

Automatic guide-wire detection for neurointerventions using low-rank sparse matrix decomposition and denoising

Markus Zweng^{1,3}, Pascal Fallavolita¹, Stefanie Demirci¹, Markus Kowarschik³,
Nassir Navab^{1,2}, Diana Mateus^{1,4}

¹ Computer Aided Medical Procedures, Technische Universität München, Germany

² Computer Aided Medical Procedures, Johns Hopkins University, USA

³ Siemens Corporate Research, Germany

⁴ Institute of Computational Biology, Helmholtz Zentrum München, Germany.

Abstract. In neuro-interventional surgeries, physicians rely on fluoroscopic video sequences to guide tools through the vascular system to the region of interest. Due to the low signal-to-noise ratio of low-dose images and the presence of many line-like structures in the brain, the guide-wire and other tools are difficult to see. In this work we propose an effective method to detect guide-wires in fluoroscopic videos that aims at enhancing the visualization for better intervention guidance. In contrast to prior work, we do not rely on a specific modeling of the catheter (*e.g.* shape, intensity, *etc.*), nor on prior statistical learning. Instead, we base our approach on motion cues by making use of recent advances in low-rank and sparse matrix decomposition, which we then combine with denoising. An evaluation on 651 X-ray images from 5 patient shows that our guide-wire tip detection is precise and within clinical tolerance for guide-wire inter-frame motions as high as 6mm.

1 Introduction

Neuro-interventional procedures are minimally invasive surgeries designed for the treatment of pathologies in the cerebrovascular system. Examples of neuro-interventions are endovascular aneurysm coiling, embolization of fistulae, or for stenosis, intracranial angioplasty and stenting. During such procedures surgeons insert a guide-wire through the patient's femoral artery in order to navigate the tools through the catheter up to the brain. The navigation is performed under image guidance, in particular angiography, where a sequence of fluoroscopic X-ray images show the motion of the interventional tools inside the patient at video frame-rate (typically 7.5 or 15 fps). Navigation is a complex task for several reasons: i) the sensible anatomy in the brain, ii) the low contrast, resolution and signal-to-noise ratio resulting from limiting the radiation dose, iii) the fact that X-ray images are only two-dimensional projections of three-dimensional structures, causing structure super-imposition and occlusions (*cf.* Fig 1-top left), and v) motion blur. In this work we aim at assisting the surgeon by automatically




Fig. 1. Example of results for one frame, cropped and scaled for clarity. **(top left)** original frame. **(top right)** results for the proposed method (RPCA / FNLN). **(bottom left)** Estimated foreground \mathbf{E} . **(bottom right)** Estimated foreground $\hat{\mathbf{E}} = \mathbf{E} < 0$.

detecting and visualizing the guide-wire (see Fig 1 top right). Such guide-wire detection could be further used in higher level visualizations, *e.g.* for 3D guidance using 3D models of the vessels acquired preoperatively with a rotational X-ray angiography [1, 3], which is our long-term goal.

Prior work addressing enhanced tool visualization for navigation are based either on a frame-wise detection or on tracking [8]. We focus hereon on detection techniques which are more suitable for a fully automated solution and can serve as input or reinitialization for tracking. Most of current detection techniques exploit the prior knowledge on the line like structure of the guidewire. For instance, Petkovic *et al.* [15] designed a filter to enhance line-like structures observed through X-Ray images, while Bismuth *et al.* [4], presented a curvilinear structure enhancement using a polygonal path image. These methods are general but may have problems detecting the tortuous guide-wire tips. Furthermore, they lead only to an enhanced image possibly containing other undesired structures. To address these issues, Lessard *et al.* [12] combined filtering with segmentation, whereas Honnorat *et al.* [9] relied on steerable filters regularized by tensor-voting. However, due to the difficulty of selecting the structures corresponding to the guidewire both methods still rely on the manual selection of

segmentation seeds for the former, or the guide-wire endpoints for the later, in order to initialize the algorithm. To overcome the difficulties above, there has been an increasing interest in automatic guide-wire detection by means of learning techniques, *e.g.* [2, 10, 14], which have shown promising results at the cost of a prior training stage and large amounts of manually labelled data needed in order to generalize to different image settings.

