# TRAINING A CNN FOR GUIDEWIRE DETECTION

*Donghang Li, Adrian Barbu*

Florida State University
Department of Statistics
Tallahassee, FL 32306, USA

## ABSTRACT

Guidewire detection is an important and challenging problem in image-guided interventions. The guidewire is barely visible in fluoroscopic sequences, since it is thin and the image has poor quality. Most recent methods for guidewire localization have a first level of pixel-wise detection based on a trained classifier on hand-crafted features. A Convolutional Neural Network (CNN) could in principle learn its own features, however training a CNN for guidewire detection has proved to be difficult because the wire is very thin and can have any orientation. In this paper we present a method to train a Fully Convolutional Neural Network for guidewire detection, and highlight what challenges are encountered during training for this particular problem. We also introduce the Spherical Quadrature Filters (SQF) for guidewire detection and show how they can be used to improve the training data. Experiments show that the trained CNN outperforms many popular approaches such as the Frangi filter, the SQF and a trained classifier based on hand-crafted feature. Furthermore, we observe that a CNN approach that uses the SQF to obtain better aligned training examples further improves the detection accuracy.

***Index Terms***— guidewire detection, fluoroscopy, convolutional neural networks, spherical quadrature filters

## 1. INTRODUCTION

Guidewire detection is a challenging problem with wide applications in coronary angioplasty interventions. During the intervention, a catheter is inserted through the femoral artery all the way to the heart, and a guidewire is used to guide different tools beyond the catheter, inside the heart. Then the cardiologist inserts a balloon into the obstructed coronary artery, inflates it to widen the narrowing, and places a stent there to keep the blood vessel open. All these operations are monitored by the cardiologist using real-time X-ray (fluoroscopy) images. The fluoroscopy images are usually low-dose to limit the amount of radiation received by the patient, which makes them noisy and the guidewire poorly visible.

Examples of fluoroscopic images of guidewire are shown in Figure 1. As the figure reveals, the guidewire is thin and hardly visible. Thus, robust detection of the guidewire could help the cardiologist have a better visualization and possibly further reduce the radiation dose administered to the patient.

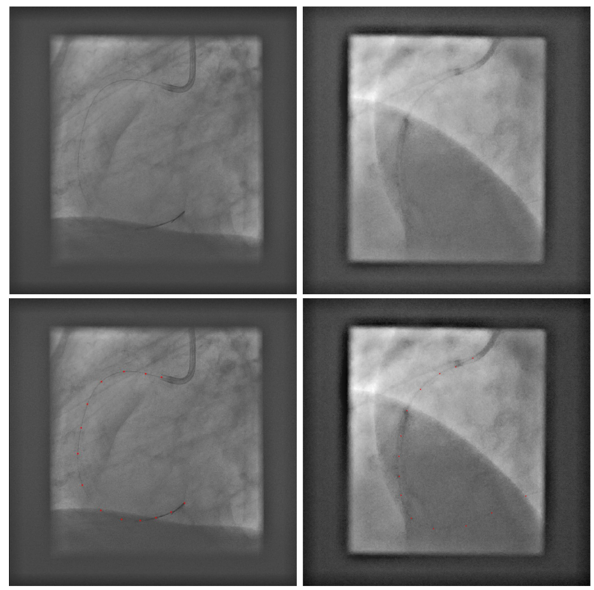

**Fig. 1**. Examples of the fluoroscopic images of guidewire(first two images) and examples of fluoroscopic images with annotations(last two images).

To detect the guidewire one needs to first obtain a low level detection layer that tells how likely is the guidewire to pass through any pixel of the image. As it will be discussed in more detail in the related work section below, guidewire detection works have two main types of approaches to obtain this first level of pixelwise guidewire detection. The first approach is filter based, which uses a predefined filter to obtain a filter response map. The second approach is learning-based, and uses a learning algorithm (Boosting, Random Forest, etc) together with hand-crafted features (e.g. Haar or rotated Haar) to obtain a per-pixel probability map. A third approach would be to train a CNN (Convolutional Neural Network) for this purpose, which will learn its own features using the training data. We could not find any work that trains a CNN for detecting the guidewire pixels, which is why it will be investigated in this paper.

