Lennart Bargsten*, Daniel Klisch, Katharina A. Riedl, Tobias Wissel, Fabian J. Brunner, Klaus Schaefers, Michael Grass, Stefan Blankenberg, Moritz Seiffert, and Alexander Schlaefer

Deep learning for guidewire detection in intravascular ultrasound images

Abstract: Algorithms for automated analysis of intravascular ultrasound (IVUS) images can be disturbed by guidewires, which are often encountered when treating bifurcations in percutaneous coronary interventions. Detecting guidewires in advance can therefore help avoiding potential errors. This task is not trivial, since guidewires appear rather small compared to other relevant objects in IVUS images. We employed CNNs with additional multi-task learning as well as different guidewire-specific regularizations to enable and improve guidewire detection. In this context, we developed a network block which generates heatmaps that highlight guidewires without the need of localization annotations. The guidewire detection results reach values of 0.931 in terms of the F1-score and 0.996 in terms of area under curve (AUC). Comparing thresholded guidewire heatmaps with ground truth segmentation masks leads to a Dice score of 23.1 % and an average Hausdorff distance of 1.45 mm. Guidewire detection has proven to be a task that CNNs can handle quite well. Employing multi-task learning and guidewire-specific regularizations further improve detection results and enable generation of heatmaps that indicate the position of guidewires without actual labels.

Keywords: Multi-task learning, Coronary artery, Vessel, Heatmap, Regularization, Segmentation.

https://doi.org/10.1515/cdbme-2021-1023

***Corresponding author: Lennart Bargsten:** Hamburg University of Technology, Institute of Medical Technology and Intelligent Systems, Hamburg, Germany, E-mail: lennart.bargsten@tuhh.de
Daniel Klisch, Alexander Schlaefer: Hamburg University of Technology, Institute of Medical Technology and Intelligent Systems, Hamburg, Germany
Katharina A. Riedl, Fabian J. Brunner, Stefan Blankenberg, Moritz Seiffert: Department of Cardiology, University Heart & Vascular Center Hamburg, Hamburg, Germany
Tobias Wissel, Michael Grass: Philips Research - Hamburg, Germany
Klaus Schaefers: Philips Research - Eindhoven, The Netherlands

1 Introduction

Intravascular ultrasound (IVUS) is a widely used imaging modality for planning and performing percutaneous coronary interventions. Automated analysis of IVUS images can help streamlining the clinical workflow by avoiding time-consuming manual delineations of important structures to be quantified or by enabling detection of hardly visible objects. This includes segmentation of lumen and vessel wall [1, 5, 9, 10] as well as detection of calcifications or stents [2, 7, 8]

When treating bifurcations, two guidewires are often used, resulting in a guidewire appearing in IVUS sequences. These can lead to artifacts which complicate interpretation of image content. Automatic guidewire detection would therefore provide valuable information to physicians. Furthermore, image processing algorithms could produce meaningless results due to the presence of a guidewire, as its bright appearance could lead to false findings in other tasks, e.g., calcium or stent segmentation. To prevent such behavior by taking targeted countermeasures, it would be necessary to detect guidewires automatically beforehand.

Our contribution is twofold. First, we show that guidewire detection is feasible by means of convolutional neural networks (CNNs). Second, we developed a novel CNN block that transforms a feature map into a single output neuron for leveraging guidewire detection with multi-task learning using auxiliary lumen and vessel wall segmentation labels. This method additionally enables producing heatmaps which localize guidewires without localization labels for training. To further enhance heatmap generation, we make use of regularization techniques that take into account guidewire brightness and geometry.