In this work we propose to detect tools in fluoroscopic sequences by relying on their motion, which is a less explored type of prior knowledge. A first approach in this direction was proposed by Spiegel *et al.* [16], where a simple background subtraction is applied using as mask an initial frame where no guide-wire is present. Although such detection is fast and easy to implement, it assumes a static background; so no patient motion, illumination variations nor image-parameter changes are allowed. Inspired by advances in sparse and low-rank decomposition of matrices we propose here a two step approach to detect the guidewire during neuro-interventional procedures. In the *first step*, we use a robust background subtraction method based on Robust PCA (RPCA) [6] to detect candidate guidewire regions. Several reasons motivate our choice for RPCA: (i) robustness to variations in image brightness, (ii) ability to handle a dynamic background, and (iii) capability of accounting for the partly-correlated noise present in X-ray sequences. In the *second step*, we investigate different methods to filter remaining noise in the foreground estimation. We evaluate our solution on clinical data and show its robustness for inter-frame guidewire motion of up to 6 mm.

2 Method

The input to our method is the fluoroscopic video used for guiding a neuro-intervention. We assume the main source of visible motion in the video is caused by the interventional tools, while the remaining part of the images will be close to static and considered as background. The tools can then be separated by subtracting the background from the current image of interest. To enforce robustness against subtle motions and imaging condition changes, the background is estimated w.r.t. the previous k frames in a sliding window fashion. More formally, we denote the video frames by $\mathbf{f}_i \in \mathbb{R}^s$, $i \in \{1, \dots, N\}$, with N the length of the video and $s = m \times n$ the size of the image. For each frame of interest \mathbf{f}_i , the goal is to obtain a binary mask $\mathbf{M}_i \in \{0, 1\}^s$ indicating which pixels belongs to the guidewire. Our method is composed of two steps. First, for each time-window Robust Principal Component Analysis (RPCA) [6] is applied to detect candidate guidewire regions for the current frame (*cf.* Sec 2.1). Then, the remaining noise is filtered out (*cf.* Sec. 2.2).

2.1 Low-rank background subtraction via RPCA

To estimate the background of current image \mathbf{f}_i , we consider a time window including the last k frames and stack them in the columns of the data matrix

$\mathbf{D} = [\mathbf{f}_{i-k}, \dots, \mathbf{f}_i]$, where $\mathbf{D} \in \mathbb{R}^{s \times k}$. Given the neuro-interventional sequences have video frame-rates, it is expected that the contiguous frames composing \mathbf{D} are highly correlated. If the correlation is linear, then the matrix \mathbf{D} is low-rank and it is possible to estimate the background masks by finding a rank- q matrix \mathbf{L} approximating the data matrix \mathbf{D} , *i.e.* by optimizing $\min_{\mathbf{L}} \|\mathbf{D} - \mathbf{L}\|$ s.t. $\text{rank}(\mathbf{L}) \leq q$, where $q \leq \text{rank}(\mathbf{D})$. This minimization can be solved using ordinary PCA over the data matrix \mathbf{D} and retaining only the most significant components to build \mathbf{L} .

However, moving parts primary caused by the guide-wire, would cause errors when using linear methods for estimating a low-rank representation. To handle such cases, where gross but sparse corruptions are present, Candès *et al.* [6] proposed the following Robust PCA method. Assuming the data matrix \mathbf{D} can be modeled as the sum of a low rank matrix \mathbf{L} and some additive sparse error \mathbf{E} , *i.e.* $\mathbf{D} = \mathbf{L} + \mathbf{E}$, RPCA is able to exactly recover⁵ \mathbf{L} by solving the following convex optimization problem:

