

Curvilinear structures segmentation in fluoroscopic images

Nicolas Honnorat – Régis Vaillant – Nikos Paragios

N° 7414

October 2010



*Rapport
de recherche*

Curvilinear structures segmentation in fluoroscopic images

Nicolas Honnorat ^{*†‡} – Régis Vaillant[‡] – Nikos Paragios ^{*†}

Theme :
Équipe-Projet Galen

Rapport de recherche n° 7414 — October 2010 — 21 pages

Abstract: Fluoroscopic imaging provides means to assess the motion of the internal structures and therefore is of great use during surgery. In this report we propose a novel bottom-up approach for the segmentation of curvilinear structures in these images.

The main challenge to be addressed is the lack of visual support due to the low SNR where traditional ridge-based methods fail. Our approach combines machine learning techniques, unsupervised clustering and linear programming / discrete optimization. In particular, numerous invariant to position/rotation classifiers are combined to detect candidate pixels of curvilinear structures. These candidates are grouped into consistent geometric segments through the use of a state-of-the-art unsupervised clustering algorithm. The complete curvilinear structure is obtained through a linking of these segments using the elastica model in a linear programming framework in a first variant, and using discrete optimization in a second one. Very promising results were obtained on angioplasty guide wire segmentation in cardiac interventional images.

Key-words: Boosting, clustering, Linear Programming, local search, ordering, steerable filters, linear structures, fluoroscopic images, guide wire

* Laboratoire de Mathématiques Appliquées aux Systèmes, École Centrale de Paris, France

† Equipe GALEN, INRIA Saclay - Île de France, Orsay, France

‡ General Electric Healthcare, Buc, France

Segmentation de structures linéaires dans les images fluoroscopiques

Résumé : La fluoroscopie, qui permet de visualiser les mouvements des organes et des outils chirurgicaux à l'intérieur des patients, est très importante en chirurgie. Dans cet rapport, nous proposons une nouvelle approche bottom-up permettant de segmenter les structures linéaires dans les images fluoroscopiques. Le principal défi réside dans le niveau de bruit important de ces images, qui tient en échec les techniques de détection de lignes traditionnelles. Notre approche combine des techniques de Machine Learning (apprentissage) une méthode de clustering non supervisée et des techniques de programmation linéaire et d'optimisation discrète. Plus précisément : de nombreux détecteurs invariants par translation/rotation sont combinés pour détecter les pixels susceptibles d'appartenir aux structures curvilinéaires ; ces candidats sont regroupés en segments cohérents au moyen d'une technique de clustering non supervisée ; et les structures entières sont finalement obtenues en reliant les segments en se basant sur des critères comme l'Elastica, les liaisons optimales étant obtenues en résolvant un problème de programmation linéaire dans le cas de la première variante que nous proposons, et en résolvant un problème d'optimisation discrète pour une seconde variante. Des résultats très prometteurs ont été obtenus pour la segmentation des guides utilisés en angioplastie cardiaque.

Mots-clés : Boosting, clustering, Programmation linéaire, optimisation discrète, filtres steerable, structures linéaires, images fluoroscopiques, guide

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 2 | Robust curvilinear structures segmentation | 5 |
| 2.1 | Low-level detection | 5 |
| 2.2 | Clustering towards Local Segments Detection | 6 |
| 2.3 | Structure delineation as a Linear Program | 7 |
| 2.4 | Experimental validation | 9 |
| 2.5 | Discussion | 10 |
| 3 | Extraction of a single Guide Wire | 11 |
| 3.1 | Improving the detection | 11 |
| 3.2 | Segment ordering | 11 |
| 3.2.1 | Image-based costs provided by Fast Marching | 12 |
| 3.2.2 | geometric costs | 13 |
| 3.2.3 | ordering with outlier rejection | 13 |
| 3.3 | Experimental Validation | 14 |
| 3.4 | A simple improvement | 16 |
| 4 | Discussion | 16 |

1 Introduction

Curvilinear structures extraction is a well studied problem in medical imaging as well as in computer vision. Image segmentation of open-structures, vessel segmentation or tree-extraction of the retina are examples where these methods are of great interest.

This task remains challenging when sought structures are ridges and not contours (because context helps poorly) or when the available image support is not sufficient even for humans for the automatic delineation of these structures. Segmentation of devices used for electrophysiology or cardiovascular interventions is a typical example where one has to deal with such a situation: one can cite thin structures, as well as low and dramatically variable SNR, making part of the structures of interest (guide-wire) invisible.

Conventional edge-detection methods often considered in the field of computer vision fail short in this context. Existing methodologies are based on a two-stage procedures: the first step aims to detect candidate pixels and the second focuses on the complete recovery of the structure of interest.

