

An Augmented Reality System to Guide Liver Punctures

S. Nicolau¹, X. Pennec², L. Soler¹, N. Ayache², J. Marescaux¹

- ¹ IRCAD-Hopital Civil, Virtual-surg, 1 Place de l'Hopital, 67091 Strasbourg Cedex
{stephane.nicolau, jerome.schmid, luc.soler}@ircad.u-strasbg.fr
<http://www.virtual-surg.com> Phone number: +33 388119095
- ² INRIA Sophia, Epidaure, 2004 Rte des Lucioles, F-06902 Sophia-Antipolis Cedex
{Stephane.Nicolau, Xavier.Pennec, Nicholas.Ayache}@sophia.inria.fr
<http://www-sop.inria.fr/epidaure/Epidaure-eng.html>

Abstract. Liver punctures are a difficult operative task that require a precise needle positioning in the center of an invisible target (tumor or healthy tissue in case of biopsy). This article presents an Augmented Reality system for hepatic therapy guidance that superimposes in real-time 3D reconstructions (from CT acquisition) and a virtual model of the needle on external views of a patient. The superimposition of reconstructed models is performed with a 3D/2D registration based on radio-opaque markers stuck on the patient skin. The characteristics of the problem (accuracy, robustness and time processing) led us to develop automatic procedures to extract and match the markers and to track the needle in real-time. Experimental study confirmed that our algorithms are robust and reliable. Preliminary experiments conducted on a human abdomen phantom showed that our system is highly accurate (needle positioning error within 3mm) and enables the surgeon to reach a target in less than 40 seconds on average. Our next step will be to perform an in-vivo evaluation.

1 Introduction

Liver punctures are a medical act that enable to realize little invasive biopsies and to treat patients with unresectable primary or metastatic hepatic cancers (radio-frequency tumor ablation, cryotherapy) [11]. The guidance procedure to reach the target with the needle is usually performed visually using intraoperative two dimensional (2D) cross-sections of the patient obtained with either ultrasound (US) or computed tomography (CT) acquisition. Because of the difficulty to locate the target's center in three dimensions (3D), the needle positioning is not very accurate and the targeting procedure is very time consuming, since it requires many trials. Real-time superimposition of images reconstructed in 3D from CT acquisition onto a real patient, so-called Augmented Reality (AR), may improve the accuracy and decrease complications in interventions such as biopsy needle placement. Such AR guidance systems are routinely used in the field of neurosurgery and orthopedic surgery [4–6, 10, 12, 13, 17]. However, there are few applications on abdominal and thoracic zones. Mourgues [14] superimposes coronary arteries on endoscopic images with an accuracy of about 5 pixels and Lango [8] provides preliminary results in laparoscopy with an accuracy of about 1 cm. Our purpose is to build a guidance system for the needle positioning

that superimposes 3D reconstructions of the liver and the planned target onto peroperative video images of the patient abdomen. To provide a significant help to surgeons, the overall superimposition error has to be less than 5 mm and the computation time of the data processing has to be under 10 minutes.

In our setup, a CT-scan of the patient is acquired just before the intervention, and an automatic 3D-reconstruction of his skin, his liver and the target is performed [18]. Two cameras (jointly calibrated) are viewing the patient’s skin from two different points of view. The patient is intubated during the intervention, so the volume of gas in his lungs can be controlled and monitored. Then, it is possible to fix the volume at the same value during a few seconds repetitively and to perform the needle manipulation almost in the same volume’s condition than the one obtained during the preliminary CT-scan. Balter [3] and Wong [19] indicate that the mean tumor repositioning at exhalation phase in a respiratory-gated radiotherapy context is under 1 mm. Thus, it is reasonable to assume that a rigid registration is sufficient to register accurately the 3D-model extracted from the CT with the 2D video images. We argue in [15] that a rigid 3D/2D point based registration is more adapted to our problem than iconic or surfacic registration. To carry out this registration, we stick twenty five radio-opaque markers on the abdomen skin (cf. phantom in Fig. 1). Since an interactive procedure to localize and match the markers in the CT-scan and video images takes several minutes, we need to realize these tasks automatically.