In this paper, we are interested in training a CNN for guidewire detection. Because the wire is thin and hardly visible, it is difficult to train the CNN as the loss function is flat near the random initialization. We will show how to overcome this issue, through a better initialization obtained from training on a single image.

Another guidewire specific issue is that imprecisions in the annotation make the positive examples misaligned. We will also show how to obtain better aligned training data using Spherical Quadrature Filters (SQF) [1]. The Spherical Quadrature Filters (SQF) [1] are a type of steerable filters derived analytically to obtain maximal responses to edge, line or wedge structures. Examples of log-Gabor and Cauchy SQFs are shown in Figure 2. Steerable filters, first introduced in [2], are oriented filters obtained from a basis using predefined weights that depend on the rotation angle. Moreover, the oriented filter response can be computed using the same predefined weights from the response maps obtained by the basis.

The SQFs have been used in [3] for person identification from grayscale images of the ear and in [4] for detecting faint streaks (space debris) in astronomical images. The ear images have edge/ridge structures, and that is why the SQF were a good fit, but the ear images have no noise. In this paper, we introduce another potential application of SQF, guidewire detection in fluoroscopy images. To our knowledge, we are the first to apply the SQF for this problem.

This paper brings the following contributions:

- It shows how to train a Fully Convolutional Neural Network (FCNN) for guidewire detection and how to escape the flat energy landscape present near a random initialization.

- It introduces the Spherical Quadrature Filters (SQF) for guidewire detection, which work better than the popular Frangi filters.

- It shows how to address another challenge in training a CNN, which is due to the imprecision in the manual annotation of the thin guidewire. For that, it show how to use the SQF response map to obtain better aligned examples.

Our experiments reveal that the CNN trained with SQF-aligned examples is the best, followed by the CNN and then the SQF. Furthermore, all three methods introduced in this paper greatly outperform the existing guidewire detection methods such as the Frangi filters and a trained classifier with hand-crafted features.

### 1.1. Related Work

All guidewire detection methods rely on a first level of pixel-wise guidewire detection that applies either a predefined filter or a trained classifier to all locations of the image to obtain a pixelwise guidewire response map.

Filter-based methods include [5] and the Frangi Filter [6], are widely used to detect vessel-like structures. Both of these methods are based on the sorted eigenvalues $(\lambda_1, \lambda_2)$ of the Hessian matrix at every pixel. The Frangi filter was used in [7] as the data term for guidewire tracking in X-ray videos. The filter-based approaches are attractive due to their simplicity and interpretability. However, the Frangi filter was compared in [8] with a voting-based approach that integrates many candidate curves through each pixel, and the Frangi filter results were clearly inferior to the voting-based approach.

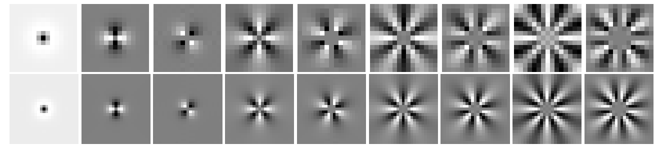

**Fig. 2**. Spherical Quadrature Filters [1] of order 0,2,4,6,8. Top: log-Gabor filters. Bottom: Cauchy filters.

Learning-based methods include [9, 10, 11, 12, 13, 14, 15]. A hierarchical model for guidewire localization was introduced in [10], where longer and longer parts of the guidewire were detected using a Probabilistic Boosting Tree (PBT) [16] and Haar or other hand-crafted features. The PBT and hand-crafted features were also used in [11] for semiautomatic guidewire localization with user constraints, in [12] for guidewire tracking and in [15] for vessel detection. The catheter was detected by a learning-based framework using Boosting and Haar features in [9], and a comparison with the Frangi filter showed that the learning based approach obtained smaller tracking errors than the filter based approach. Their work was aimed at detecting the catheter, while the guidewire is much thinner and more challenging. In [13], a Boosted classifier was trained on ridge and edge features to detect the guidewire pixels. In [17] the guidewire was detected using a Random Forest classifier trained on hand-crafted features.