The detection part can be considered either with or without preprocessing. Pre-processing aims to improve the discriminability with respect to the background of these structures. It can be achieved either through filtering using dedicated operators [1] or coherence enhancing diffusion [2]. One can also imagine filtering a feature space being derived from these images like for example the method proposed in [3] that is using Tensor Voting.

State of the art with respect to detection includes the use of conventional edge-detection methods [1], Hessian matrix [2] and the consideration of particular class of image operators being able to enhance these structures [4, 3] like for example steerable filters. The use of vesselness [5] that was considered in the context of arteries detection has been also considered [6] as well as *phase congruency* [7] and dedicated filter banks [8]. More recently, boosting was introduced in this context using Haar features as classification space [9]. Invariance with respect to the orientation of these structure is the main limitation of the method since numerous filters are required to capture all possible orientations.

Once the detection step is completed, the next step consists of linking them. This can be done naturally by chaining methods as suggested in [1, 4, 3, 6]. The use of snakes-splines on the feature images has been also investigated [2, 7] in the case of single structure. These methods are less interesting for multiple curvilinear structure detection (and when numerous confusing structures are present). More recently the idea of hierarchical grouping guided by a classification at each stage was considered, that is also a very flexible framework to introduce prior knowledge.

In this report, we introduce a novel approach combining detection using boosting [10], with a mid-level grouping scheme based on clustering and complete reconstruction through optimization. In comparison with other methods which are prone to suffer from local minima [2, 9], which do not reconstruct undetected portions of guide wire [1, 4] or which are greedy [8] our approach is the first one designed to achieve robustness at every level.

The reminder of this report is organized as follows: in section 2 we discuss a first approach solving the reconstruction as a linear program and present the detection and clustering part of our approach ; providing experimental results and a discussion of the limitations of the method. The next section describes a

second approach suited for single structure extraction achieving better results. Discussion concludes the report.

2 Robust curvilinear structures segmentation

Our algorithm consists of three stages: a low-level stage detecting pixels belonging to the structures of interest, a grouping stage extracting parts of the structures from the point cloud obtained, and a reconstruction stage linking these primitives together.

2.1 Low-level detection

AdaBoost [10] is an efficient classification method that is based on the following principle: select and linearly combine weak classifiers towards the creation of a strong one.

The use of such a learning framework requires the definition of a set of operators and a training set with negative and positive examples. We have adopted a feature space that is derived from the class of separable steerable filters proposed in [11].

These operators have been designed to be optimal in the sense of Canny-like criteria and are well suited for ridges and edges detection allowing clear, well localized and without echo detections.

Despite certain improvement thanks to the orientation invariant approach, still the method suffers from numerous outliers detections as shown in [Fig. (2, 4)].

This can be partially addressed through the use of a robust to outliers version of boosting, namely the Gentle AdaBoost algorithm [12]. Because our learning sets are very inhomogeneous (GW pixels represent less than 0.25% of whole images) we replaced the re-weighting formula of Gentle AdaBoost by the following:

$$w_j \leftarrow w_j e^{-y_j h(x_j)} e^{y_j (asymmetry)}$$

(where w_j are the weights of the data x_j of label y_j and h is the weak classifier chosen during the current AdaBoost step)

in order to put more emphasis on samples of the structures like in asymmetric AdaBoost [13] ; and we fixed the asymmetry parameter by validation to get the best possible classifier.

If M is the number of weak classifier combined, we have to shift the classifier output by $-M(asymmetry)$ because our asymmetric Gentle Adaboost finally try to minimize :

$$E \left(e^{-y_i (F(x) - M(asymmetry))} \right)$$

[Fig. (1)] shows the output of our classifier applied on whole images. The use of such a powerful classification technique along with an appropriate orientation invariant filter bank improves the detection rate but still suffer from numerous erroneous detections.




Figure 1: (i) input image (ii) output of the classifier

Curvilinear structure extraction from such an information space where the outliers are more numerous than the true detections is almost impossible. One can overcome this limitation through a local clustering of points towards forming individual line segments with proper handling of outliers. It is natural to assume that errors are not equally distributed and several clusters will correspond to inconsistent with weak support line segments. The main challenges of such a process is that one has to deal with numerous outliers, and the fact that the number of segments/clusters is unknown in advance.

