

Automatic Detection of the Uterus and Fallopian Tube Junctions in Laparoscopic Images

Kristina Prokopetc, Toby Collins, and Adrien Bartoli

Image Science for Interventional Techniques (ISIT),
UMR 6284 CNRS, Université d’Auvergne, France
prokopec.kristina@mail.ru
{toby.collins, adrien.bartoli}@gmail.com

Abstract. We present a method for the automatic detection of the uterus and the Fallopian tube/Uterus junctions (FU-junctions) in a monocular laparoscopic image. The main application is to perform automatic registration and fusion between preoperative radiological images of the uterus and laparoscopic images for image-guided surgery. In the broader context of computer assisted intervention, our method is the first that detects an organ and registration landmarks from laparoscopic images without manual input. Our detection problem is challenging because of the large inter-patient anatomical variability and pathologies such as uterine fibroids. We solve the problem using learned contextual geometric constraints that statistically model the positions and orientations of the FU-junctions relative to the uterus’ body. We train the uterus detector using a modern part-based approach and the FU-junction detector using junction-specific context-sensitive features. We have trained and tested on a database of 95 uterus images with cross validation, and successfully detected the uterus with Recall = 0.95 and average Number of False Positives per Image (NFPI) = 0.21, and FU-junctions with Recall = 0.80 and NFPI = 0.50. Our experimental results show that the contextual constraints are fundamental to achieve high quality detection.

1 Introduction

An ongoing research objective in medical imaging is to perform inter-modal registration of organs during laparoscopic surgery. The main motivation is to provide Augmented Reality (AR) by visualizing the position of important sub-surface structures such as tumors and blood vessels. This has the potential to significantly improve intraoperative resection planning. The registration problem falls into two main categories depending on whether the non-optical modality is captured preoperatively *e.g.* [11, 12, 5, 18, 15] or simultaneously and intraoperatively *e.g.* [17]. The registration problem is considerably more challenging in the first category because the transform between modalities is not usually rigid. This is due to changes in the organ’s shape between capture times, and caused mainly by the patient lying in different positions, abdominal insufflation and interventional incisions. All the methods for registering laparoscopic and preoperative

images of an organ use anatomical landmarks, which are locations on the organ that are visible in both modalities. A limitation of the above methods is that the landmarks are found manually by a human operator. This is not ideal because it requires the operator to be on hand during surgery and is not practical for locating landmarks in laparoscopic videos. The development of systems to automatically locate landmarks is therefore an important research direction. A second important problem that is also overlooked is organ detection. In previous work the organ is assumed to be visible in the laparoscopic images, so the detection problem is avoided. However, a fully-automatic registration system should detect when the organ is visible, and then instantiate registration. Automatic organ detection also has other important applications, including surgical video parsing and video summarization.

In the context of uterine laparoscopic surgery, it was recently shown that FU-junctions are good landmarks, which are normally formed either sides of the uterus body (Fig. 1). However in [5] FU-junctions were detected manually, and the uterus was assumed to be visible in all laparoscopic images. We present a system for fully automatic detection of the uterus and FU-junctions (with all parameters trained), which brings us a step closer to automatic AR to assist uterine surgeries such as myomectomy and endometriosis.

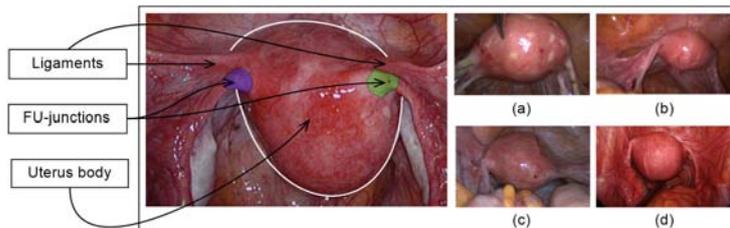


Fig. 1: Laparoscopic images of the uterus. FU-junctions are shown in blue and green for left and right respectively. The detection difficulty comes from ligament junctions, variation in the Fallopian tube orientations and their width. Images (a-d) illustrate inter-patient appearance variation.

2 Background and Related Work

Registering preoperative images in laparoscopic surgery. Existing methods for tackling this problem follow a common pipeline. First the organ is semi-automatically segmented in the preoperative image and a mesh model of its surface is constructed. A deformable model is also constructed to model the non-rigid 3D transform that maps points in the organ to their positions in the laparoscope’s coordinate frame. Most methods require stereo laparoscopic images [11, 12, 18] because these can provide intraoperative 3D surface information.