In this article, we present the technical part of our AR system. Firstly, we describe the algorithms we developed to extract and match automatically the markers from the CT-scan and video images. Secondly, we present the criterion used to find the rigid transformation that relates the CT-scan and the camera reference frame. Then, we explain how we track the needle in real-time to superimpose its virtual model. Finally, we show our preliminary experiments made on our abdominal phantom.




Fig. 1. (a) Video image of the phantom on which 25 radio-opaque markers were stuck. (b) 3D reconstruction of the phantom from CT slices. (c) Organ view in transparency. The liver appears under the skin.

2 2D Marker extraction from video images

The principle of the marker localization in the video images is based on a color analysis and the assumption that the skin covers the main surface area in the video images. In a real intraoperative image, the operative field may modify this assumption. However, it has a homogeneous green color and will cover the border part of the image. Therefore it will be easily identified. The first step is to find the

skin in the image. Working on HSV images, the mean color of the skin MC_{skin} is given, in the hue histogram, by the mode position with the highest integral area, and not necessary the highest value. Indeed, a small region, more homogeneous than the skin, can provide a mode with a lower standard deviation but with a higher maximum value. Then, a color thresholding is computed around the skin hue. Considering that the skin hue distribution is approximately Gaussian, a good choice for the threshold values T_{min} and T_{max} is:

$$T_{min} = MC_{skin} - 3 \cdot \sigma_{skin} \quad \text{and} \quad T_{max} = MC_{skin} + 3 \cdot \sigma_{skin}$$

where σ_{skin} is the standard deviation of the skin mode. To evaluate σ_{skin} , we can compute the full width at half max (FWHM) of the skin hue distribution, which is theoretically equal to $2.3548 \cdot \sigma_{skin}$ under the Gaussian assumption. However, as the histogram is noisy and not totally smooth, it is preferable to evaluate σ_{skin} from a smoothed version of the original hue histogram. The optimal histogram is defined as the iteratively smoothed original histogram with a Gaussian filter until its number of modes becomes stable. We call σ_{opt} the standard deviation of the smoothed hue distribution of the skin and Max_{opt} (resp. Max_{skin}) its maximum value (resp. the maximum value of the original skin hue distribution). Since smoothing does not modify the integral area of the distribution mode, σ_{skin} is related to σ_{opt} by: $\sigma_{skin} \cdot Max_{skin} = \sigma_{opt} \cdot Max_{opt}$

After this step, as we suppose that the skin covers the largest area in the image, we extract the largest connected component. Figure 2b shows the results after the color thresholding and the largest component extraction. Then, since the color of markers is very different from skin color (the markers are green), it can be considered that all the connected components entirely enclosed in the detected skin are potential markers. To characterize the real markers among all these components (cf. Figure 2c), we first remove those whose surface is empirically too small to be a marker (under 10 pixels of surface). Then, we compute the median of component surfaces S_{median} and we get rid of the components whose surface is not in the interval $[0.5 \times S_{median} ; 2.0 \times S_{median}]$. Finally, we discard the components whose inertial moment is too small (characterizing the flat components).

The remaining components (cf. Figure 2d) are considered as markers and their barycenter is computed. To evaluate the performance of this algorithm, we conducted an experimental study on 10 subjects. We stuck around 20 markers on each person and we took pictures of their abdomen in 4 different positions with 2 cameras. From these 80 images, the mean computation time was around 1 second and we had 5% of false negatives and 2% of false positives whose location were never on the skin. Visual inspection indicates that the detection accuracy is around 2 pixels on average. This 2D extraction algorithm is then fast and robust enough to be easily used in a matching algorithm involving 2D and 3D points.

3 3D Marker extraction from the CT-scan

In order to localize the markers in CT-scan images (cf. Figure 3a), we essentially use mathematical morphological tools. In a first step, we threshold the 3D image to extract both patient and radio-opaque markers. CT-scans being calibrated




Fig. 2. (a) Original image of the phantom. (b) Result of the color thresholding and largest component extraction. (c) The potential markers are all the components in white. (d) Final extraction of markers into the detected skin. The real markers are characterized among the components with respect to their surface and shape.