$$\min_{\mathbf{L}, \mathbf{E}} \|\mathbf{L}\|_* + \gamma \|\mathbf{E}\|_1 \text{ s.t. } \mathbf{D} = \mathbf{L} + \mathbf{E}, \quad (1)$$

where $\|\cdot\|_*$ denotes the nuclear norm (enforcing low-rank), $\|\cdot\|_1$ is the L1 norm modeling the sparsity, and γ is a positive weighting parameter controlling the sparsity (of the foreground). From the different solvers for Eq. 1 we use the inexact augmented Lagrangian multiplier method (iALM) [13]. After solution, the low rank component \mathbf{L} gives us an estimate of the background, while the moving parts of the image and the noise will be encoded in the noise component \mathbf{E} . Actually, the foreground estimate for the current frame corresponds to the (matrix version of the) last column of \mathbf{E} , which we denote here \mathbf{E}_k . More particularly we are interested in its negative range, which contains the motion information of the dark structures of the interventional tool. To ease later notation, let us additionally denote:

$$\hat{\mathbf{E}}(\mathbf{x}) = \begin{cases} |\mathbf{E}_k(\mathbf{x})|, & \forall \mathbf{E}_k(\mathbf{x}) < 0 \\ 0 & \text{otherwise} \end{cases}$$

where $\mathbf{x} \in \Omega$ is a pixel taken from the image domain $\Omega \subset \mathbb{R}^2$. Finally, $\hat{\mathbf{E}}$ is normalized to lie in the range $[0, 1]$, such that intensities belonging to the guide-wire are close to 1. After this first step, the current foreground estimate $\hat{\mathbf{E}}$ will still contain noise that we target to remove in the second stage of our method described next.

2.2 Denoising

Since our low-rank background estimation can only cover up noise that is correlated between frames we need a further step to remove remaining artifacts. To do so, we compare the following four different methods.

⁵ As long as the error matrix \mathbf{E} is sufficiently sparse w.r.t. the rank of \mathbf{L}

Thresholding and Cluster Removal (TCR). This method, chosen for its simplicity and computational efficiency, consists of a histogram-based thresholding heuristic followed by the removal of small connected components. First, the discrete histogram of intensities for the current frame is estimated. We assume that the catheter information is in the tail of the histogram (as most pixels belonging to background have values closer to 1 after the normalization and dominate the distribution). Therefore a threshold t_{TCR} was heuristically determined by finding the position h_{tail} of the first bin exceeding a minimum of counts and defining $t_{\text{TCR}} = h_{\text{tail}} + \frac{1}{3}(1 - h_{\text{tail}})$. Only the values below t_{TCR} are retained. Finally, each pixel \mathbf{x} is assigned to a connected component $\omega_{\mathbf{x}}$ and components with fewer than c_{TCR} elements are removed. For a pixel $\mathbf{x} \in \Omega$ this results in:

$$\mathbf{M}_{\text{TCR}}(\mathbf{x}) = \begin{cases} 1 & \forall \left(\hat{\mathbf{E}}(\mathbf{x}) \leq t_{\text{TCR}} \right) \wedge (\#\omega_{\mathbf{x}} \geq c_{\text{TCR}}) \\ 0 & \text{otherwise} \end{cases}$$

Hessian-based filtering (HES). Hessian-based methods are popular for enhancing lines for guide-wire tracking and detection [4]. We used the approach presented in [16], computing the eigenvalues ($|\lambda_1| \leq |\lambda_2|$) of the image Hessian at each pixel, and determining that a line like structure is present when the following constraints are true: $|\lambda_1| \approx 0$, $|\lambda_1| \leq |\lambda_2|$, and $|\lambda_2| > t_{\text{HES}}$, where t_{HES} denotes a user-specified threshold. The Hessian is computed on the basis of image second derivatives over smoothed images, where a smoothness parameter σ_{HES} determines the scale of the used Gaussian kernel. A search is done over a range of scales, $\mathbf{r}_{\text{HES}} = [r_{\text{min}}, r_g]$, where r_{min} is a small value and r_g is an estimate of the guide-wire radius. At the end, the guidewire mask for a pixel is computed for the optimal scale $\mathbf{r}_{\text{HES}}^*$ leading to the highest vesselness value λ_2^* . In sum,

$$\mathbf{M}_{\text{HES}}(\mathbf{x}) = \begin{cases} 1 & \text{if } |\lambda_1^*| \approx 0, |\lambda_1^*| \leq |\lambda_2^*|, \text{ and } |\lambda_2^*| > t_{\text{HES}} \\ 0 & \text{otherwise} \end{cases}$$

