of patch-based approaches difficult.

The work presented in this paper uses a network architecture that is able to deal with a varying number of landmarks, since, in principle, a varying number of sutures can be placed. This work uses a U-Net based architecture with a *single* foreground heatmap as output. Thresholding this heatmap allows us to detect a non-predefined number of points in endoscopic images, which represent the positions where the sutures enter the tissue.

2 Materials and methods

2.1 Network Architecture

This work uses a U-Net-based [8] architecture with a depth of 4. After each 3×3 -convolution, batch normalisation is applied. The padded 3×3 -convolutional layers use the *ELU* activation function. The final layer is a 1×1 -convolutional layer with a *sigmoid* activation function. The first convolutional layer has 16 filter maps, while the bottleneck layer has 256 filter maps. The number of filters is doubled/halved after every maxpooling/upsampling in the downsampling/upsampling path. The dropout rate is set to values between 0.3 and 0.5 with the lowest dropout rate in the first downsampling/last upsampling block. Unlike the original U-Net [8], our architecture applies zero-padding in the convolutional layers. Therefore, the input size has to be divisible by 4^{depth} , which equals 16 in this work. The loss function is of the form $MSE - SDC$, where MSE is the mean squared error function and SDC , the Sørensen dice coefficient.

To preserve the aspect ratio of the images, a width of 512 pixels and a height of 288 pixels was chosen as the input size. The input images are RGB-images

with 3 channels. The two-channel output masks contain one channel for the suture landmarks and one for the background. This design is more efficient than heatmaps approaches, which use one channel for each landmark. The predicted output heatmap is converted into a binary mask by thresholding. Then, the centre of mass for each region is determined as point of interest.

2.2 Dataset

Two instances of the same neural network approach proposed in this work were trained separately on two different data sub-collections (Tab. 1, A.1 and B) consisting of endoscopic images of mitral valve repair. One data collection contains very heterogeneous *intraoperative* images extracted from videos which were recorded during surgery (different view angles, light intensities, number of sutures, surgical instruments, prosthesis etc.). Different amounts of frames were annotated from these datasets, therefore an additional hold-out test set was created from surgeries with significantly less frames (A.2). The other *simulation* data collection was extracted from videos of operations performed on an artificial mitral valve replica made of silicone, which was first introduced in [2] and used in surgical training applications.

The endoscopic videos were recorded from an Image S1 stereo camera (Karl Storz SE & Co. KG, Tuttlingen, Germany) in full HD resolution or larger at 25 fps. Frames were saved in top-down format (left image top, right image bottom). Relevant scenes were identified before extracting the frames from the videos and afterwards every 120th frame was extracted. In scenes with rapid changes, every 10th frame was extracted and in scenes with only few changes, every 240th frame was extracted. These characteristics were identified manually.

This work uses the two subimages of the stereo recording as separate input images to the network, thus the final number of frames for training and testing is twice as large as given in Tab. 1. The intraoperative dataset used for training consists of 2654 frames, which were extracted from 5 surgeries (mean $530,8 \pm 213,6$ frames per surgery). An additional balanced intraoperative test dataset was annotated, containing 200 frames extracted from 4 surgeries (mean 100 ± 0 frames per surgery). The simulator datasets consists of 2708 frames extracted from 10 simulated surgeries (mean $270,7 \pm 77,7$ frames per surgery).

The frames were manually annotated with the tool *labelme* [9]. During annotation, corresponding suture points in the left and right image of the stereo pair are joined with a line. If a suture point is only visible in one of the subimages it is marked as a point. The ground truth heatmaps were then created by placing 2D Gaussians at the position of the landmarks. For a down-scaled input size of 512×288 , a variance σ of 1 was used.

During training the images are randomly augmented with a probability of 80% using tensorflow functions: rotation of $\pm 60^\circ$, pixel shifting in a range of $\pm 10\%$, mask pixel shifting in a range of $\pm 1\%$, shearing in a range of ± 0.1 , brightness in a range of ± 0.2 , contrast in a range from 0.3 to 0.5, random saturation in a range from 0.5 to 2.0 and hue in a range of ± 0.1 . Additionally the images were flipped horizontally and vertically with a probability of 50%.

Table 1. Three sub-collections from two different domains were used. Note that we treat left and right stereo frame independently in our work, therefore training frames are doubled. The additional intraop test set has significantly lower number of labeled frames, hence we decided to not include it directly in the cross validation (CV).

domain	usage	# stereo frame	# suture endpoints	# surgeries
A.1 intraop	5-fold CV (train, test)	1,327	26,937	5
A.2 intraop	test	200	4,305	4
B simulator	5-fold CV (train, test)	1,354	33,893	10

2.3 Evaluation

During method development, which involved hyperparameter tuning, training and validation was performed on dataset A.1. The best neural network had a total of 2.1M trainable parameters. The initial learning rate was set to 10^{-3} with a decay factor of 0.1. After fixing these parameters, final evaluation was conducted on A.1, A.2 and B, involving two 5-fold cross validations (CV) on different splits of dataset A.1 and B. Furthermore, an additional disjoint test set was used (A.2) to further assess method generalizability. To prevent data leakage, dataset splitting was always carried out on the level of the surgeries. The results are given for the epoch ≤ 200 with the lowest value of the validation loss. In total, 10 different models were trained, one for each fold of each application (intraop and simulation).

A suture 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. If a point in the produced mask or the ground truth mask is assigned multiple times, only the matched pair with the least distance is kept. 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). To evaluate the precision (positive predictive value, PPV) and sensitivity (true positive rate, TPR) are used. Precision and sensitivity are defined as $PPV = TP / (TP + FP)$ and $TPR = TP / (TP + FN)$. The values of PPV and TPR are displayed over a threshold from 0.05 to 1.0 to show the influence of thresholding the heatmap on the detection rate.