(Hounsfield Units), a CT independent threshold value can be chosen (-800 HU for the phantom and -300 for a human). In a second step, since some artefacts can appear (like the CT table), we extract the largest connected component from the binary images to get only the patient, and we fill the internal cavity out (cf. Figure 3b). Then, in order to erase the markers, we perform an opening with an adapted structuring element (a sphere with a radius 3 times larger than the marker diameter) (cf. Figure 3c). Eventually, a simple subtraction between the opened image and its previous version provides the markers.

The opening erases the markers but smooths the surface as well, therefore, residual pieces of surface can remain (cf. Figure 3d). We get rid of them by opening the components found with a small structuring element and by characterizing them with their volume and shape (the same procedure as for video markers). Since the radio-opaque markers are designed to give a high response in CT images, we compute their intensity weighted barycenter \vec{G} to take into account partial volume effects:

$$\vec{G} = \frac{\sum_{\vec{P} \in Comp} I(\vec{P}) \cdot \vec{P}}{\sum_{\vec{P} \in Comp} I(\vec{P})}$$

where \vec{P} is a 3D position of a voxel in the component and $I(\vec{P})$ is the voxel intensity. An evaluation performed on 6 CT-scans of the phantom and 3 of a pig, on which were stuck about 25 markers, showed that there was less than 5% of false negatives and 2.3% of false positives. Most of them corresponded to the pig nipples or to pipes lying along its body. Their presence had no incidence since they were easily identified and discarded during the robust matching step.




Fig. 3. a) Original CT-slice. A radio-opaque marker is circled. b) Original image after the thresholding and the filling. c) Opening of the largest component. Markers have disappeared. d) Subtraction of image (c) to image (b) provides the markers.

4 Video marker matching

The matching between the 2D markers extracted from the two video images is made with epipolar geometry constraints. As our cameras are jointly calibrated, it is possible to define a relationship between a point in one image and its correspondent in the other one. If \mathbf{m} is a point in the image \mathbf{P} , its correspondent \mathbf{m}' in the image \mathbf{P}' lies on a segment whose location depends on the optical center \mathbf{C} and \mathbf{C}' of the cameras, the position of \mathbf{m} and the range of depth to which the phantom belongs. Because of imperfect camera calibration and noise, \mathbf{m}' is not perfectly on the predicted segment, therefore we select the points that are very close to it. If there is only one point, we consider the match as reliable, otherwise it is ambiguous (see Figure 4). In that case, it is discarded at this step (these markers will be reconsidered in the next step).

Finally the matched markers are reconstructed in 3D in the camera reference frame. If erroneous 3D reconstructions due to false point matches appear, they are discarded by the subsequent video/CT matching procedure designed to be robust in the presence of such outliers. Experiments realized with 210 image pairs of the phantom showed that on 25 potential matches, an average of 15 were correct, 8.7 were ambiguous and 1.3 were wrong.

5 Video/CT-scan marker matching

A set of 3D markers being reconstructed in the camera reference frame, we need to match them with the corresponding 3D points extracted from the CT-scan. We decided to use a prediction/verification algorithm [2] with additional distance constraints. This algorithm provides us not only a matching estimation, but a transformation initialization as well. The principle is to choose a triplet from the reconstructed marker set and to find a triplet from the CT-scan markers that has the same dimension (the triangle sides have approximately the same length). If no CT-scan triplet is found, a new triplet from the reconstructed marker set is chosen. When a CT-scan triplet is found to have the same dimension,




Fig. 4. Example of epipolar matching on a video image pair. On the right image are drawn the epipolar segments that correspond to the outlined markers in the left image. One can see that only one marker lies on L1 and L2 (at the bottom). In that case, the match is considered as reliable. L3 and L4 being close to 3 different markers, this case is ambiguous: P3 and P4 are not matched with another marker.

we compute the rigid transformation that best matches the 2 triplets [1]. To verify this prediction, we apply this transformation to the extracted CT-scan markers and project them in the video images. Since the CT (resp. video) marker extraction algorithm provided 5% (resp. 5%) of false negatives, we accept the transformation as correct if more than 85% of the projected points are close to the extracted 2D video points. Otherwise, the 2 triplets are not homologous despite the dimension similarity (this case is illustrated on Figure 5): we have to choose a new triplet in the CT-scan marker set. In fact, to limit computation time, it is much more interesting to sort the triplet so that we first choose those whose number of potential corresponding triplets is small (typically, we begin with the triplets of largest dimension).