Recently methods have been proposed for monocular laparoscopes [5]. The registration problem is considerably more challenging with monocular laparoscopes. However the application is broader because the overwhelming majority of laparoscopic surgery is performed with monocular laparoscopes. All methods require a suitable deformation model to constrain the organ’s shape. These have included biomechanical models [12, 11], 3D splines or affine transforms [5]. Organs which have been studied include the liver [12], kidney [11] and uterus [5]. A limitation with all the above methods is that they assume the organ is visible in the laparoscopic images and that there is a manual operator on hand to locate anatomical landmarks.

Detecting objects in optical images. Detecting objects in optical images is a long-standing problem in computer vision that spans several decades of research. In recent years Deformable Part Models (DPMs) have emerged as the best-performing general-purpose object detector [3, 9]. DPMs work by modeling the shape variation of an object class with a set of simple parts that are linked with geometric constraints. Each part models the appearance of the object within a local region. The parts can move to handle geometric variation caused by shape and viewpoint changes. DPMs currently are the best performing detectors in the Pascal Challenge dataset [8], and have been used successfully in other areas of medical imaging such as lung nodule classification [20] and fetal nuchal translucency [7]. However their application to organ detection in laparoscopic images has not yet been investigated.

Junction detection in optical images. There are three main classes of methods for junction detection in optical images. The first are corner-based methods which measure ‘cornerness’ using the image structure tensor [13]. Junctions are then detected as image points with high degree of cornerness. The second are contour-based methods which detect junctions as intersection of image contours [2]. The third are template-based methods which model junctions with a set of templates that correspond to specific junction geometries such as ‘Y’ or ‘T’-shaped, and are learned from natural images [19]. We found that the above classes of methods are not suitable for detecting FU-junctions (Fig. 2). This is for two reasons: *(i)* they are not discriminative enough to separate FU-junctions from other junctions, such as vascular bifurcations, so they give many false positives and *(ii)* they cannot handle well the appearance variation of FU-junctions (Fig. 1).

3 Detection Framework

We propose a learning-based fully-automatic system to detect the uterus and FU-junctions. This is based on four concepts: *(i)* the uterus can be detected prior to FU-junction detection. *(ii)* FU-junctions are too difficult to be detected with generic corner detectors such as [13, 2, 19], so they should be detected with a learned model. *(iii)* FU-junctions are always located close to tube-like structures, so we can filter out many incorrect FU-junction locations if they exist far

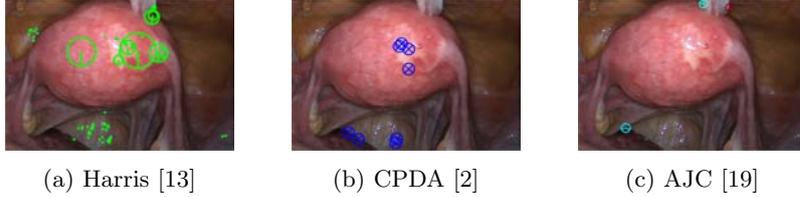


Fig. 2: Failure of generic junction detectors to detect FU-junctions.

from tube-like structures. *(iv)* There exist contextual constraints between the uterus body and FU-junctions. We use two types of contextual constraints. The first models the conditional probability of an FU-junction occurring at a position in the image given the uterus center. Given a uterus detection we can eliminate pixel locations which have low conditional probability giving us Regions of Interest (ROIs) for the locations of FU-junctions. The second contextual constraint encodes the fact that FU-junctions are on the uterus surface, which means there should usually exist a path in the image that connects them to the uterus center which does not cross an object boundary.

Automatically detecting the uterus and FU-junctions is not an easy problem to solve due to large inter-patient anatomic variability (both in shape and texture) (Fig. 1). We restrict the scope of the problem to images of the uterus before resection. This means that the uterus has not changed topologically by surgical incisions. We also assume the uterus is not significantly occluded by surgical tools. In uterine surgery the laparoscope is nearly always held in upright position, so our detectors do not need to be invariant to high degrees of rotations about the laparoscope’s optical axis.