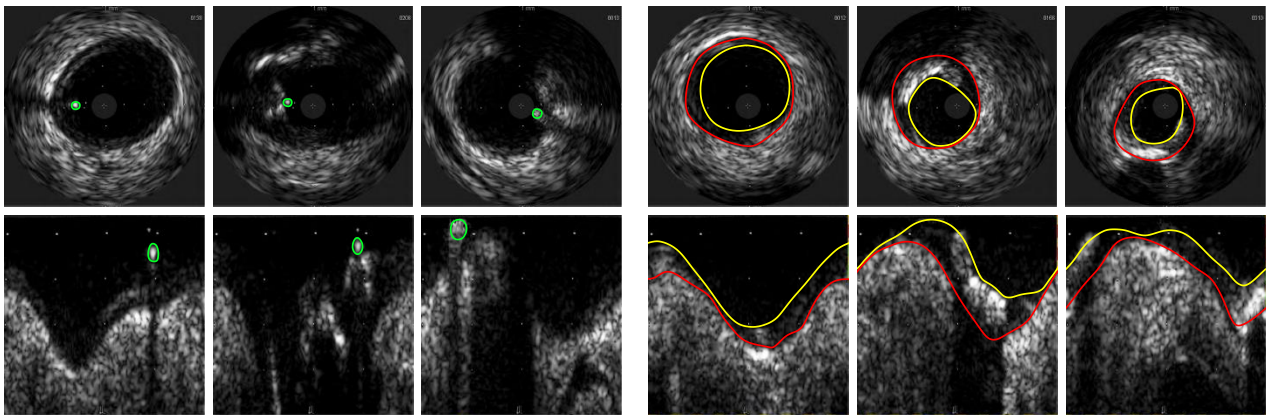

Figure 1: Exemplary images from the guidewire (left) as well as the lumen and vessel wall (right) segmentation dataset. Lumen is indicated by yellow color, vessel wall by red color. The first row shows images in Cartesian view. The second row shows the same images transformed into polar coordinates.

2 Methods and material

2.1 Datasets

We collected three distinct IVUS datasets. The first set consists of 22,213 images from 23 patients with corresponding binary guidewire detection labels (guidewire present or not). The second dataset consists of 165 images from 9 patients with corresponding guidewire segmentation masks. The third set comprises 410 images from 22 patients with corresponding segmentation annotations delineating lumen and vessel wall. We will call this dataset “lumen-wall” throughout the rest of this work. The annotations were made by an experienced cardiologist. All images have a size of 500×500 pixels and were acquired with a 20 MHz phased array Eagle Eye Platinum probe (Philips Healthcare, San Diego, USA). Figure 1 depicts some exemplary images from both segmentation datasets in Cartesian view (first row) and transformed into polar coordinates (second row).

2.2 CNN architectures

For detection without auxiliary data (baseline), we used an encoder-CNN with two residual blocks [4] per downsampling stage. For multi-task learning with auxiliary lumen and vessel wall segmentation data as well as generating guidewire heatmaps, we used an encoder-decoder CNN similar to U-Net [6] but comprising residual blocks (HeatPool CNN). Sketches of the architectures are depicted in Figure 2. In the case of multi-task learning with auxiliary segmentation labels, the network outputs four masks. Three for lumen and vessel wall segmentation and a single mask for guidewire detection. The former three masks were obtained by applying softmax. The latter mask serves as a guidewire heatmap and is transformed into a single neuron by the HeatPool block explained in the next section. For visualization, the heatmap is transformed with a sigmoid function and thresholded at 0.5.

2.3 Heatmapping via pooling

To leverage guidewire detection with auxiliary images and corresponding segmentation masks, we developed a simple CNN block that transforms a feature map into a single detection neuron. This block comprises an average pooling layer followed by a max pooling layer, both with window size and stride equal to the square root of the spatial input feature map size. Experiments revealed that these feature maps highlight regions near the object which the network is trained to detect. Because it generates some kind of heatmap via two pooling layers, we call this block HeatPool. See Figure 2 for a sketch of the HeatPool block (bottom right corner). For visualization, the heatmaps were transformed with a sigmoid function and thresholded at 0.5.

2.4 Regularizations

In order to improve the placement of highlighted regions in the guidewire heatmap, we added auxiliary regularization terms to the overall loss function. These make use of gray value distributions of input images and the geometry of guidewires.