Fast non-local means filtering (FNLM). Non-local means (NLM) filtering methods, introduced by Buades *et al.* [5], are known to achieve good denoising while preserving textures and fine structures, even for high noise levels. Instead of depending on the guidewire width as the HES method, NLM depends on an estimate of the standard deviation of the image noise σ_{FNLM} . Given the noisy foreground estimation $\hat{\mathbf{E}}$, the denoised value for pixel \mathbf{x} is computed as a weighted average over a search region $R_{\mathbf{x}} \subset \Omega$ around \mathbf{x} :

$$\hat{\mathbf{E}}_{\text{FNLM}}(\mathbf{x}) = \sum_{\mathbf{y} \in R_{\mathbf{x}}} w(\mathbf{x}, \mathbf{y}) \hat{\mathbf{E}}(\mathbf{y}) \quad (2)$$

The weights $w(\mathbf{x}, \mathbf{y}) = -\exp \frac{\sum_{\mathbf{t} \in \{0, \dots, p\}^2} |\mathbf{f}(\mathbf{x}+\mathbf{t}) - \mathbf{f}(\mathbf{y}+\mathbf{t})|^2}{h}$, with a h as a filtering parameter reflect the similarity of two patches of size $p \times p$ centered at \mathbf{x} and \mathbf{y} . This definition enforces that higher weights are given to pixels lying on patches with similar structure to the current one. In this work we used a fast variant

of NLM (FNML) that allows for an efficient computation of the weights [7]. The final mask \mathbf{M}_{FNLM} is computed by thresholding the denoised image with a user-defined parameter t_{FNLM} .

Markov-Random-Field (MRF). Inspired by the work in [17] we use a MRF to enforce the guide-wire continuity. A graph is built, where the vertices represent the pixels of the image $\mathcal{V} = \{\mathbf{x}_1, \dots, \mathbf{x}_s\}$ and edges $\mathcal{E} = \{e_{\mathbf{x}\mathbf{y}}\}$ link pairs of neighboring nodes, that is, $e_{\mathbf{x}\mathbf{y}}$ exists if \mathbf{y} lies within a square neighbourhood $\mathbf{y} \in \mathcal{N}_{\mathbf{x}}$ around \mathbf{x} . The binary foreground mask \mathbf{M}_{MRF} for the guidewire is obtained by optimizing the following energy according to the Ising model [11]:

$$\min_{\mathbf{M}_{\text{MRF}}} \sum_{\mathbf{x} \in \mathcal{V}} u(\mathbf{M}_{\text{MRF}}(\mathbf{x})) + \sum_{e_{\mathbf{x}\mathbf{y}} \in \mathcal{E}} \alpha |\mathbf{M}_{\text{MRF}}(\mathbf{x}) - \mathbf{M}_{\text{MRF}}(\mathbf{y})| \quad (3)$$

The first term describes the unary potential of a pixel being foreground or background. u is set to depend on the intensity value of the foreground estimation:

$$u(\mathbf{M}(k)) = \begin{cases} \beta \hat{\mathbf{E}}(k) & \text{if } \mathbf{M}(k) = 0 \\ (1 - \hat{\mathbf{E}}(k))^2 & \text{if } \mathbf{M}(k) = 1 \end{cases} \quad (4)$$

where the parameter β controls the balance between the two unary costs. The second term in Eq. 3 describes the pairwise potential imposing constraints on the mask value for neighboring pixels; it penalizes discontinuities and discouraging noise. The parameter α , controlling the strength of this neighboring dependency was chosen to be constant and Eq. 3 was minimized using graph cuts [11].