Fig. 5. Illustration of a bad triplet matching. One can see in the light circles three 3D markers from the CT model that are well superimposed with three 2D markers of the video image. Since almost all the other 3D model markers are not close to video markers (in dark circles), it means that the two triplets are not representing the same markers.

When a correct transformation is found, we keep all the 3D/2D correspondences provided by the verification step to compute a more accurate transformation (see next section). In this process, markers that belong to ambiguous 2D/2D video matches are implicitly reconsidered thanks to the verification. The

matching experiments made with the 210 image pairs of the phantom used in the previous section always provide a correct transformation, and 100% of the possible matches were found. The whole matching process took on average less than 2 seconds and 3 seconds at the most. Eventually, these results show that our procedure of extraction and matching is sufficiently robust and reliable for clinical use.

6 Computation of the best rigid transformation

We could have directly used the rigid transformation obtained between the 2 matched triplets, but it only relies on three 3D matches. Since we need the best possible accuracy, it is desirable to take the information brought by all 3D/2D matches into account. The standard least square 3D/2D criterion supposes that the noise corrupting the 3D data is negligible. This assumption being false in our case, we optimize the maximum likelihood EPPC criterion that we developed in [16] and that is optimal in the case of noise on both 2D and 3D markers extraction. This criterion is more complex than the standard one:

$$EPPC(T) = \sum_{i=1}^N \frac{\|\tilde{M}_i - M_i\|^2}{2 \cdot \sigma_{3D}^2} + \sum_{k=1}^M \sum_{i=1}^N \xi_i^k \cdot \frac{\|\tilde{m}_i^{(k)} - P^{(k)}(T \star \tilde{M}_i)\|^2}{2 \cdot \sigma_{2D}^2}$$

where T is the sought transformation, P^k is the projective function of the k^{st} camera, \tilde{M}_i are the extracted 3D CT-scan markers, $\tilde{m}_i^{(k)}$ the extracted video markers from the camera k , σ_{2D} and σ_{3D} the estimated standard deviation of the noise corrupting the CT and video markers extraction, M_i the exact location of the CT markers and ξ_i^k a binary variable equal to 1 if M_i is observed by the k^{st} camera and 0 otherwise. Note that M_i are not known and are consequently estimated just as the transformation T (M_i are called hidden variable). We perform an alternated minimization w.r.t. the two groups of variables, starting from a transformation initialization of the M_i with the \tilde{M}_i . The algorithm is stopped when the distance between the last two estimates of the transformation becomes negligible. An experimental study realized with the phantom [16] showed that a superimposition accuracy of 2 mm was reached for targets within the liver.

7 Superimposition of the virtual needle on the real one

As the needle is no longer visible as soon as the radiologist introduces it under the skin, its virtual model needs to be superimposed in real-time onto the video images. This means that we have to track its location and orientation in space in the camera reference frame. To realize this, we attach an oriented square marker to the external part of the needle, and we localize in real-time its corners on video images using a standard version of the ARToolkit library [7] that we have improved. Then, knowing the size of the square, we are able to localize it in the camera reference frame by minimizing the classical 3D/2D SPPC criterion [16]:

$$SPPC(T) = \sum_{k=1}^S \sum_{i=1}^N \xi_i^k \cdot \frac{\|\tilde{m}_i^{(k)} - P^{(l)}(T \star \tilde{M}_i)\|^2}{2 \cdot \sigma_{2D}^2}$$

where \tilde{M}_i are the position of the 4 corners in the square marker frame, $\tilde{m}_i^{(k)}$ are the pixel position of the 4 corners in the k^{st} video image and T is the

transformation between the square marker frame and the camera reference frame. Calibrating the relative needle position with respect to the square marker with the pivot method [9], we are finally able to superimpose the virtual model on the real one on video images. An accuracy evaluation, realized in simulated clinical conditions, proved that the average superimposition error of a standard needle is below 1 mm.