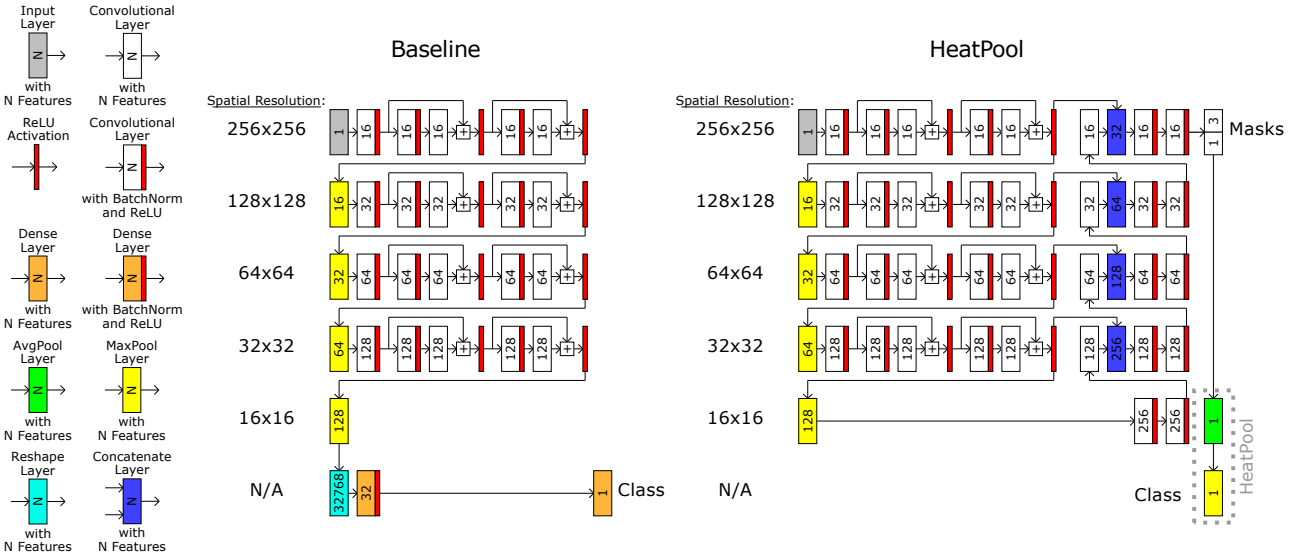

Figure 2: Sketches of CNN architectures used for guidewire detection. The HeatPool CNN can be used with and without multi-task learning by auxiliary segmentation data.

2.4.1 Brightness regularization

In IVUS images, guidewires pretty much always appear as bright spots (see Figures 1 and 3). We therefore developed a loss function for penalizing heatmaps which highlight dark regions of the image. Since we standardize images before feeding them into the CNNs, all pixels with gray values below average are negative. This can be used to define the following loss function for image locations x:

$$\mathcal{L}_b = -\text{mean}_x \min\{p(x) \cdot I(x), 0\} \quad (1)$$

with heatmap value $p(x) \in [0,1]$ and input image gray value $I(x)$.

2.4.2 Convex regularization

Since the guidewire appears convex in IVUS images, the area highlighting the guidewire in the generated heatmaps should also be convex. A set X is convex if and only if

$$a + (b - a) \cdot t \in X \quad (2)$$

for all points $a, b \in X$ and $t \in [0, 1]$. This means that for any two points in X all points on the line connecting the two points belong to X as well. In the case of our heatmap, we can define the following loss as a differentiable approximation of the convexity condition:

$$\mathcal{L}_c = \text{mean}_{a,b,t}\{p(a)p(b)(1 - p(a + (b - a)\, t))\}. \quad (3)$$

However, calculating this expression for all possible a, b and t is way too expensive. We therefore used an approximation by calculating 16 random combinations of a, b and t for locations x in the heatmap with values $p(x) > 0.1$. This approach is meaningful since $p(x)$ appears to be near zero for most of the image.

2.5 Training and evaluation

We split the datasets as follows:

- Guidewire detection: 19459 images from 20 patients for training and 2754 images from 5 patients for testing
- Lumen-wall segmentation: 262 images from 18 patients for training and 57 images from 4 patients for testing

The guidewire segmentation dataset was used only for assessing the heatmaps, not for training.

We trained all methods with eight-fold cross-validation and tested the best performing network of every fold with the distinct test sets. No patient appeared in both training and test set. If ground truth data was not available during multi-task training, the corresponding term in the loss function was set to zero. For example, when an image from the guidewire detection dataset was fed into the network, the loss term for