3 Results

Some visual examples are provided in Fig. 1(a) and Fig.1(b). In intraoperative images, green sutures are better recognized than white ones because they can be better distinguished from the background. During surgery (not during simulation), white sutures often appear red because they are soaked with blood which makes it even harder to distinguish them from the background.

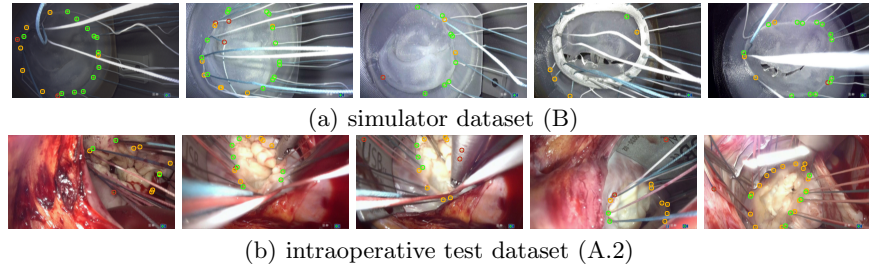


Fig. 1. Example predictions on the two domains. Green circles represent true positives (TP), red circles show false positives (FP) and orange circles show false negatives (FN).

The model achieved a mean PPV of $67.99 \pm 7.69\%$ and a mean TPR of $29.03 \pm 7.74\%$ on the intraoperative dataset (A.1) for a threshold of $t = 0.8$ during testing on the respective hold-out folds (Fig. 2(a)). All models from the CV were also applied to the separate test set (A.2) and comparable results were obtained with a mean PPV of $66.68 \pm 4.67\%$ and a mean TPR of $24.45 \pm 5.06\%$ (Fig. 2(b)), meaning that it generalizes beyond the surgeries where hyperparameter tuning was performed on. When training the model on the simulator dataset (B) with the same settings, a mean PPV of $81.50 \pm 5.77\%$ and a mean TPR of $61.60 \pm 6.11\%$ is achieved during CV (Fig. 2(c)). As can be seen in the plots, the estimated TPR and PPV are not sensitive with regard to the chosen threshold. Performance differences between the folds are expressed by the bars in Fig. 2.

4 Discussion

The developed approach is able to detect a varying number of landmarks, which represent the position of sutures stitched through the mitral annulus. Evaluation on a large dataset of two domains revealed that the neural networks detects sutures in most of the scenes, but had difficulties in cases where the mitral annulus is partly occluded by the prosthetic ring, the surgical tools or tissue (Fig. 1(a) 4th image). Reflections or embossings on tools has also led to an increase in FPs (Fig. 1(b) 3rd image). The best detection results were achieved when the




Fig. 2. The mean PPV and TPR of all folds with variable mask threshold. The bars show the value of the fold with the lowest and highest value respectively.

camera was positioned above the mitral valve with good illumination. Comparing both domains, the network performance was better on the simulation domain, where the camera angle is often more favorable and the silicone appearance is more homogeneous in comparison to intraoperative scenes with various tissue textures and blood. However, the performance differences could be also explained by the higher numbers of simulated surgery instances used during training in the simulator domain.

In this work, two neural networks were trained separately on two datasets from different domains with the same architecture and hyperparameter settings. Future work includes incorporating image-to-image translation by generative adversarial networks [10] to adapt between the domains and to allow for joint landmark detection in both domains.

Acknowledgement. The research was supported by the German Research Foundation DFG Project 398787259, DE 2131/2-1 and EN 1197/2-1 and by Informatics for Life funded by the Klaus Tschira Foundation.

References

1. Carpentier A, Adams D, Filsoufi F. Carpentier’s Reconstructive Valve Surgery. Saunders; 2010.
2. Engelhardt S, Sauerzapf S, Preim B, et al. Flexible and comprehensive patient-specific mitral valve silicone models with chordae tendinae made from 3D-printable molds. *Int J Comput Assist Radiol Surg.* 2019;14(7):1177–1186.
3. Engelhardt S, De Simone R, Zimmermann N, et al. Augmented reality-enhanced endoscopic images for annuloplasty ring sizing. In: *Augmented Environments for Computer-Assisted Interventions.* Springer International Publishing; 2014. p. 128–137.
4. Engelhardt S, Kolb S, De Simone R, et al. Endoscopic feature tracking for augmented-reality assisted prosthesis selection in mitral valve repair. In: *Proc SPIE, Medical Imaging: Image-Guided Procedures, Robotic Interventions, and Modeling.* vol. 9786; 2016. p. 402–408.
5. Gilbert A, Holden M, Eikvil L, et al. Automated left ventricle dimension measurement in 2D cardiac ultrasound via an anatomically meaningful CNN approach. In: *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis.* Springer International Publishing; 2019. p. 29–37.
6. Jin H, Liao S, Shao L. Pixel-in-pixel net: towards efficient facial landmark detection in the wild. *arXiv:200303771v1 [csCV].* 2020;.
7. Noothout JMH, De Vos BD, Wolterink JM, et al. Deep learning-based regression and classification for automatic landmark localization in medical images. *IEEE Trans on Med Imag.* 2020; p. 1–1.
8. Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. In: *MICCAI.* vol. 9351 of LNCS. Springer; 2015. p. 234–241.
9. Wada K. labelme: image polygonal annotation with python; 2016. <https://github.com/wkentaro/labelme>.
10. Engelhardt S, Simone RD, Full PM, et al. Improving surgical training phantoms by hyperrealism: deep unpaired image-to-image translation from real surgeries. In: *MICCAI.* Springer International Publishing; 2018. p. 747–755.