We outline the full proposed detection process in Fig. 3. This consists of two main steps: *(i)* uterus detection and *(ii)* FU-junction detection. We use a trained DPM model to detect the whole uterus, its center and its bounding box. We then proceed to detect the FU-junctions using contextual constraints and a number of processing steps which reduce the search space for FU-junction locations. We then compute local and contextual features for all candidate locations and perform classification with a sparse linear SVM.

3.1 The Uterus Detector

Given an input laparoscopic image (Fig. 3 (a)) we use a trained DPM model to detect the uterus body. This is achieved with an open-source implementation of [10] and a set of annotated uterus images (details of the dataset are given in §4.1). The detector scans the image at multiple scales and positions and returns bounding boxes (Fig. 3 (b)) around positive detections and their corresponding detection scores. We select the bounding box with the highest detection score τ_u , and if τ_u is greater than an acceptance threshold τ'_u the detection is kept (Fig. 3 (c)), otherwise it is rejected (details for computing τ'_u are in §4.1). We use $u_w \in \mathbb{R}$, $u_h \in \mathbb{R}$ and $\mathbf{u}_p \in \mathbb{R}^2$ to denote the uterus bounding box width,

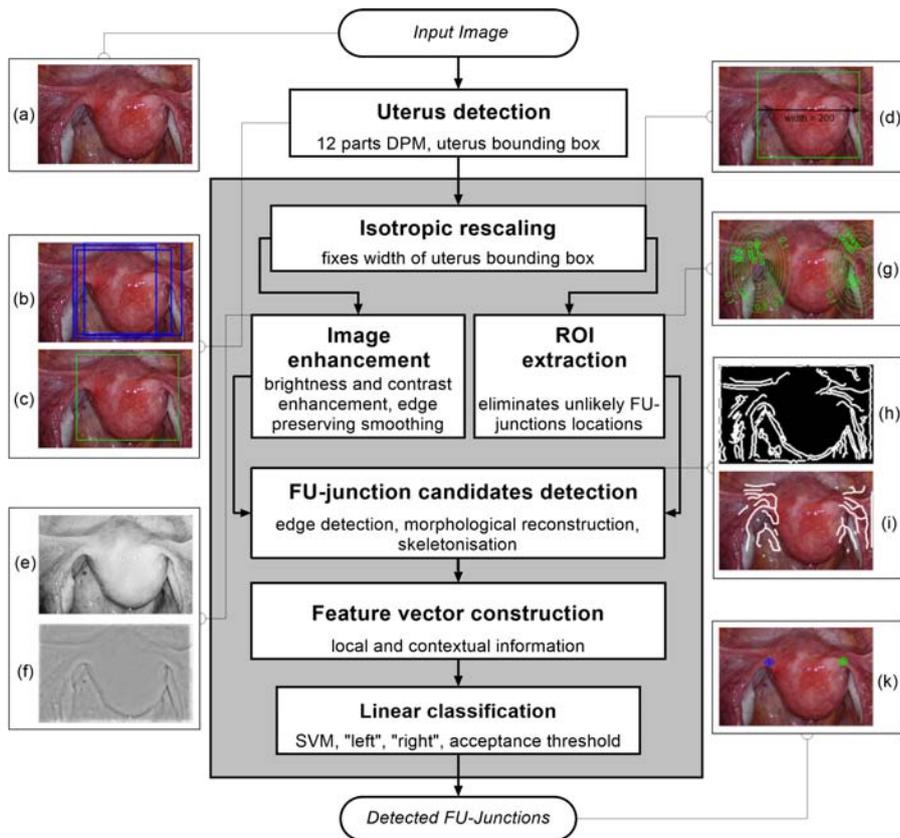


Fig. 3: Diagram of the main pipeline of the proposed detection process.

height and center outputted from the DPM uterus detector. We then proceed to detect the FU-junctions.

3.2 The FU-junction Detector

Step 1: Isotropic rescaling. First the image is isotropically rescaled so the bounding box of the uterus has a default width of $u_w = 200$ pixels (Fig. 3 (d)). This fixes the scale of the uterus and allows us to detect FU-junctions without requiring detection at multiple scales. This has the benefit of increasing computation speed and reducing false positives.

Step 2: Image enhancement. We enhance the image with contrast stretching on the red channel (Fig. 3 (e)). We perform coarse illumination correction to remove uneven illumination with low pass filtering. We then perform edge preserving smoothing using the guided filter method from Matlab (Fig. 3 (f)). We use only the red channel because it is mostly insensitive to the uterus' natural texture

variation (unlike the green and blue channels [4]). This means that strong edges in the red channel are highly indicative of object boundaries.