3 Experimental Validation

The validation was performed on fluoroscopic sequences of five patient datasets with a total of 651 frames of size $s = 512 \times 512$. The sequences were acquired with a frame rate of 15 fps. Two of the datasets were acquired using a biplanar X-ray system. The tip position and guide-wire centerlines were manually annotated by two experienced observers.

Implementation details

- For the background estimation we set the time window size to $k = 20$, which corresponds to 1.333s at the current frame-rate.
- The sparsity parameter of RPCA to $\gamma = 1/(\sqrt{\max(m, n)})$, the tolerance of the iALM solver to 10^{-4} , and its maximum iterations to 1000.
- For TCR, the number of histogram bins is set to 1000, and the threshold t_{TCR} is computed by finding the first histogram bin containing 250 elements. Also, the minimum allowed size for a connected component is $c_{\text{TCR}} = 8$.
- For HES, we set the scale range to $\mathbf{r}_{\text{HES}} = [0.3, r_g]$, with r_g given by the specifications of the guide-wire, and fix t_{HES} to the standard deviation of the foreground image $\sigma_{\hat{\mathbf{E}}}$.

- For FNLM, we set the filtering parameter as $h = 2\sigma_{\mathbf{E}}$, the patch size to $p = 11$, and the size of the search region $R_{\mathbf{x}}$ to 43×43 . h was also used as final threshold t_{FNLM} .
- Finally, for MRF, a parameter search was performed on a subset of the data. According to the F1-measure we set $0.06 \leq \alpha \leq 0.16$, and $8 \leq \beta \leq 10$.

Validation / Results Quantitative results of our evaluation are presented in Table 1. All of the method combinations show as desired high accuracies (ACC) and very low false positive rates (FPR). This is in part due to an imbalance of the positive and negative classes as there is a significant larger amount of pixels in the background (negatives) in comparison to those depicting the catheter (positives). Furthermore, the positive predicted values (PPV), not taking into account the negative class, are also high which is associated to a good precision. The apparently low TPR values are a consequence of the videos containing relatively long periods of no guide-wire motion, where our motion-based assumption does not hold. Indeed, the performance depends on the amount of motion, as well as how much of it is captured by the current time-window. We claim that during such static frames the navigation assistance is less relevant.

On the contrary, the most important for the surgeon is being able to see the moving tip, which has been a challenge for prior techniques. Therefore, we also measure the distances of the detected tips to the ground truth, which results in values of around $d_{\text{tip}} \approx 2.5\text{mm}$ for at least 3 of the methods. Moreover, if we focus on the most interesting frames containing tip motion, that is, considering frames where there ground-truth tip moves up to 6mm and removing frames for which the tip is static, then, the tip is accurately detected (e.g. $d_{\text{tip}}^* = 1.28 \pm 1.18$ mm). To better illustrate the discussion above we also plot in Fig. 2 the influence of the underlying motion on the estimated tip distance. The crosses on the vertical axis (tip motion = 0) indicate no motion of the ground truth tip w.r.t. the previous frame, leading to the RPCA not being able to detect the guidewire, and thus resulting in high tip distance estimates. On the other hand, errors in distance estimates also increase when the tip moves very fast (tip motion > 6mm), mainly due to motion blur. In these frames tips appear washed out, and have weak intensities that prevent their detection.