A CNN approach to detecting the guidewire location was presented in [18]. The method uses the Region Proposal Network to place a bounding box around the guidewire, but it does not obtain pixel-wise response map. In contrast, the method proposed in this paper is aimed at detecting the guidewire pixels using a CNN.

### 1.2. Overview of the Spherical Quadrature Filters

The Spherical Quadrature Filters (SQF) [1] are obtained by the convolution of a generalized Hilbert transform kernel and an isometric filter. The $n$-th order SQF has the following form in the spatial domain:

$$SQF^{(n)}(x,y) = G(x,y) * \left( \frac{n}{2\pi} \frac{(-(x+\mathrm{i}y))^n}{\|x+\mathrm{i}y\|^{2+n}} \right) \quad (1)$$

where "$*$" denotes convolution, and in the Fourier domain

$$S\hat{Q}F^{(n)}(\mathbf{u}) = \hat{G}(\mathbf{u}) \cdot \left( i \frac{\mathbf{u}}{\|\mathbf{u}\|} \right)^n \quad (2)$$

where $x, y \in \mathbb{R}, \mathbf{u} \in \mathbb{R}^2, n \in \mathbb{N}^*$ and $G(x,y)$ is a bandpass isometric filter. In this paper we will generate the SQFs using bandpass filters such as the log-Gabor filters [19]

$$G_l(\omega) = n_c \exp\left( -\frac{\log^2(\omega/\omega_0)}{2\log^2(\sigma)} \right), \quad (3)$$

Gaussian derivative filters [19]

$$G_d(\omega) = \begin{cases} n_c \omega^{(\omega_0\sigma)^2} \exp(-(\sigma\omega)^2), & \text{if } \omega \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

and Cauchy filters [19]

$$G_{Cauchy}(\omega) = \begin{cases} n_c \omega^{\omega_0\sigma} \exp(-\sigma\omega), & \text{if } \omega \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

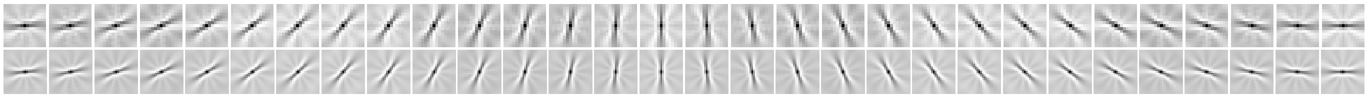

**Fig. 3**. Steered SQF Filters. Top: log-Gabor filters of rank 9. Bottom: Cauchy filters of rank 11.

where $\omega_0 \in \mathbb{R}$ is the peak tuning frequency, and $\sigma \in \mathbb{R}$ such that $\omega_0\sigma \geq 1$. Eq. (3), (4) and (5) are defined in the frequency domain. For more details see [1].

Observe that except for the order 0 SQF, the higher order SQFs come in pairs as the real and imaginary part of eq. (1) or (2). For ridge detection, we only need the even order (symmetric) SQFs, and we will use all the even SQFs of order $n < r$ where $r$ is an odd number. Observe that for any odd number $r > 0$ there are exactly $r$ SQFs of even order $n < r$, and we will call them the SQF of rank $r$.

The SQF bank of rank $r$ can be steered to an angle $\theta$ by dot product multiplication with the following weight vector:

$$\mathbf{w}(\theta)=[-1, -\cos(2\theta_k), \sin(2\theta_k), ..., -\cos(2r\theta_k), \sin(2r\theta_k)].$$

An example of steered log-Gabor SQFs of rank 9 and steered Cauchy SQFs of rank 11 is shown in Figure 3.

One could directly obtain the filter response map by using the SQFs or could use a Convolutional Neural Network for detecting the guidewire pixels as described in the next section.