Step 3: ROI extraction. We filter out highly improbable locations for the left and right FU-junctions. For each pixel $\mathbf{p} \in \mathbb{R}^2$ in the image we compute the conditional probability $P_L(\mathbf{p}|\mathbf{u}_p) \in \mathbb{R}^+$ of the left junction occurring at \mathbf{p} given \mathbf{u}_p . This is a contextual constraint that we model with a Gaussian Mixture Model (GMM):

$$P_L(\mathbf{p}|\mathbf{u}_p) \stackrel{\text{def}}{=} \sum_{k=1}^K w_k^L G(\mathbf{p} - \mathbf{u}_p; \boldsymbol{\mu}_k^L, \boldsymbol{\Sigma}_k^L) \quad (1)$$

where K is the number of GMM components and $\{w_k^L, \boldsymbol{\mu}_k^L, \boldsymbol{\Sigma}_k^L\}$ are the GMM parameters. We keep \mathbf{p} as a left junction candidate if $P_L(\mathbf{p}|\mathbf{u}_p) \geq c$, where c is a small probability threshold. For the right FU-junction we also use a GMM to model the conditional probability $P_R(\mathbf{p}|\mathbf{u}_p)$ of the FU-junction occurring at \mathbf{p} . To train the GMM parameters we exploit the fact that the FU-junctions have strong bilateral symmetry about the uterus body (Fig. 1). Because the laparoscope is normally in upright position this implies the FU-junctions are horizontally symmetric. We therefore propose to simplify the model with $\mu_k^R(1) = -\mu_k^L(1)$, $w_k^R = w_k^L$ and $\boldsymbol{\Sigma}_k^R = \boldsymbol{\Sigma}_k^L$. The advantage of doing this is that we effectively double the amount of training data. This is because each training example can now be used to train P_L and P_R by reflecting its position horizontally relative to \mathbf{u}_p . Training is performed with the standard K-means/EM algorithm on the training set. We set c using a training dataset (see §4.1) at the 99% percentile cut-off point. We select K automatically such that it minimizes the cross-validation error using a hold-out training set (see §4.1). We then compute two ROIs (Fig. 3 (g)), R_l and R_r for the left and right FU-junctions respectively, with

$$R_l(\mathbf{p}) = \begin{cases} 1 & \text{if } P_L(\mathbf{p}|\mathbf{u}_p) \geq c \\ 0 & \text{otherwise} \end{cases} \quad R_r(\mathbf{p}) = \begin{cases} 1 & \text{if } P_R(\mathbf{p}|\mathbf{u}_p) \geq c \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Step 4: Detecting FU-junction candidates. We then detect candidate FU-junction locations using the ROIs from Step 3 (Fig. 3 (h)). This uses the fact that FU-junctions occur close to the medial axis of the Fallopian tubes. We find tube like structures by performing edge detection on the enhanced image computed in Step 2, using the Canny detector with automatic thresholding. Because we use the enhanced image strong edges are highly indicative of object boundaries. We then compute a skeleton S of the edge-map within the region $R_l \cup R_r$ (Fig. 3 (h)) using the implementation of Contour-Pruned Skeletonization from [14], where $S(\mathbf{p}) = 1$ if \mathbf{p} is on the skeleton and $S(\mathbf{p}) = 0$ otherwise. As we see from Fig. 3 (i) the skeleton can be computed quite robustly despite of imperfect edge map. We take all pixels for which $S(\mathbf{p}) = 1$ as a candidate FU-junction locations.

Step 5: Feature vector computation. For each candidate location \mathbf{p} we compute three types of local features (we denote these by x_h , x_θ and x_w). The first x_h