	ACC	PPV	FPR	TPR	$d_{\text{tip}} \pm \sigma_{\text{tip}}$	missed tips	$d_{\text{tip}}^* \pm \sigma_{\text{tip}}^*$
RPCA / TCR	99.30	70.65	0.15	16.16	2.03 ± 2.49	12.66	1.11 ± 1.20
RPCA / HES	99.38	88.11	0.05	13.33	2.54 ± 3.19	34.82	1.64 ± 1.77
RPCA / FNLM	99.37	88.71	0.06	13.78	2.24 ± 2.51	17.88	1.28 ± 1.18
RPCA / MRF	99.33	82.59	0.09	11.90	2.92 ± 2.57	28.68	1.49 ± 1.52

Table 1. Quantitative evaluation: True Positive Rate (TPR), False Positive Rate (FPR), Accuracy (ACC), Positive Predictive Value (PPV/Precision), avg. tip distance d_{tip} , missed tips within ROI (80×80 px), d_{tip}^* avg. tip distance for the cases where the ground truth tip moved up to 6mm between the last and the current frame (see Fig. 2). Distances are given in mm, other values in percentages.

Regarding the comparative performance of the different denoising techniques, we recommend FNLM given its good accuracy (ACC) and precision (PPV). Even if HES gives similar results and has less artifacts (FPR) it has twice as many missed tips than FNLM. Additionally FNLM has the advantage of being independent of the tool width. While the performance of FNLM in terms of tip distance is slightly worse than that of TCR, FNLM produces less than half of the outliers (FPR) compared to TCR, which is for the application more relevant. In case of tension stress on the guide-wire, fast motions can occur resulting in artifacts showing a second line (Fig. 1-bottom-left). Note that our method is able to differentiate such motion (see full foreground estimation Fig. 1-bottom-right) and correctly detect the guide-wire, whereas other detection methods would produce false positives or need further processing.

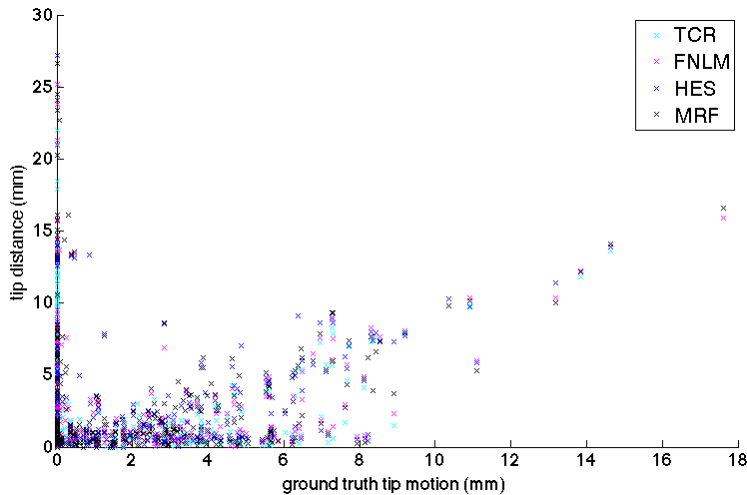


Fig. 2. Estimated tip distance vs. ground truth tip motion. The ground-tip motion is measured as the Euclidean distance of the ground-truth tip position w.r.t. the previous frame. Static tips corresponding to ground-tip motion of 0 are not well captured by our approach, as shown by the estimated distances variate along the vertical axis. This is to be expected as our method relies on motion cues. Better estimates are obtained for the four different methods, when the tool effectively moves.

4 Conclusion

Navigating the guide-wire through the vessels and positioning its tip at the operation site is a crucial step in neuro-interventional procedures. Guide-wire tracking and detection methods have focused on the delineation of the tip or the whole guide-wire. Our approach concentrates instead on the motion information obtained from the last k frames, is independent of the interventional tool being

used, and provides valuable information to the surgeon as during navigation the motion of the tip and changes in tension on vessels depict the most important pieces of information that could be overlaid.

One current limitation of the method is in detecting parts of the guide-wire that are static. This simply because our main assumption to detect the tool is that the tool moves. Although we have discussed that in a real application this assumption will be probably enough for navigation, possible ways to overcome it would be by combining our approach with an appearance model, in order to complete continuous line structures.