2.2 Clustering towards Local Segments Detection

Let us consider the N points $p_i = (x_i, y_i)$ being detected during the classification process. Such a segmentation can be formulated as follows: find a set of cluster centers $c_i = (cx_i, cy_i)$ (the cardinality of this set is unknown and denoted as $|C|$), and a global labeling \mathcal{L} (individual labels coming from the set $1, |C|$) for all detected points p_i such that if all points with the same label are considered, then they form a local line segment. This can be cast as a minimization process, of the following nature:

$$\min_{c_i, |C|, \mathcal{L}} \sum_{i=1}^{|C|} f(c_i) + \beta \sum_{j=1}^N \delta(\mathcal{L}(j) - i) g(c_i, p_j) \quad (1)$$

where f is the penalty of a candidate point to be a cluster center, $\mathcal{L}(j)$ is the label indicating the optimal cluster for the point j and $g(c_i, p_j)$ is the cost of associating p_j to c_i . It is obvious that the $\delta(\mathcal{L}(j) - i)$ will only measure the cluster cost with respect to the optimal attribution label. The minimization of this objective function with respect to the number, centers and cluster labels of the set of detected points will provide a set of local segments of the curvilinear

structure. The definition of the functions f and g , as well as the optimization of the objective function are the main challenges to be addressed.

We consider f to be a constant and no-particular preference is given among the set of detected points. In order to define the cost of attributing a point p_j to the cluster c_i , particular attention is to be paid. We consider the following: (i) (c_i, p_j) are forming a line segment l_{cp} , (ii) c_i and the axis of orientation given by the optimal steerable filter form a line l_c , and (iii) p_j and the optimal steerable filter axis do also form a line l_p . In the ideal case, if the point p_j belongs to the cluster c_i , then the three line segments will coincide.

We consider two criteria: (i) The actual geometric distance between the two points, $d(c_i, p_j)$ since we would expect capturing local straight line segments of curvilinear structures. (ii) The distances $d_1(c_i, p_j) = d(c_i, p_j) \max(\tan(\theta_1), t_{max})$ and $d_2(c_i, p_j) = d(c_i, p_j) \max(\tan(\theta_2), t_{max})$ where θ_1 and θ_2 are the angles between the line being formed from the two points and the ones given by steerable filters orientations. These distances are combined to a single metric as follows:

$$g(c_i, p_j) = d^2(c_i, p_j) + \alpha(d_1^2(c_i, p_j) + d_2^2(c_i, p_j)) \quad (2)$$

The optimization of this objective function is done through a graph clustering method able to found by itself the optimal number of clusters. We refer to the reader to what is so-called clustering via LP-based stabilities introduced [14]. Experimental results of this grouping are shown in [Fig. (2)].

The output of the clustering process will provide collections of points corresponding to local line segments. One now has to order these segments towards forming the entire structure. First, from the subgroup of points, local line segments are detected in a robust manner. We adopt the MSAC algorithm finding the best line to fit to each cluster - and kept only the part of this line where most of the cluster points project themselves. This variant of RANSAC presented for first time in [15] (simply using skipped-mean loss function instead of zero-one loss function) is indeed known to perform better in practice. Some experimental results of this clustering are shown in [Fig. (2)]. The next and final step of the process consists of parsing these line segments toward forming a consistent curvilinear structure. Three challenges are to be addressed, the existence of segments corresponding to noise, the miss-detections due low contrast and the case of segments corresponding to other curvilinear structures than the guide. We will focus on the first two aspects.

2.3 Structure delineation as a Linear Program

Let us now consider the problem of forming a continuous curvilinear structure out of these segments.

Let (s_m^1, s_m^2) be the extremities of the line segment s_m , and consider the extended set of tips

$$\mathcal{Q} = \{s_1^1, s_1^2, \dots, s_m^1, s_m^2, \dots, s_{|C|}^1, s_{|C|}^2\}$$

Let consider the set of boolean labels $l_{u,v}$ - for u and v in \mathcal{Q} equal to 1 when extremities u and v have to be linked to reconstruct the structures of interest. Recovering the global curvilinear structure can be considered as an optimization problem, or:




Figure 2: (i) clustering results and (ii) corresponding line segments

$$\begin{aligned} \min \quad & \sum_{u,v \in \mathcal{Q} \times \mathcal{Q}} l_{u,v} h(u,v) + (1 - l_{u,v}) \epsilon \\ \text{s.t.} \quad & \begin{cases} \sum_v l_{u,v} \leq 1 & \forall u \in \mathcal{Q} \\ l_{u,v} \in \{0, 1\} & \forall u, v \in \mathcal{Q} \times \mathcal{Q} \end{cases} \end{aligned}$$

where $h(u, v)$ is a cost for linking extremities u and v , ϵ penalizes not linked pairs of segment extremities and constraints prevent extremities being linked twice. This problem can be solved using a standard LP solver.