are HOG features [6] to encode image gradient patterns around FU-junctions. We extract HOG features within a local window of w pixels using default HOG parameters, giving x_h 81 dimensions. We have conducted experiments with different window sizes and found a default of $w = 15$ pixels works well. The second local feature $x_\theta \in [0, \pi]$ encodes the orientation of the Fallopian tube as it enters the uterus (Fig. 1). This is computed from the skeleton edge map, by fitting a line to the 5 nearest-neighbors in the skeleton edge map and keeping its slope. The third feature $x_w \in \mathbb{R}^+$ encodes the width of the Fallopian tube as it enters the uterus. This is approximated by twice the distance between \mathbf{p} and the closest edge in the edge map. The reason why we use both HOG and edge-based features is that they complement one another. HOG features do not require computing edge or skeleton maps, which makes them very robust particularly when the contrast between the uterus and background structures is low (even after enhancement). However, HOG features also include gradient information from background structures within the HOG window. On the other hand, edge-based features require edge detection, which makes them less robust. Nevertheless, the benefit of using edge-based features is that if the edges have been computed well, then the edge features encode only the shape of the FU-junction and not structures in the background. We compute two types of contextual features (we denote these by x_g, x_c). The first x_g is computed from the position and direction of \mathbf{p} relative to the uterus center \mathbf{u}_p in the rescaled image:

$$x_g = [d_x, d_y, d_x^2, d_y^2, \alpha]^\top, \quad [d_x, d_y] \stackrel{\text{def}}{=} \mathbf{p} - \mathbf{u}_p, \quad \alpha \stackrel{\text{def}}{=} \text{atan}(d_x/d_y) \quad (3)$$

The second contextual feature x_c encodes the fact that FU-junctions lie on the uterus. Assuming uterus is not occluded by a tool, this means there should exist a path in the image between points \mathbf{p} and \mathbf{u}_p that does not cross the bounding contour of the uterus (Fig. 1). To evaluate this exactly we would need to segment the uterus, which is hard to achieve automatically. Instead we exploit the fact that the uterus body is mostly convex. This means that with high probability the straight line segment between \mathbf{p} and \mathbf{u}_p will not cross the bounding contour of the uterus. In our dataset this assumption holds in all cases, including pathological cases such as uteri with fibroids. We evaluate x_c as the number of times the line segment between \mathbf{p} and \mathbf{u}_p crosses an edge in the edge map. Typically we find that when \mathbf{p} is a correct junction location then $x_c = 0$, however this is not always the case because some spurious edges may exist in the edge map which are caused by high-contrast texture variation.

Step 6: Linear classification. The features are combined into a feature vector which is passed to two trained classifiers. We use one classifier for the left and one for the right FU junctions. We use linear SVM classifiers with an L1 sparse prior, which are known to work well for detectors with HOG features and small datasets of order $\mathcal{O}(10^2)$. We then take the candidates with the highest detection scores for the left and right FU-junctions, and output positive detections (Fig. 3 (k)) if their scores are above an acceptance threshold τ'_j (details for computing τ'_j are in §4.1).

4 Dataset, Training and Performance Evaluation

4.1 Dataset and Training

We have not find any large publicly available collection of laparoscopic uterus images. We therefore constructed the dataset from various sources. This has a total of 95 uterus images from 38 different individuals. 45 images were collected from internet image search engine queries; 26 of which were obtained from 3 publicly available surgical videos. The image resolution of these varied from 311×217 to 1377×1011 pixels. We collected 50 images from 13 videos of different individuals recorded with monocular HD laparoscopes at a partner local hospital. The image resolution of these varied from 553×311 to 817×556 . 77 images in the database were of healthy uteri and 18 were of abnormal uteri with fibroids. All images were annotated manually with uterus bounding box and junction locations. We obtained a negative dataset of 100 images from the 13 videos where the uterus was not visible. These were randomly chosen frames in the time period from insufflation to when the surgeon begun incising the uterus. We divided the dataset into training and test sets using k -fold cross validation with $k = 7$. To guarantee that we measure patient-independent performance, we ensured that images of the same patient were not in training and test sets. At most 4 images of each individual were put in the test set, which was done to keep test performance results balanced across the population. The detection thresholds τ'_u and τ'_j were computed for each fold as the best ‘cut-off’ point on the recall vs. NFPI curve that was closest to $[0, 1]$ (Fig. 4 (b) and Fig. 6).

4.2 Uterus Detection

To evaluate the performance of the uterus detector we adopted the PASCAL VOC challenge protocol to generate Receiver Operating Curves (ROC). A predicted bounding box was considered a true positive if it overlapped more than 50% with the ground-truth bounding box, otherwise it was considered a false positive. Two types of performance have been computed. The first is recall vs. precision and the second is precision vs. Number of False Positives Per Image (NFPI). The most important free parameter of the DPM detector is the number of parts, which we varied from 1 to 12. The evaluation curves shown in Fig. 4 illustrate a general performance gain with increased number of parts. For a precision of 0.90, the recall of the 12-parts modes was 0.86, and the recall of the 6-part model was 0.78. We show some representative detection results in Fig. 5. Typical correct detections are shown in the five top left images while the bottom-right shows a failure due to it being mostly out of frame.