To conclude, we have proposed a robust guidewire detection approach that takes advantage of state-of-the-art developments in denoising and low-rank and sparse decomposition. Our work introduces a novel perspective for road-mapping in neuro-interventions that considers motion as the most important cue for tool detection. The method is precise, fully automatic, and does not require specific models of the tools, nor a training stage. It uses the image statistics over a sliding time-window of a given number of frames, thus being robust to rigid patient motion or imaging parameter changes. Finally, the algorithm requires few parameters and can be implemented to eventually achieve real-time performance. In sum, our method deals with important requirements for its real application in interventional neuroradiology.

References

1. S.A.M. Baert, E.B. van der Kraats, and W.J. Niessen. Three-dimensional guide-wire reconstruction from biplane image sequences for integrated display in 3-D vasculature. *IEEE Transactions on Medical Imaging*, 22(10):1252–8, October 2003.
2. A. Barbu, V. Athitsos, B. Georgescu, S. Böhm, P. Durlak, and D. Comaniciu. Hierarchical Learning of Curves Application to Guidewire Localization in Fluoroscopy. In *IEEE Comp. Vision and Pattern Recog. (CVPR)*, 2007.
3. C. Baur, F. Milletari, V. Belagiannis, N. Navab, and P. Fallavollita. Automatic 3d reconstruction of electrophysiology catheters from two-view monoplane c-arm image sequences. In *Int. Conf. on Information Processing in Computer-Assisted Interventions*, 2015.
4. V. Bismuth, L. Vancamberg, and S. Gorges. A comparison of line enhancement techniques: applications to guide-wire detection and respiratory motion tracking. *SPIE Conference Series*, 7529, 2009.
5. A. Buades, B. Coll, and J.M. Morel. A non-local algorithm for image denoising. In *IEEE Comp. Vision and Pattern Recog. (CVPR)*, 2005.
6. E.J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, June 2011.
7. J. Darbon, A. Cunha, T. Chan, S. Osher, and G. Jensen. Fast nonlocal filtering applied to electron cryomicroscopy. In *IEEE Int. Conf. on Biomedical imaging: from nano to Macro*, pages 1331–1334, 2008.
8. H. Heibel, B. Glocker, M. Groher, N. Paragios, N. Komodakis, and N. Navab. Discrete tracking of parametrized curves. In *IEEE Comp. Vision and Pattern Recog. (CVPR)*, 2009.
9. N. Honnorat, R. Vaillant, and N. Paragios. Graph-Based Geometric-Iconic Guide-Wire Tracking. In *MICCAI*, volume 6891, pages 9–16. Springer, 2011.

10. Nicolas Honnorat, Régis Vaillant, and Nikos Paragios. Robust guidewire segmentation through boosting, clustering and linear programming. In *IEEE Int. Conf. on Biomedical imaging: from nano to Macro*, pages 924–927, 2010.
11. V. Kolmogorov and R. Zabini. What energy functions can be minimized via graph cuts? *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 26(2):147–159, Feb 2004.
12. S. Lessard, C. Lau, G. Roy, D. and Soulez, and J.A. de Guise. Wires segmentation in fluoroscopic images during cerebral aneurysm endovascular intervention. In *IEEE Int. Conf. on Biomedical imaging: from nano to Macro*. IEEE, 2008.
13. Z. Lin, M. Chen, and Y. Ma. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. *Mathematical Programming*, 2010.
14. F. Milletari, V. Belagiannis, N. Navab, and Pascal Fallavollita. Fully automatic catheter localization in c-arm images using l1-sparse coding. In *MICCAI*, 2014.
15. T. Petković and S. Lončarič. Guidewire tracking with projected thickness estimation. In *IEEE Int. Conf. on Biomedical imaging: from nano to Macro*, pages 1253–1256, 2010.
16. M. Spiegel, M. Pfister, D. Hahn, V. Daum, J. Hornegger, T. Struffert, and A. Dörfler. Towards real-time guidewire detection and tracking in the field of neuroradiology. *SPIE Conference Series*, pages 726105–726108, 2009.
17. X. Zhou, C. Yang, and W. Yu. Moving object detection by detecting contiguous outliers in the low-rank representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(3):597–610, March 2013.