In our case, the function $h(u, v)$ was the outcome of a classifier that was also trained. The feature space of this classifier was a combination of geometric and photometric features. In terms of geometry, we considered features derived from the euclidian distances between u and v , the corresponding segments centers, the square of these two distances and min and max lengths of the segments. High level geometric features are considered as well, like the Elastica defined in [15], which is computed between extremities or between segments centers, and multiplied or not by the corresponding distances and square distances - in order to introduce scale-invariant and scale-dependent curvature constraints. The photometric features were also considered to re-introduce low-level information in the process: min and max mean intensity along the segments and terms insuring likely links: the means of the output of the detection along cubic splines linking extremities or centers of the segments (and tangents to segments) - multiplied or not by corresponding distances and square distances to take the length of the cubic curve into account or not. The choice of such a classification strategy opposite to costs directly defined on the above spaces is mainly due to the parameter-free nature of the process, since the corresponding weights of the different measurements are learned thanks to boosting.

We finally fit splines robustly on linked segments: beginning with splines passing through segments tips and centers, we iteratively assigned detected pix-

els near from splines to their nearest splines and fit cubic-splines (penalizing the number of node of the spline) on their novel set of supporting points again. This iteratively robust fitting (based on a zero-one loss function) allowed us to get slightly better splines, by re-introducing low level information.

2.4 Experimental validation

Experiments were carried out on a base of 15 sequences of 10 images 1000x1000 acquired during interventions on 13 different patients. Clinical experts have manually marked guide wires and guide wire tips (made of a different and more opaque to X-ray material). We took 75 images in the training set almost equally distributed in the 11 first sequences only. The testing database includes 28 images distributed in an equal manner across the data sets. The remaining 47 images were used for validation.

The validation metric was inspired from [9] and corresponds to the proportion of pixels of GW at more than 2 pixels of the detected structures (*missed detections*) and the proportion of pixels of the detected structures at more than 2 pixels of the GW (the *false detections*).

We first compared our classifier with steerable filters alone and with the filter bank proposed in [8]. We adopted classifiers combining separable steerable filters designed for ridge and edges detection [11] at different scales. The responses at 90° and 45° of the orientation giving the highest response as features for classification (in addition to the highest response of the filters) were considered. The classifier used here combine the ridge detector of order 2 for $\mu = 0$ and scale 1.0 and the ridge detectors of order 4 for $\mu = 0.25$ and scales 0.8, 1.6 and 3.0 (see [11] for details).

The best steerable ridge detector was the one of scale 1.6 and the best oriented filter was the one of length 30 and width 3. We present the corresponding ROC curves in [Fig. (3)]. Note that only results for small false positive rates are interesting because "negative" pixels are hundreds of times more numerous than "positives" ones.

We then evaluated the different stages of the method both in terms of missed detections and false detections. We have considered the top 5,000 detected points from a dense 5×5 grid, and the same set of parameters for the LP-Clustering and the MSAC was considered for all images. Applying these parts to training images, we built a set of segments pairs for training the classifier used for structure delineation. We finally fixed ϵ to get a good trade off between false detection and missed detections. [Fig. (4)] presents our results. Before fitting splines, we removed sequences of length smaller than 90 pixels.

It is expected that the method will fail to recover parts that have been widely ignored from the classifier. Our method links otherwise well detected part of structures correctly, losing few detected pixels. On the other hand we observe a significant improvement over false-detections: starting from a proportion of 1 to 5 between the the guide-wire detected portion and the false positives, this ratio is dropped down to 1.1 at the end of the process - and even 0.47 (68.2% of good detections for 31.8% of false detections) when static confusing structures are not taken into account. A significant part of these structures and of the remaining false detections could be removed using a detector of GW tip like in [9]. However, this was not within the scope of our approach.




Figure 3: Detection performances (on 5.000.000 points)

| stages | MD | FD | FD-I |
|-----------------|---------|---------|---------|
| detection | 25.78 % | 83.52 % | 66.46 % |
| segments | 27.02 % | 76.54 % | 57.51 % |
| linked segments | 25.91 % | 64.68 % | 44.32 % |
| splines | 27.66 % | 52.36 % | 31.80 % |

Figure 4: MD : *missed detections*, FD : *false detections* and FD-I : *false detections* when immobile structures are not taken into account

2.5 Discussion

This first work was presented in [16]. Despite the use of AdaBoost, false detections are too numerous (there are roughly 4 times more false detections than right detections) and most of them cannot be discarded during clustering or segment linking because the main confusing structures such as rib borders are curvilinear.

Most of the confusing structures are however static and GW are still in movement. Introducing features able to discriminate between static and moving structures could improve the low-level detection, as shown before.

Besides, though a global solving the decision for segment linking is mainly took locally: the last step will never link extremities u and v such that

$$h(u, v) > \epsilon$$

Because in practice some parts of the GW are invisible due to noise, we would prefer a linking procedure allowing a locally higher linking cost if this link bind two long parts of the structure together.