4.3 FU-junction Detection

We compared the performance of our FU-junction detector against two other approaches. The first was a context-free version of our detector where we excluded the contextual features (we named this Context-free). The purpose was

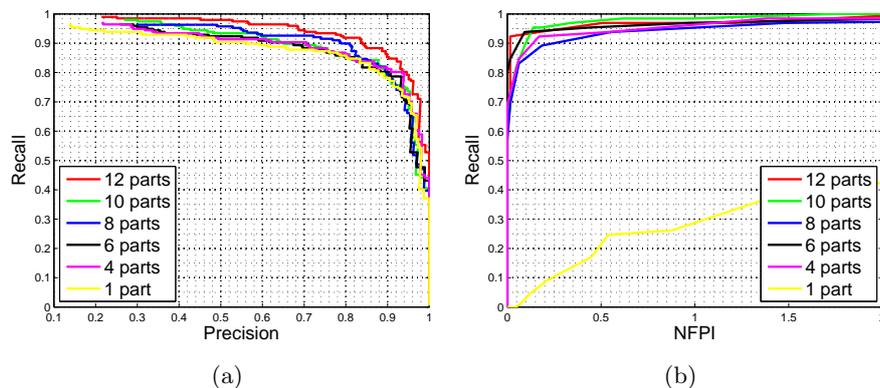


Fig. 4: Uterus detection performance. The Precision/Recall curves are shown in (a) and the NFPPI/Recall curves are presented in (b) with different number of parts.

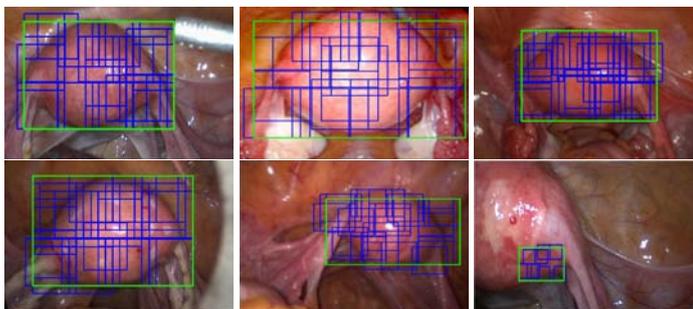


Fig. 5: Examples of uterus' detections. Bounding boxes of the uteri are shown in green and the bounding boxes of the parts are shown in blue.

to reveal the benefit that contextual features had on the problem. The second was the DPM detector from [9] (we named this DPM) that was trained on FU-junctions (and not on the whole uterus). We tested different numbers of parts for DPM and show results for the best number (which was 6). A detection was a true positive if its central point was within the FU junction's ground-truth bounding box, otherwise it was a false positive. We show the recall *vs.* NFPPI curves in Fig. 6. The performance of Context-free and DPM is comparable. One can see a dramatic improvement by our proposed method (*i.e.* when the contextual features are included). For a recall of 0.80 our method achieves a mean NFPPI for the left and right junctions of 0.47 and 0.53 respectively. The performance plateaus at a recall of approximately 0.93%. Therefore in 7% of cases the FU-junctions are so difficult to detect that they cannot be found without having 5 or more false positives. We show some example detections from our method in Fig. 7. The examples show results with normal uteri (second row) and abnormal

uteri with fibroids (first row). The images show the ability to handle significant variation in orientation of the Fallopian tubes. In Fig. 7 bottom left we show a test image where only the right FU-junction was visible (the left FU junction was occluded by the uterus body). A failure is given in the bottom right image, where there was confusion with the round ligament junctions.

In a second experiment we took each positive test image and computed the distance of the best-scoring detection to the ground-truth position. The purpose was to see how well the approaches could localize FU-junctions when they were forced to make a hard decision (*i.e.* the point where the detection score was maximal). Because the test images had different resolutions we rescaled the images to a default width of 640 pixels before computing the distances. The results are shown in Fig. 6 and summary of statistics is given in table 1. For a distance of 25 pixels our method had a recall of 0.73 and 0.64 for the left and right FU-junctions respectively. If we consider the application of registering the uterus, it therefore makes sense for our detector to return a small number of high-scoring detections rather than return the single highest-scoring detection. The set can be used for registration because the correct detection may be determined during registration with *e.g.* softassign [16]. We see that our proposed method performs the best in all statistics except the minimum distance (although it is still under a pixel).