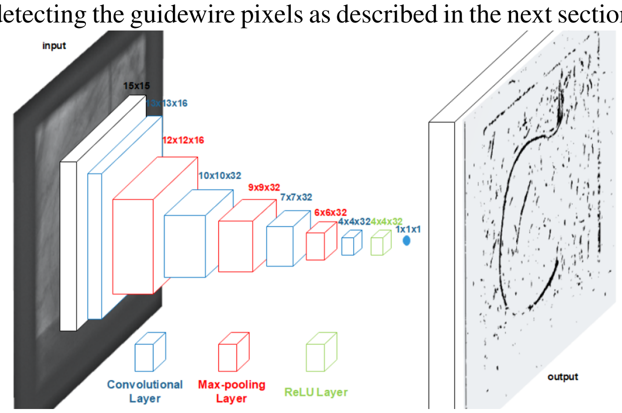


**Fig. 4**. The Fully Convolutional Neural Network (FCNN) used in this paper with input patches of size $15 \times 15$.

## 2. TRAINING A FULLY CONVOLUTIONAL NETWORK FOR GUIDEWIRE DETECTION

Assume we are given $n$ training patches $(\mathbf{x}_i, y_i), i = 1, ..., n$ where $\mathbf{x}_i \in \mathbb{R}^{p^2}$ is the image of a patch of size $p \times p$ either centered on the guidewire (a positive example) or away from the guidewire (a negative), and $y_i$ is the label. The labels are $y_i = -1$ for negative patches and $y_i = 1$ for patches centered on the guidewire.

**CNN architecture.** We implemented a Fully Convolutional Neural Network (FCNN) for guidewire detection. The network (Figure 4) is composed of 5 convolutional layers, the first three layers are followed by $2 \times 2$ max-pooling with stride 1, while the fourth layer is followed by ReLU (Rectified Linear Unit). The last convolutional layer obtains the binary guidewire/non-guidewire response.

For a receptive field of size $15 \times 15$, the first convolutional layer has 16 filters of size $3 \times 3$, and the next three layers have

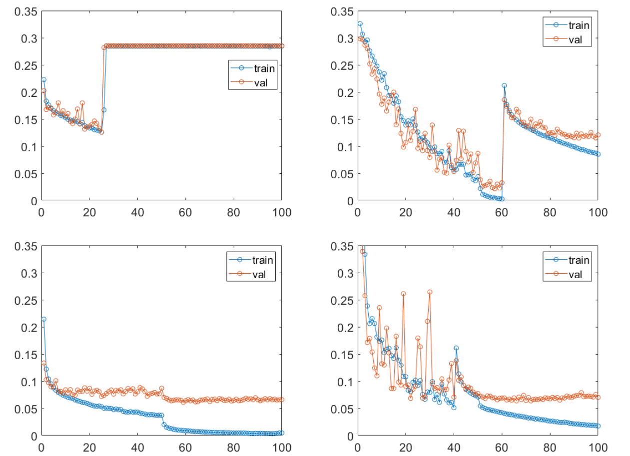


**Fig. 5**. The plot of loss function for 100 epochs of a patch size $25 \times 25$. Top left: training loss with all training examples. Top right: training loss with our approach. Bottom Left: training loss using all NMS-based examples. Bottom right: training loss of NMS examples with our approach

32 filters of size $3 \times 3$ each. The last layer is $4 \times 4$. For a receptive field of size $25 \times 25$, the first layer has 16 filters of size $5 \times 5$, the following layers have 32 filters of size $5 \times 5$, and the last layer is $6 \times 6$.

For training we used the Lorenz loss [20]

$$\ell(u) = \log(1 + \max(1 - u, 0)^2) \tag{6}$$

due to its ease of training and robustness to outliers.

**Training examples.** As positive examples we used image patches at distance at most 1 pixel from the annotation, while negative examples were at distance at least 8 pixels from annotation.

**Training initialization.** All weights were initialized with random Gaussian values with std 0.01. The initial learning rate was 0.01 with mini-batch size 32. The learning rate was multiplied by 0.8 and the minibatch was doubled every 50 epochs, for a total of 300 epochs.