We thus decided to improve detection and develop an another linking procedure.

3 Extraction of a single Guide Wire

3.1 Improving the detection

There are mainly three kinds of confusing structures : rib borders, which are thin and little absorbent like GW body, staples which absorbent like GW tips (but most of the time slightly larger) and anatomical noise (narrow and noisy structures of aleatory size and orientation constituting most of the background of the fluoroscopic images).

Fortunately, the most annoying structures (thin staples and rib borders) are static in most of the sequences, and GW always in movement (because of cardiac motion). Computing features based on the difference between the current (analyzed) frame and the last one allows therefore to discard them in many sequences. Rib borders are moreover regions of higher intensity variance ; and erroneous detections due to noise can be reduced when introducing features computed directly on the output image of the imaging system (fluoroscopic images are indeed corrupted by an additive gaussian noise which is significantly disformed when taking the log of the intensity to compute the image of X-ray absorption. Extracting features in the raw image can therefore help taking noise level into account, whereas absorption images are required to detect the structures).

In order to improve the low-level detection, we finally added three types of features to our set of features : variances computed on the difference image between the actual and the preceding frame in order to discard immobile structures (rib borders and staples), variances computed on the raw image (in order to discard remaining rib borders) and we computed several Differences of Gaussians in both raw image and absorption image to try reducing false detections due to anatomical noise.

By computing ROC curves with 100 000 points taken in the test images, we observed a significant improvement of the performances of the classifier : the last one exhibited an area under ROC curve restrained to false positive rates inferior to 1% of nearly 32 % whereas we get a classifier with an area of 46.5% by adding variances computed in squares of size 3 and 5 (in difference image and raw image) and three Differences of Gaussian (scale 1.8 minus scale 1.4, scale 4.0 minus scale 2.0 and scale 5.0 minus scale 4.0). As shown in [Fig. (5)] results are qualitatively clearly better (rib borders can now be discarded, barely visible structures like GW in catheters are now a little more detected...).

3.2 Segment ordering

The main drawback of the linking procedure presented in 2.3 is that the segment linking is decided locally, though a global resolution: when deciding to link two extremities, the algorithm does not advantage linking that would elongate a delineated structure or that would link two parts of a structure. As a result, this algorithm is always reluctant to connect distant extremities and does not link well the different detected parts of the GW.




Figure 5: output of the first classifier against output of the new one : rib borders which were detected more than GW body respond now less than GW. Moreover, the GW part in the catheter is now detected, which helps segmenting the beginning of the GW body

We thus developed another approach accepting locally higher costs if they lead a globally better optima: an algorithm finding both which line segment to reject and how to order the remaining one to get the GW. This approach is suited for the delineation of a single GW, but could be applied more than one time when necessary.

This novel algorithm requires a linking cost $C(a, b)$ for each pair of extremities a and b - like the preceding one 2.3. We computed it as a linear combination of an image-based cost estimated by Fast Marching and a geometric cost.

In the following subsections, we describe the costs we took and the way we performed ordering with outliers rejection.

3.2.1 Image-based costs provided by Fast Marching

As image-based cost for linking two line segment extremities a and b we computed the minus log-likelihood of the most likely path linking a and b , where likelihoods are derived from the classifier output.

Assuming that pixels along the path are independent, this criterion can be computed as a sum over the pixels of the better path, and can therefore be computed by fast marching. Denoting with $h(x)$ the detection score at location x and using a sigmoid function to compute pixel likelihood, we finally used the following cost map $c(x)$ as fast marching input:

$$c(x) = -h(x) + \log(e^{h(x)} + e^{-h(x)})$$

It is indeed well known that taking the sigmoid of 2 times the classifier output $h(x)$ for a sample x gives an estimation of the corresponding likelihood that x is a positive sample (this operation is sometimes called *logistic correction* and is directly inspired from the statistical view of boosting of Friedman et al. [12] as explain for example in [17]).




Figure 6: Minimal geodesic costs determined using the fast marching algorithm in different images of the data base. Colors allow to differentiate maps and cost are proportional to the color intensity

Linking costs above an arbitrary threshold were not computed and approximated with the product between the euclidian distance between the tips and the maximal linking cost normalized according to the euclidian distance between tips found by fast marching. [Fig. (6)] shows several maps computed with the algorithm presented in [18].

3.2.2 geometric costs

As geometric costs, we simply used the scale invariant Elastica criterion defined in [15] computed between segment center in order to get results robust with respect to small inter-crossing of segments. The idea was that a and b should be linked by a path of low curvature.

3.2.3 ordering with outlier rejection

Notations

Let s_n be the segment number n provided by the clustering stage. Its extremities are denoted with an and bn : $s_n = (an, bn)$.