Left FU-Junction						Right FU-junction					
	<i>mean</i>	<i>median</i>	<i>min</i>	<i>max</i>	<i>std</i>		<i>mean</i>	<i>median</i>	<i>min</i>	<i>max</i>	<i>std</i>
Proposed	27.16	10.26	0.25	381.47	56.45	Proposed	25.46	16.12	0.80	117.76	24.95
Context-free	77.96	40.69	0.97	479.17	95.61	Context-free	44.23	24.41	0.52	415.80	65.66
DPM	51.60	23.63	2.12	477.25	78.87	DPM	54.99	29.84	0.31	373.20	82.73

Table 1: Comparison statistics (in pixels) for the three methods in the second experiment with the best method highlighted.

5 Conclusion and Future Work

We have presented an automatic system for detecting the uterus and FU-junctions in laparoscopic images. This work brings us an important step closer to fully automatic inter-modal registration. The average detection time with our current implementation is approximately 8 seconds in unoptimized Matlab code, but with an efficient parallelized implementation can be reduced dramatically because many operations are easily parallelized. With the inclusion of a tool detector *e.g.* [1] the assumption about absence of tool occlusion can be relaxed. We also want to extend the database which will improve performance. Another direction is to extend the detector to stereo images, and it will be valuable to know if the depth data helps detection performance. The possibility to exploit multiple images and/or motion information is also promising for further research.

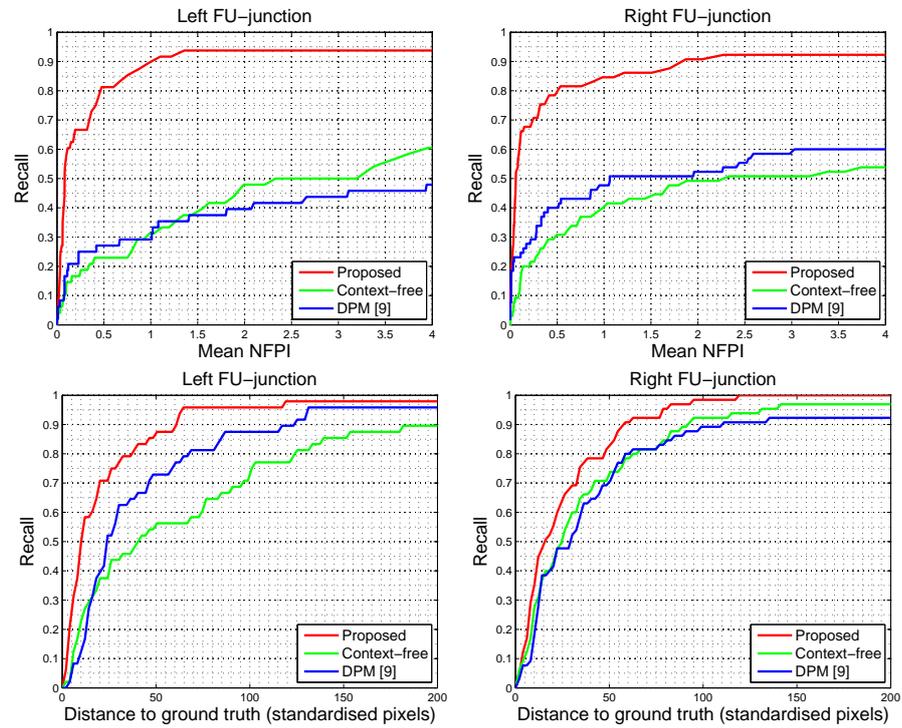


Fig. 6: FU-junction detection performance.

