

# Detecting a Needle in C-arm X-ray Images

Matthias Scharrer\*

Institute for Computer Graphics and Vision  
Graz University of Technology  
Graz / Austria

## Abstract

This paper presents an algorithm to detect a rigid straight biopsy needle in multi-view C-arm X-ray images and reconstruct its tip and orientation in 3D-space. Several well known computer vision techniques are applied to achieve this goal. The processing pipeline that we describe consists of five stages containing a denoising and preprocessing stage with Gaussian filtering, computing the Hessian tensorfield and a Radon transform for ridge detection, needle tip detection, the reconstruction of the data into 3D-space with Direct Linear Transform and improving robustness by backprojection. The processing steps are described in detail, a short overview is given about the surrounding application and finally we present the evaluation results of the experiments on real X-ray imaging data.

**Keywords:** X-ray, multiview geometry, Radon transform, Hessian matrix, line detection

## 1 Introduction

During prostate biopsy a probe head and a core needle are inserted to excise a tissue sample from the prostate for histological examination to find out if cancer or other abnormal cells are present or to remove or destroy tumors in the prostate. Since this happens *in vivo* with least possible intrusion, the surgeon has no direct visual opportunity to get feedback on the needle's exact position. Due to this, the surgeon can not be certain about the success of taking the sample without retracting the probe or if the right positions were penetrated.

To overcome this problem several X-ray photographs are taken with a C-arm device from different angles.

A good approach would be to determine a model of the patient's organs with a computed tomography as a preliminary stage to the biopsy, fit this model to the reconstruction from the inter-operative X-ray images [9] and reconstruct the needle within this model. This will permit the surgeon to relocate a specific position of the needle and, if necessary, to adhere to a fixed predefined pattern of taking a number of samples.

Recent work has already covered the registration of the pre-operative CT volume dataset and inter-operative X-ray

images [5]. In this paper we present an algorithm to reconstruct the needle's position and orientation in a robust way from several views.

### 1.1 Outline

In Section 2 we give a short overview of similar work that has been done considering the task of detecting a needle. In Section 3 we describe the basics of the algorithms and principles of the methods that we apply.

The processing pipeline stages consist of the following steps: detecting the needle ridge in each image, finding its tip, reconstructing this information in 3D-space and rejecting outliers that would lead to misdetection. The stages of the algorithm are covered in Section 4 along with a short description of the overall application that this algorithm is a part of, and the results are discussed in Section 5.

In Section 6, at last, we will point out the main issues that have arisen during the experiments and provide some future prospects.

## 2 Related work

A lot of research has been put recently into robotic steerable needle insertion as this seems to be very promising to improve methods in minimal invasive therapy and surgery in near future. Though most of the papers focus on the robotic and mechanical aspects of this task, needle detection seems to play an increasingly important role.

In [4], an optically controlled flexible needle steering device is presented, where needle detection in X-ray images plays an important role in the feedback loop of the steering control. It is based on detecting a characteristic shape at the basis of the needle with a normalized cross correlaton of a template image and the X-ray image. The needle tip itself is found afterwards by following the low gradient area that starts at the detected basis feature. Knowing the length of the needle, the shape detected can easily be verified. Additionally, the shape of the needle is fitted by a polynomial that smoothes the line. This approach seems sufficiently robust for testing environments, but as the authors state it is subject to noise, occlusion by beads and low contrast of the image intensities.

In [3] a steering aid for flexible needles is developed that is intended to operate as an image enhancer to the physi-

---

\*matthias.scharrer@student.tugraz.at

cian or fully automatically. In contrast to the method described by Glozman and Shoham [4], the image of the needle is taken on an ultra sound basis, where the transducer has to be aligned to the image plane manually once. The needle detection algorithm takes two points on the approximate needle axis determined manually as an initial value and then performs the detection. The detection is based on noise filtering, gradient based edge detection, Hough transform and polynomial fitting of the needle. Although the needle detection is very robust, it is stated that the detection of the needle tip is not very accurate due to the noise in ultrasound images, e.g. speckles and signal drop-out.

To further improve needle tip detection it is suggested to actively oscillate the tip with some higher frequency. This is not interfering with the steering of the needle and permits the ultra sound device to detect the source of the frequency easier.

### 3 Background Techniques

Throughout this work several well established methods are applied to accomplish the given task and in this Section we want to describe the principles of these shortly.