We denote with ϕ_n a boolean variable equal to 0 if s_n is considered as an outlier. We introduce the boolean variables ω_n indicating in what sense the segment is part of the GW: 0 if an and bn appears in that order when parcouring the GW from its begin to its end, and 1 if bn appears before an .

Let θ_n be the number of the segment at place n in the GW of s_n if s_n is part of the GW, and θ_n^{-1} the index of the segment at place n in the delineated GW.

We define the costs

$$C_{\omega_m, \omega_n}(m, n) = \begin{cases} C(bm, an) & \text{if } \omega_m = 0 \text{ and } \omega_n = 0 \\ C(bm, bn) & \text{if } \omega_m = 0 \text{ and } \omega_n = 1 \\ C(am, an) & \text{if } \omega_m = 1 \text{ and } \omega_n = 0 \\ C(am, bn) & \text{if } \omega_m = 1 \text{ and } \omega_n = 1 \end{cases}$$

algorithm

Finding the optimal sequence of segments can be considered as an optimization problem. In simple words we solved the following minimization problem:

$$\min_{\theta_i, \omega_i, \phi_i} \sum_{m=1}^{|\mathcal{C}|} \sum_{n=1}^{|\mathcal{C}|} \phi_m \phi_n \delta(\theta_n - \theta_m - 1) C_{\omega_m, \omega_n}(m, n) + \beta \sum_{m=1}^{|\mathcal{C}|} (1 - \phi_m)$$

where product $\phi_m \phi_n$ guarantees that both s_m, s_n segments are retained and the term $\delta(\theta_n - \theta_m - 1)$ that the segment s_n succeeds s_m . The second term penalizes the attribution of the outlier label to a segment.

Starting from a configuration where all the segments are considered to be inliers and have been linked greedily (one segment has been arbitrary chosen to be the first one, and for all n , the segment chosen to be at place $n + 1$ is the nearest from the segment at place n that had not been chosen previously) our algorithm reduces the objective function iteratively until convergence. At each step, inliers and outliers are ordered, and all the possible permutations of a subsequence of inliers with a subsequence of outliers that reduce the objective function are applied.

The ordering algorithm we use for the first task try to find by local search the order minimizing the sum of the links made between successive segments. In other words, it tries to find the 'Shortest Spanning Path' among the segments by local search. The ordering of the outliers provides in this manner interesting subsequences for the permutation step without affecting the objective function, whereas the ordering of the inliers reduces the objective function. Given that all tasks reduce the objective function, convergence is guaranteed.

We remind that local search is an heuristic minimizing an energy depending on a configuration by passing from the current configuration to a neighbouring one while it allows reducing the energy. Particular attention is to be paid when choosing the neighborhood definition. In particular, both the speed of convergence and the quality of the local minimum found depends on its size.

We chose the neighborhood defined by all the configurations reached (from the current one) when applying one of the following operations: (i) reversion of a subsequence of segments (ii) shift of a subsequence (iii) reversion followed by a shift of a subsequence. These operations are illustrated in [Fig. (7)].

3.3 Experimental Validation

We performed first the following experiment : a classifier was trained using 7 images from the 10 first sequences and 5 from the next one, and the test was performed using 3 image per sequence (different from the one used for training).

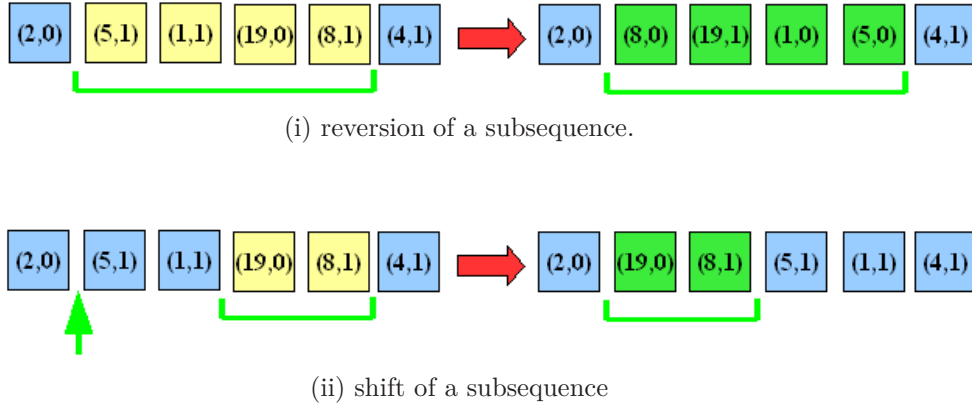


Figure 7: two operations being part of the neighborhood, described with couples $(\theta_n^{-1}, \omega_n)$. The reversion of a subsequence requires a modification of all the ω_n of the segments being part of the subsequence

We got the first results of the following tables.