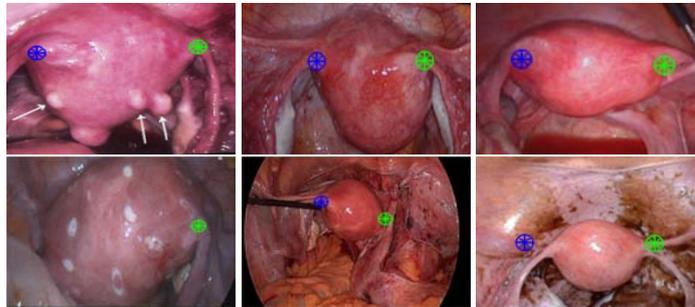


Fig. 7: Examples of detected FU-junctions. The Left FU-junction is shown in blue and right in green. Arrows in the top left image show multiple small fibroids.

Acknowledgements. This research was funded by the EU FP7 ERC research grant 307483 FLEXABLE.

References

1. M. Allan, S. Ourselin, S. Thompson, D. J. Hawkes, J. Kelly, and D. Stoyanov. Toward detection and localization of instruments in minimally invasive surgery. *IEEE Transactions on Biomedical Engineering*, pages 1050–1058, 2013.
2. M. Awrangjeb and G. Lu. Robust image corner detection based on the chord-to-point distance accumulation technique. *IEEE Transactions on Multimedia*, 2008.
3. G. Bouchard and B. Triggs. Hierarchical part-based visual object categorization. In *CVPR*, 2005.
4. T. Collins, D. Pizarro, A. Bartoli, M. Canis, and N. Bourdel. Realtime wide-baseline registration of the uterus in laparoscopic videos using multiple texture maps. In *Augmented Reality Environments for Medical Imaging and Computer-Assisted Interventions*. 2013.
5. T. Collins, D. Pizarro, A. Bartoli, M. Canis, and N. Bourdel. Computer-assisted laparoscopic myomectomy by augmenting the uterus with pre-operative MRI data. In *ISMAR*, 2014.
6. N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
7. Y. Deng, Y. Wang, and P. Chen. Automated detection of fetal nuchal translucency based on hierarchical structural model. In *ICBMS*, 2010.
8. M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results.
9. P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *Pattern Analysis and Machine Intelligence*, 2010.
10. R. B. Girshick. *From Rigid Templates to Grammars: Object Detection with Structured Models*. PhD thesis, University of Chicago, 2012.
11. G. Hamarneh, A. Amir-Khalili, M. Nosrati, I. Figueroa, J. Kawahara, O. Al-Alao, J.-M. Peyrat, J. Abi-Nahed, A. Al-Ansari, and R. Abugharbieh. Towards multi-modal image-guided tumour identification in robot-assisted partial nephrectomy. In *MECBME*, 2014.
12. N. Haouchine, J. Dequidt, I. Peterlik, E. Kerrien, M.-O. Berger, and S. Cotin. Image-guided Simulation of Heterogeneous Tissue Deformation For Augmented Reality during Hepatic Surgery. In *ISMAR*, 2013.
13. C. Harris and M. Stephens. A combined corner and edge detector. In *In Fourth Alvey Vision Conference*, 1988.
14. N. R. Howe. Contour-pruned skeletonization.
15. R. Plantefeve, N. Haouchine, J. P. Radoux, and S. Cotin. Automatic Alignment of pre and intraoperative Data using Anatomical Landmarks for Augmented Laparoscopic Liver Surgery. In *International Symposium on Biomedical Simulation ISBMS*, 2014.
16. A. Rangarajan, H. Chui, and F. L. Bookstein. The softassign procrustes matching algorithm. In *Information Processing in Medical Imaging*, 1997.
17. T. Simpfendorfer, M. Baumhauer, M. Muller, C. N. Gutt, H.-P. Meinzer, J. J. Rassweiler, S. Guven, and D. Teber. Augmented reality visualization during laparoscopic radical prostatectomy. *Journal of Endourology*, 2011.
18. L.-M. Su, B. P. Vagvolgyi, R. Agarwal, C. E. Reiley, R. H. Taylor, and G. D. Hager. Augmented reality during robot-assisted laparoscopic partial nephrectomy: toward real-time 3d-ct to stereoscopic video registration. *Urology*, 2009.
19. G.-S. Xia, J. Delon, and Y. Gousseau. Accurate junction detection and characterization in natural images. *International Journal of Computer Vision*, 2014.
20. F. Zhang, Y. Song, W. Cai, M.-Z. Lee, Y. Zhou, H. Huang, S. Shan, M. Fulham, and D. Feng. Lung nodule classification with multilevel patch-based context analysis. *IEEE Transactions on Biomedical Engineering*, 2014.