8 Evaluation of the overall system

All the components being validated independantly, we now have to assess the accuracy of the complete AR guidance system (described in Fig. 6). Four targets were modeled with radio-opaque markers stuck on the fake liver inside the phantom. Seven participants performed each 10 consecutive needle targetings of the model tumors (cf. Fig. 7 a). During the positioning, the operator placed the needle and stopped his movement when he thought that he had reached the tumor center. After each trial, the time required to position the needle was recorded, and the accuracy of hits was verified by an independent observer using an endoscopic camera introduced into the phantom and focusing on the targets (cf. Fig. 7 b). Accuracy and time results are shown in Table 1.




Fig. 6. Example of the augmented view provided to the user. (a) Virtual view of a camera positionned at the needle tip and oriented toward its axis. (b) The 3D reconstruction of the liver and the virtual needle are displayed on the video frame. (c) Virtual view displaying the real position of the needle with respect to the model of the abdominal phantom. We indicate in its corner the virtual distance in mm that separates the needle tip from the target (in this case, a marker center).

The results indicate clearly that our system provides an accuracy that fulfills our accuracy constraint (5 mm), because the worst average accuracy obtained is below 3 mm. In addition, it allows to reach the target very quickly (average time under 30 sec.) with respect to the usual time needed for a standard percutaneous intervention (10 minutes).




Fig. 7. a) Setup of the experiment: the user is positioning the needle, tracked by a stereoscopic system, thanks to the guidance interface. b) An endoscopic view is displayed behind the user. It enabled to assess visually the correctness of each needle targeting.

	Average distance (mm) \pm std.	Minimum distance	Maximum distance	Average time (sec.) \pm std.
Engineer 1	0.95 ± 0.67	0	2	25 ± 5.6
Engineer 2	1.7 ± 0.97	0	3	14 ± 2.0
Engineer 3	1.8 ± 0.84	0	3	18 ± 5.5
Surgeon 1	2.2 ± 0.57	1	3	32 ± 12.2
Surgeon 2	2.9 ± 1.25	0	5	22 ± 3.1
Surgeon 3	1.3 ± 1.16	0	3	32 ± 3.7
Professor	0.84 ± 0.48	0	1	32 ± 6.4
All	1.8 ± 0.7	-	-	23.8 ± 7.3

Table 1. Accuracy and time results obtained by each user. The average distance, which is always below 3 mm, fulfills our accuracy constraints (5 mm). Moreover, the time needed is, by far, under 1 minute, whereas an expert needs routinely 10 minutes for such an intervention.

9 Conclusion

In this paper, we have presented an AR system, based on radio-opaque marker registration, to guide the needle positioning during liver punctures. To fulfill time constraints, we have developed procedures to extract and match the markers in the video and the CT-scan images. Experimental studies proved that these algorithms are robust and fast enough for clinical application. In order to evaluate the efficiency of the whole system, we conducted needle positioning experiments on targets embedded into our human abdominal phantom. Results show that our system of AR guidance is highly accurate (1.8 mm RMS) and enables to locate hidden targets quickly (less than 40 seconds on average), whereas a standard percutaneous intervention generally takes about 10 min. One other advantage of this system is the simple material needed. Indeed, only a PC, a video acquisition card, two cameras and a sterilizable square flag are required, unlike the other existing systems that need devoted tracking materials [8, 10, 12] (such as Optotrak©, Flashpoint©) or a laser scanner [5].

In order to reach a clinical validation and use, we are currently preparing experiments on human real patients to verify that our system fits real operative constraints (low organ mobility and good illumination conditions). We are pretty confident that the system will be fully validated in the next six months.