In order to address the concern that some sequences provide images used for both training and test, we then performed a cross-validation : we trained 15 classifiers using 75 images from 14 sequences only (10 images from one sequence and 5 images from the others) we used this classifier to segment GW on 3 images of the sequence not used for training and we got the second list of results, which are quite similar to the first ones. Features and parameters for the training and the test were the same for the 15 runs : K was fixed to 30.0, 2800 points were selected in each image.

In both case, we took the Elastica criterion 4.5 times more into account than the image-based cost when linking segments. K was fixed around 15.0 in the first case. We used the same parameters for the MSAC : a scale equal to 3.0, a minimum number of points per cluster of 12, a shortest length of segments of 25.0 and a maximal cost for Fast Marching of 14.0.

| stages | MD | FD |
|---------------------|-------|-------|
| detection | 13.2% | 49.5% |
| segments extraction | 20.6% | 40.9% |
| segments linking | 23.2% | 8.7% |
| splines fitting | 22.2% | 7.6% |

missed detections(MD) and false detections(FD) after each step, computed with 45 test images (3 per sequence)

| stages | MD | FD |
|---------------------|-------|-------|
| detection | 15.4% | 53.2% |
| segments extraction | 19.3% | 53.0% |
| segments linking | 23.8% | 14.3% |
| splines fitting | 25.8% | 13.2% |

MD and FD after each step, computed with the same 45 test images (3 per sequence) obtained through cross-validation

This second work was presented in [19].

The three new types of features allow to build far better classifiers. Because segment linking discard now all the segments which are not part of the main structure, there remain now very few False Detections. The remaining Missed Detections are mainly due to a limited amount of images where the end of the GW body is not detected at all (due to the noise...). In these configuration, there are no FD, but only the first half of the GW is detected (which causes 40

% to 60% of MD). In the following section, we present a simple improvement that helps managing these configurations.

3.4 A simple improvement

We used our training algorithm with the same features and the same training database to build a classifier detecting only GW tips, and we defined the following measure for each line segment L :

$$m(L) = \sum_{x \in L} -\log(1 - p_{classifier}(x))$$

where $p_{classifier}$ is the *logistic correction* of the classifier output :

$$p_{classifier} = \frac{e^{h(x)}}{e^{h(x)} + e^{-h(x)}}$$

If we assume, like for the Fast Marching part, that pixels are independent, then this measure approximates the minus log likelihood of the segment not to be part of the GW tip.

By penalizing more (by a constant additive cost) the segments L whose $m(L)$ is above a fixed threshold, we got the same results than before for the images where the GW was correctly linked, while forcing the algorithm to jump until GW tip in the other sequences.

We finally obtained an average MD over the 45 test images near 18%, for thresholds around 20.0 and additive costs around 30.0.

4 Discussion

In this report, we have presented a bottom-up approach for thin linear structures segmentation designed to achieve high robustness.

We used classifiers combining responses of improved steerable filters [11] with variances and feature allowing to discriminate moving structures from static ones ; and that are built by a variant of AdaBoost [10] to perform low-level detection. Our grouping scheme works in two times: pixels likely to be part of the same structure segments are grouped together by a clustering method, and then segments are robustly fitted to the groups by MSAC [20]. Finally, we formalized the reconstruction of complete structures from set of line segments either as an LP problem or as an ordering problem ; and we robustly fitted cubic-splines on the sequence of line segments found. [Fig. (8)] summarizes the operations of our algorithm.

Futur work will focus on GW tracking. Knowing the GW movement is indeed of great interest for many applications, such as cardiac motion monitoring.

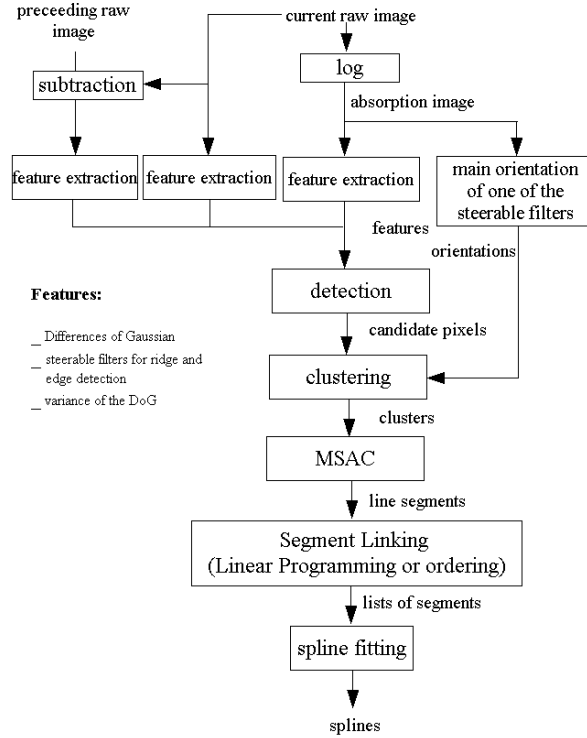


Figure 8: operations performed by our algorithm

Acknowledgments

This work was supported by ANRT (grant 1062/2008) and GE Healthcare. The authors thank N. Komodakis for providing the implementation of the clustering method