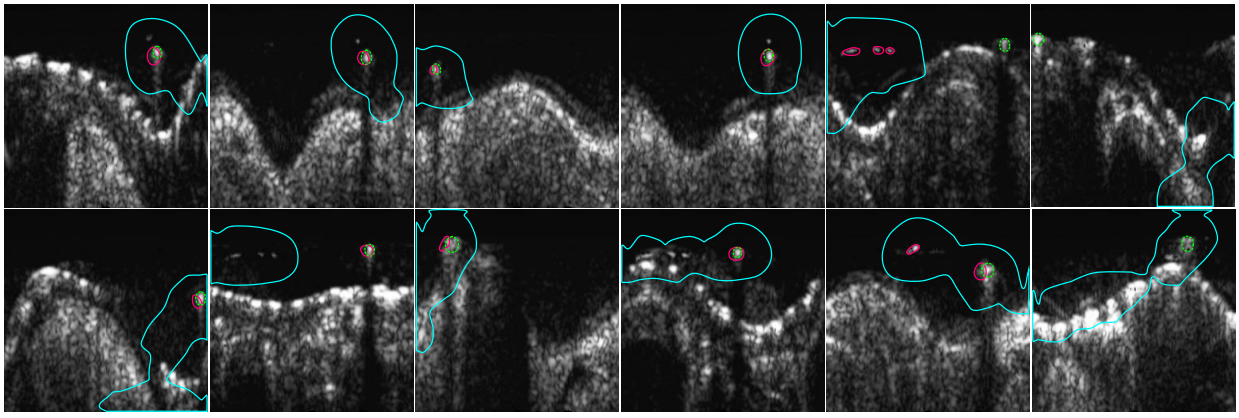

Figure 3: Exemplary thresholded guidewire heatmaps and corresponding ground truth segmentation masks for the case of HeatPool CNN with combined brightness and convex regularization. Heatmap segmentations are depicted magenta. Ground truth segmentation contours are shown as green dashed lines. Results of a similar heatmapping technique [11] are shown in cyan.

Table 1: Guidewire recognition results measured by F1-score and AUC. Multi-task learning and regularization methods for heatmaps are not applicable to the baseline detection model.

model	m-t	reg.	F1-score [%]	AUC [%]
baseline	n/a	n/a	77.6 ± 8.6	99.2 ± 0.5
HeatPool	no	no	88.4 ± 16.7	98.2 ± 2.7
		brightness	72.1 ± 32.5	96.8 ± 5.4
		convex	89.3 ± 7.7	98.2 ± 3.5
		both	85.6 ± 5.7	98.8 ± 1.4
	yes	no	82.5 ± 23.0	95.3 ± 10.7
		brightness	$\mathbf{93.1 \pm 5.3}$	$\mathbf{99.6 \pm 0.6}$
		convex	89.4 ± 6.5	99.0 ± 1.2
		both	89.2 ± 14.8	99.6 ± 0.6

lumen-wall segmentation was set to zero because no ground truth segmentation mask was available for this image. All input images were resized to 256×256 pixels and transformed into polar coordinates. We performed on-the-fly data augmentation by means of random shifts and horizontal flips. For both detection and segmentation, we employed the cross-entropy loss function extended with the presented regularization terms in certain cases.

For evaluating guidewire detection performance, we used the F1-score as well as the area under the receiver operating characteristic curve (AUC). For assessing guidewire localization of the resulting heatmaps, we used the guidewire segmentation dataset which provides ground truth segmentation masks. We calculated the Dice coefficient as a measure of overlap and the average Hausdorff distance [3] as a measure of edge alignment between ground truth and thresholded heatmaps. The threshold was set to 0.5.

3 Results and discussion

Detection results based on F1-score and AUC are shown in Table 1. The HeatPool CNN almost always outperforms the baseline CNN in terms of the F1-score. An exception is the case of brightness regularization without multi-task learning, which also seems to be rather unstable leading to a large variance. The best performing approach is HeatPool with multi-task learning and brightness regularization which outperforms the baseline CNN relatively by 20 % in terms of the F1-score. The AUC is quite constant over all methods but approaches with multi-task learning reach larger values on average. HeatPool without regularization tends to be unstable resulting in larger variances. HeatPool's improvements in F1-score due to a larger network capacity cannot be ruled out completely. However, adding multi-task learning and regularizations improve the results further.