Training the FCNN directly from the random initialization does not work because the guidewire is very thin and the energy landscape becomes flat near the random initialization. Indeed, as shown on the top left side of Figure 5, the loss becomes flat at around 0.285 after epoch 26. In this case, results show that every pixel of the response map is considered detected. To overcome this problem we started by training the first 40-60 epochs using the training examples from only one sequence as shown on the top right side of Figure 5. After that, training was done on all training examples.

**NMS-based alignment.** Another issue we observed was that the annotation was not precise enough to obtain a precise alignment of the positive examples. As a result, the false positive rate, while better than the other methods, was still

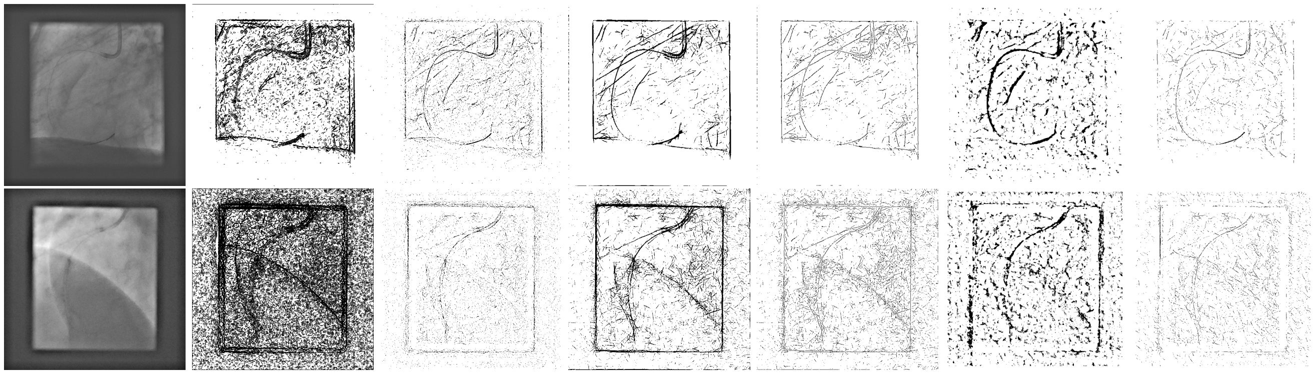

**Fig. 6**. Guidewire detection examples of 2 frames. From left to right: input image, Frangi filter [6], PBT and Haar features[10, 11, 12, 15], Cauchy SQF [1] of rank 11, SQF NMS image, FCNN $25 \times 25$, FCNN NMS detection result.

rather high. To obtain a better alignment we used the Cauchy rank 11 SQF maximum response map on which we performed non-maximal suppression (NMS) in the direction of the image gradients. An example of the SQF NMS map is shown in Figure 6. Then we used as training examples only patches centered on the NMS response map. The positives were at distance at most 2 pixels from annotation, the negatives were at least 8 pixels from annotation.

## 3. EXPERIMENTS

**Dataset.** Experiments are conducted on a dataset of 69 fluoroscopic sequences with a total of 766 frames of different sizes in the range $[512, 700] \times [512, 1024]$. The sequences were divided into a training set containing 33 sequences with 342 frames and a test set containing 36 sequences with 424 frames. The guidewire was manually annotated in all the frames using B-splines.

For training we used positive and negative patches of size $15\times15$ or $25\times25$. In both cases the training set contains about 213,000 positives and about twice as many negatives. The training set using NMS alignment contains about 91,000 positives and as about twice as many negatives. We also implemented the approach from [15] based on about 100,000 oriented Haar features and a Probabilistic Boosting Tree (PBT) [16] and trained it on the same data.

**Comparison with other methods.** We evaluated the detection performance on the training and test images. A guidewire pixel was considered detected if there is a detection (response above the threshold) at distance at most 2 pixels from it. A detection was considered a false positive if it is at distance at least 3 pixels from the guidewire or any catheter.