Figure 9: sequences of linked segments provided by the first method

References

- [1] D. Palti-Wasserman, A. M. Bruckstein, and R. P. Beyar, "Identifying and tracking a guide wire in the coronary arteries during angioplasty from x-ray images," *IEEE Tr. on Bio. Eng.*, vol. 44, no. 2, pp. 152–164, feb. 1997.
- [2] S. A. M. Baert, M. A. Viergever, and W. J. Niessen, "Guide-wire tracking during endovascular interventions," *IEEE Tr. on Medical Imaging*, vol. 22, no. 8, pp. 965–972, aug. 2003.
- [3] E. Franken, P. Rongen, M. van Almsick, and B. M. ter Haar Romeny, "Detection of electrophysiology catheters in noisy fluoroscopy images," *MICCAI*, pp. 25–32, 2006.
- [4] S. Lessard, P. Bigras, C. Lau, D. Roy, G. Soulez, and J. A. de Guise, "Automatically driven vector for guidewire segmentation in 2d and biplane fluoroscopy," *World Academy of Science, Engineering and Technology*, vol. 54, pp. 446–450, 2009.
- [5] Alejandro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever, "Multiscale vessel enhancement filtering," in *MICCAI '98*, London, UK, 1998, pp. 130–137, Springer-Verlag.
- [6] T. Hauke Heibel, B. Glocker, M. Groher, N. Paragios, N. Komodakis, and N. Navab, "Discrete tracking of parametrized curves," *CVPR*, pp. 1754–1761, 2009.
- [7] G. Slabaugh, K. Kong, G. Unal, and T. Fang, "Variational guidewire tracking using phase congruency," *MICCAI*, vol. 10, pp. 612–619, 2007.
- [8] V. Bismuth, L. Vancamberg, and S. Gorges, "A comparison of line enhancement techniques: applications to guide-wire detection and respiratory motion tracking," in *SPIE*, Feb. 2009, vol. 7259.
- [9] A. Barbu, V. Athitsos, B. Georgescu, S. Boehm, P. Durlak, and D. Comaniciu, "Hierarchical learning of curves application to guidewire localization in fluoroscopy," *CVPR*, pp. 1–8, 2007.
- [10] Y. Freund and R. Shapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *ICML*, 1996.
- [11] M. Jacob and M. Unser, "Design of steerable filters for feature detection using canny-like criteria," *IEEE PAMI*, vol. 26, no. 8, pp. 1007–1019, 2004.
- [12] J. Friedman, T. Hastie, and R. Tibshirani, "additive logistic regression : a statistical view of boosting," *the annals of statistics*, vol. 28, no. 2, pp. 337–407, 1998.
- [13] P. Viola and M. Jones, "Fast and robust classification using asymmetric adaboost and a detector cascade," in *NIPS*, 2001.
- [14] N. Komodakis, N. Paragios, and G. Tziritas, "Clustering via lp-based stabilities," *NIPS*, 2008.

-
- [15] E. Sharon, A. Brandt, and R. Basri, “Completion energies and scale,” *IEEE PAMI*, vol. 22, no. 10, pp. 1117–1131, 2000.
 - [16] N. Honnorat, R. Vaillant, and N. Paragios, “Robust guidewire segmentation through boosting, clustering and linear programming,” *IEEE International Symposium on Biomedical Imaging*, pp. 924–927, 2010.
 - [17] A. Niculescu-Mizil and R. Caruana, “obtaining calibrated probabilities from boosting,” *UAI 05*.
 - [18] L. Yatziv, A. Bartesaghi, and G. Sapiro, “O(n) implementation of the fast marching algorithm,” *Journal of Computational Physics*, vol. 212, pp. 393–399, 2006.
 - [19] Nicolas Honnorat, Régis Vaillant, and Nikos Paragios, “Guide-wire extraction through perceptual organization of local segments in fluoroscopic images,” in *MICCAI (3)*, 2010, pp. 440–448.
 - [20] P. Torr and A. Zisserman, “Robust computation and parametrization of multiple view relations,” *ICCV*, pp. 727–732, 1998.




Figure 10: splines provided by the second method



Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399