The generated heatmaps were thresholded at 0.5 and compared to ground truth segmentation masks. The results are shown in Table 2. Although not trained with guidewire segmentation labels, a Dice score of 23.1 % is reached with multi-task learning and brightness regularization. The best average Hausdorff distance of 1.45 mm was obtained with convex regularization. The variance of many results is quite high. This shows that the quality of the generated heatmaps varies a lot, e.g., heatmaps without highlighted areas but present guidewire increase the error severely (half of image size in case of the average Hausdorff distance). Exemplary segmentations (magenta) are depicted in Figure 3 together with ground truth segmentation masks (dashed green) and results of a similar heatmapping technique [11]. It can be seen that our predicted guidewire masks are much better localized. Most of our masks fit well (first four columns), while others

Table 2: Segmentation metrics of heatmaps measured by Dice coefficient (Dice) and average Hausdorff distance (ave. HD).

regularization	multi-task	Dice [%]	ave. HD [mm]
none	no	8.6 ± 8.7	1.75 ± 0.57
	yes	13.2 ± 9.0	1.96 ± 0.38
brightness	no	9.4 ± 13.0	2.27 ± 1.25
	yes	$\mathbf{23.1 \pm 12.9}$	1.50 ± 0.35
convex	no	15.1 ± 8.9	2.15 ± 0.45
	yes	15.8 ± 8.9	$\mathbf{1.45 \pm 0.28}$
both	no	17.6 ± 14.6	1.51 ± 0.26
	yes	20.2 ± 12.8	1.81 ± 0.48

do not (fifths column) or are completely missing (last column). A major problem seems to be ambiguities with other bright areas in the lumen (column five). Hence, the presented approach needs to be improved to provide appropriate results with clinical relevance. One possibility would be to penalize heatmaps which highlight more than a single region.

4 Conclusion

We showed that guidewire detection in intravascular ultrasound (IVUS) images can successfully be performed by convolutional neural networks (CNNs). This allows other, more general IVUS image analysis algorithms to take targeted countermeasures against potential errors due to presence of guidewires. Furthermore, we presented a method which generates heatmaps indicating locations of guidewires without actual localization labels. This method works best when adding an auxiliary lumen and vessel wall segmentation task as well as guidewire-specific regularizations, which in the same way improves guidewire detection performance. However, to make this heatmapping approach clinically applicable, e.g., for localizing calcifications, some limitations have to be tackled in the future. In particular, incorrectly highlighted bright ambiguous regions.

Author Statement

Research funding: This work was partially funded by the European Regional Development Fund (ERDF) and the Free and Hanseatic City of Hamburg in the Hamburgische Investitions- und Förderbank (IFB)-Program PROFI Transfer Plus under grant MALEKA.

Compliance with ethical standards: The local institutional review board approved our retrospective single-center study and waived the requirement for informed consent.

References

[1] Bargsten L, Riedl KA, Wissel T, Brunner FJ, Schaefers K, Sprenger J, Grass M, Seiffert M, Blankenberg S, Schlaefer A. Tailored methods for segmentation of intravascular ultrasound images via convolutional neural networks. In: Medical Imaging: Ultrasonic Imaging and Tomography. 2021;11602:1–7

[2] Ciompi F, Balocco S, Rigla J, Carrillo X, Mauri J, Radeva P. Computer-aided detection of intracoronary stent in intravascular ultrasound sequences. Medical Physics 2016;43:5616–5625

[3] Dubuisson MP, Jain A. A modified hausdorff distance for object matching. In: Proceedings of the 12th International Conference on Pattern Recognition. 1994:566–568

[4] He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition. 2016:770–778

[5] Katouzian A, Angelini ED, Carlier SG, Suri JS, Navab N, Laine AF. A state-of-the-art review on segmentation algorithms in intravascular ultrasound (ivus) images. IEEE Transactions on Information Technology in Biomedicine 2012;16:823–834

[6] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention. 2015:234–241

[7] Ulli TC, Gupta D. Segmentation of calcified plaques in intravascular ultrasound images. In: Smart Computing Paradigms: New Progress and Challenges. 2020:57–67

[8] Wissel T, Riedl KA, Schaefers K, Nickisch H, Brunner FJ, Schnellbaecher N, Blankenberg S, Seiffert M, Grass M. Delineation of coronary stents in intravascular ultrasound pullback In: Medical Imaging 2021;11598:210–216

[9] Xia M, Yan W, Huang Y, Guo Y, Thou G, Wang Y. Extracting membrane borders in ivus images using a multi-scale feature aggregated u-net. In: 42nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 2020:1650–1653

[10] Yang J, Faraji M, Basu A. Robust segmentation of arterial wall in intravascular ultrasound images using dual path u-net. Ultrasonics 2019;96:24–33

[11] Zhou B, Khosla A, Lapedriza A, Oliva A, Torralba A. Learning deep features for discriminative localization. In: IEEE Conference on Computer Vision and Pattern Recognition. 2016:2921–2929