The response map obtained by any method was thresholded to obtain a binary detection image as shown in Figure 6, with the threshold chosen so that the average detection rate was about 90%. The input image on the second row is noisier than the first one in Figure 6.

The results are shown in Table 1. FCNN has the lowest false positive rate on the training set, but is outperformed by the FCNN NMS on the test set. The SQF with a Cauchy filter of rank 11 performs the best among the SQFs, and also outperforms the Frangi Filter [6] and the PBT with Haar Features. This is very good considering that training the PBT takes about 24h, and training the FCNN takes 6h.

**Table 1**. Per-image evaluation of different filter based and training based guidewire detection methods

| | Det. rate | | FP rate | |
|---|---|---|---|---|
| Method | Train | Test | Train | Test |
| Frangi Filter [6] | 89.98 | 89.96 | 26.99 | 23.74 |
| SQF [1] Gauss deriv, $f_0$=1/2, rk. 9 | 89.52 | 89.55 | 13.78 | 15.28 |
| SQF [1] Cauchy, $f_0 = 1/6$, rank 9 | 90.06 | 89.91 | 10.04 | 10.31 |
| SQF [1] Cauchy, $f_0 = 1/6$, rank 11 | 89.56 | 89.63 | 9.56 | 10.12 |
| SQF [1] Cauchy, $f_0 = 1/6$, rank 13 | 90.58 | 89.50 | 10.41 | 10.18 |
| SQF [1] log-Gabor, $f_0 = 1/6$, rk. 9 | 90.55 | 89.62 | 11.09 | 10.39 |
| SQF [1] log-Gabor, $f_0 = 1/6$, rk. 11 | 90.19 | 89.43 | 10.17 | 10.22 |
| SQF [1] log-Gabor, $f_0 = 1/6$, rk. 13 | 90.31 | 89.54 | 10.73 | 10.39 |
| 25×25 rk. 9 Cauchy SQF, $f_0$=1/6 | 89.92 | 89.98 | 10.31 | 10.93 |
| PBT and Haar features[10, 11, 12, 15] | 90.23 | 90.08 | 1.92 | 4.57 |
| FCNN, size $15 \times 15$ | 89.86 | 90.02 | 4.24 | 12.51 |
| FCNN, size $25 \times 25$ | 90.37 | 90.04 | **2.43** | 7.94 |
| FCNN NMS patches, size $25 \times 25$ | 90.06 | 90.06 | 2.53 | **4.28** |

## 4. CONCLUSION

In this paper, we presented a method to train a CNN for guidewire detection in fluoroscopic images. Training the CNN is not straightforward because the guidewire is thin, noisy and can have any orientation. Moreover, imprecision of the annotation makes the training even more difficult. To address this problem we showed how to get a better initialization with training examples from one image and how to obtain better aligned training examples using the Spherical Quadrature Filters [1]. Experiments show that the Fully Convolutional Neural Network trained with our SQF-aligned data outperformed all other methods evaluated. In terms of filter-based methods, we evaluated the Frangi filter [6] and the Spherical Quadrature Filters [1]. Compared to a trained classifier, the SQFs are more computationally efficient and can simultaneously obtain the guidewire detection and the orientation of the guidewire with uncertainty quantification. In the future, we plan to apply the SQF and FCNN to automatic guidewire localization. This is a higher level process that uses the pixelwise detection as a data term to find the most likely position of the guidewire by searching in the high dimensional space of all possible curves.

## 5. REFERENCES

[1] Ross Marchant and Paul Jackway, "Feature detection from the maximal response to a spherical quadrature filter set," in *DICTA*. IEEE, 2012, pp. 1–8.

[2] William T Freeman, Edward H Adelson, et al., "The design and use of steerable filters," *IEEE Trans. on PAMI*, vol. 13, no. 9, pp. 891–906, 1991.

[3] Tak-Shing Chan and Ajay Kumar, "Reliable ear identification using 2-d quadrature filters," *Pattern Recognition Letters*, vol. 33, no. 14, pp. 1870–1881, 2012.