References

1. K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3d point sets. *IEEE Trans. on PAMI*, 9(5):698–700, september 1987.
2. N Ayache and O. Faugeras. Hyper: A new approach for the recognition and positioning of 2d objects. *IEEE Trans. on PAMI*, 8(1):44–54, January 1986.
3. J.M. Balter, K.L. Lam, C.J. McGinn, T.S. Lawrence, and R.K. Ten Haken. Improvement of CT-based treatment-planning models of abdominals targets using static exhale imaging. *Int. J. Radiation Oncology Biol. Phys.*, 41(4):939–943, 1998.
4. D. Gering, A. Nabavi, R. Kikinis, W. Grimson, N. Hata, P. Everett, F. Jolesz, and W. Wells. An integrated visualization system for surgical planning and guidance using image fusion and interventional imaging. In *Proc. of MICCAI'99*, pages 808–819, 1999.
5. W. Grimson, G. Ettinger, S. White, W. Wells T. Lozano-Perez, and R. Kikinis. An automatic registration method for frameless stereotaxy, image-guided surgery and enhanced reality visualization. *IEEE TMI*, 15(2):129–140, April 1996.
6. A. Guézic and P. Kazanzides. Anatomy-based registration of CT-scan and intra-operative x-ray images for guiding a surgical robot. *IEEE Trans. Medical Imaging*, 17(5):715–727, October 1998.
7. Hiro. Human interface technology laboratory, <http://www.hitl.washington.edu/>.
8. T. Lango, B. Ystgaard, G. Tangen, T. Hernes, and R. Mrvik. Feasibility of 3d navigation in laparoscopic surgery., Oral presentation at the SMIT (Society for Medical Innovation and Technology) Conference. September 2002. Oslo, Norway.
9. S. Lavalley, P. Cinquin, and J. Troccaz. *Computer Integrated Surgery and Therapy: State of the Art*, chapter 10, pages 239–310. IS Press, Amsterdam, NL, in C. Roux and J.L. Coatrieux edition, 1997.
10. C. Maurer, J. Fitzpatrick, M. Wang, R. Galloway, R. Macinuas, and G. Allen. Registration of head volume images using implantable fiducial markers. *IEEE Trans. Med. Imaging*, 16:447–462, 1997.
11. J.F. McGahan and G.D. Dodd. Radiofrequency ablation of the liver: Current status. *American Journal of Roentgenology*, 176(1):3–16, 2001.
12. P. Merloz, F. Lavalley, J. Tonneti, and L. Pittet. Image-guided spinal surgery: technology, operative technique and clinical practice. *Operative Techniques in Orthopaedics*, 10(1):56–63, 2000.
13. M Mitschke, A. Bani-Hashemi, and N. Navab. Interventions under video-augmented x-ray guidance: Application to needle placement. In *Proc. of MICCAI'00*, pages 858–868, Pittsburgh, Pennsylvania USA, October 11-14 2000.
14. Fabien Mourgues. *Guidage par réalité augmentée: application la chirurgie cardiaque robotisée*. PhD thesis, Université de Nice - Sophia Antipolis, September 2003.
15. S. Nicolau, X. Pennec, L. Soler, and N. Ayache. Evaluation of a new 3D/2D registration criterion for liver radio-frequencies guided by augmented reality. In *Proc. of IS4TM'03*, LNCS 2673, pages 270–283, France, 2003. Springer-Verlag.
16. S. Nicolau, X. Pennec, L. Soler, and N. Ayache. An accuracy certified augmented reality system for therapy guidance. In *European Conference on Computer Vision (ECCV'04)*, LNCS 3023, pages 79–91. Springer-Verlag, 2004.
17. B. Peuchot, A. Tanguy, and M. Eude. Virtual reality as an operative tool during scoliosis surgery. In *Proc. of CVRMed'95*, pages 550–554, Nice, France, 1995.
18. L. Soler et al. Fully automatic anatomical, pathological, and functional segmentation from CT scans for hepatic surgery. *Comp. Aided Surg.*, 6(3), August 2001.
19. J. Wong et al. The use of active breathing control (abc) to reduce margin for breathing motion. *Int. J. Radiation Oncology Biol. Phys.*, 44(4):911–919, 1999.