[4] Radim Sara and Vojtech Cvrcek, "Faint streak detection with certificate by adaptive multi-level bayesian inference," *European Conference on Space Debris*, 2017.

[5] Shirley AM Baert, Wiro J Niessen, Erik HW Meijering, Alejandro F Frangi, and Max A Viergever, "Guide wire tracking during endovascular interventions," in *MICCAI*, 2000, pp. 727–734.

[6] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever, "Multiscale vessel enhancement filtering," in *MICCAI*, 1998, pp. 130–137.

[7] Bor-Jeng Chen, Ziyan Wu, Shanhui Sun, Dong Zhang, and Terrence Chen, "Guidewire tracking using a novel sequential segment optimization method in interventional x-ray videos," in *ISBI*, 2016, pp. 103–106.

[8] Vincent Bismuth, Régis Vaillant, Hugues Talbot, and Laurent Najman, "Curvilinear structure enhancement with the polygonal path image-application to guide-wire segmentation in x-ray fluoroscopy," in *MICCAI*, 2012, pp. 9–16.

[9] Alexander Brost, Andreas Wimmer, Rui Liao, Joachim Hornegger, and Norbert Strobel, "Catheter tracking: Filter-based vs. learning-based.," in *DAGM-Symposium*, 2010, pp. 293–302.

[10] Adrian Barbu, Vassilis Athitsos, Bogdan Georgescu, Stefan Boehm, Peter Durlak, and Dorin Comaniciu, "Hierarchical learning of curves application to guidewire localization in fluoroscopy," in *CVPR*, 2007, pp. 1–8.

[11] Philippe Mazouer, Terrence Chen, Ying Zhu, Peng Wang, Peter Durlak, Jean-Philippe Thiran, and Dorin Comaniciu, "User-constrained guidewire localization in fluoroscopy," in *SPIE Med. Imag.*, 2009, vol. 7259, pp. 72591K–72591K.

[12] Peng Wang, Terrence Chen, Ying Zhu, Wei Zhang, S Kevin Zhou, and Dorin Comaniciu, "Robust guidewire tracking in fluoroscopy," in *CVPR*, 2009, pp. 691–698.

[13] Nicolas Honnorat, Régis Vaillant, and Nikos Paragios, "Guide-wire extraction through perceptual organization of local segments in fluoroscopic images," *MICCAI*, pp. 440–448, 2010.

[14] Olivier Pauly, Hauke Heibel, and Nassir Navab, "A machine learning approach for deformable guide-wire tracking in fluoroscopic sequences," *MICCAI*, pp. 343–350, 2010.

[15] Terrence Chen, Gareth Funka-Lea, and Dorin Comaniciu, "Robust and fast contrast inflow detection for 2d x-ray fluoroscopy," *MICCAI*, pp. 243–250, 2011.

[16] Zhuowen Tu, "Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering," in *ICCV*. IEEE, 2005, vol. 2, pp. 1589–1596.

[17] Erkang Cheng, Yu Pang, Ying Zhu, Jingyi Yu, and Haibin Ling, "Curvilinear structure tracking by low rank tensor approximation with model propagation," in *CVPR*, 2014, pp. 3057–3064.

[18] Li Wang, Xiao-Liang Xie, Gui-Bin Bian, Zeng-Guang Hou, Xiao-Ran Cheng, and Pusit Prasong, "Guide-wire detection using region proposal network for x-ray image-guided navigation," in *IJCNN*. IEEE, 2017, pp. 3169–3175.

[19] Djamal Boukerroui, J Alison Noble, and Michael Brady, "On the choice of band-pass quadrature filters," *Journal of Mathematical Imaging and Vision*, vol. 21, no. 1-2, pp. 53–80, 2004.

[20] Adrian Barbu, Yiyuan She, Liangjing Ding, and Gary Gramajo, "Feature selection with annealing for computer vision and big data learning," *IEEE Trans. on PAMI*, vol. 39, no. 2, pp. 272–286, 2017.