#### 3.1 The Radon transform for line detection

The Radon transform is defined in general to integrate functions  $f$  on  $n$ -dimensional Euclidean space  $\mathbb{R}^n$  over all  $n - 1$ -dimensional hyperplanes. In our case,  $f$  that can be considered a two dimensional distribution is projected onto all lines by determining all line integrals of  $f$  on some domain  $D$  on  $\mathbb{R}^2$ .

$$\hat{f} = \int_L f(x,y) ds \quad (1)$$

This can also be thought of rotating a parallel beam projector around a problem domain and simultaneously registering the incoming beams at the opposite site of the domain, as illustrated in Figure 1. A line  $L$  can be denoted by the normal form

$$p = x \cos \phi + y \sin \phi \quad (2)$$

A straight line becomes a single point in this transformed space, so the problem of detecting a line can be further reduced to detecting a peak in two dimensional parameter space, which can be found easily.

#### 3.2 The Hessian matrix for ridge detection

The ridge detector we used is based on the Hessian matrix defined in (3), as it is supposed to be very responsive at edges and blobs [10]. The Hessian matrix is defined as:

$$\mathbf{H}(x,y) = (\nabla' \cdot \nabla)I = \begin{bmatrix} \frac{\partial^2 I}{\partial x^2} & \frac{\partial^2 I}{\partial xy} \\ \frac{\partial^2 I}{\partial xy} & \frac{\partial^2 I}{\partial y^2} \end{bmatrix}, \quad (3)$$

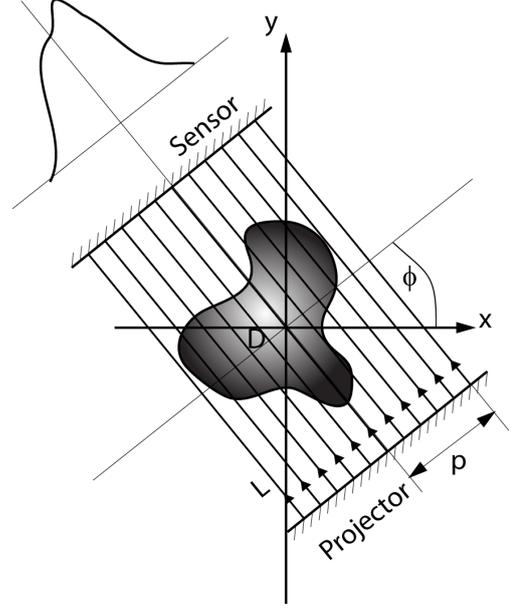


Figure 1: The beamer-sensor system the Radon transform can be thought of.

where  $\nabla$  denotes the operator  $\left(\frac{\partial}{\partial x} \frac{\partial}{\partial y}\right)^T$ ,  $I_{xx}$  and  $I_{yy}$  are the second derivatives of the image intensity  $I$  at position  $(x,y)$  and  $I_{xy}$  is the mixed derivative in both directions. After performing the eigenvalue decomposition of the Hessian matrix

$$\lambda_{1,2} = \frac{1}{2} \left[ (H_{11} + H_{22}) \pm \sqrt{4H_{12}H_{21} + (H_{11} - H_{22})^2} \right], \quad (4)$$

where  $|\lambda_1| \geq |\lambda_2|$  are the eigenvalues of  $\mathbf{H}(x,y)$  and

$$\mathbf{H}(x,y)v_i = \lambda_i v_i, i \in \{1,2\} \quad (5)$$

where  $v_{1,2}$  are the eigenvectors of  $\mathbf{H}(x,y)$ . It can be stated that if  $\nabla I v_1 = 0$ , the point belongs to a ridge ( $\lambda_1 < 0$ ) or a valley ( $\lambda_1 > 0$ ) [7].

#### 3.3 Reconstruction of a point to 3d-space

A very simple linear triangulation method would be to combine the known coordinates of the points in the image planes  $\mathbf{x}_i = P_i \mathbf{X}_i$ , which can also be written as

$$\mathbf{x}_i \times (P_i \mathbf{X}_i) = 0, \quad (6)$$

where  $\mathbf{X}_i$  are the world coordinates of  $\mathbf{x}_i$  and  $P_i = (\mathbf{p}_{i,1}, \mathbf{p}_{i,2}, \mathbf{p}_{i,3})^T$  are the respective projection matrices.

Reconstruction can be done by bringing the above into a form of homogeneous equations  $A\mathbf{X} = 0$ , where  $A$  is a

composition of two Equations (6) for each image:

$$\begin{pmatrix} u_1 \mathbf{p}_{1,3}^T & - & \mathbf{p}_{1,1}^T \\ v_1 \mathbf{p}_{1,3}^T & - & \mathbf{p}_{1,2}^T \\ u_2 \mathbf{p}_{2,3}^T & - & \mathbf{p}_{2,1}^T \\ v_2 \mathbf{p}_{2,3}^T & - & \mathbf{p}_{2,2}^T \\ & & \vdots \\ u_n \mathbf{p}_{n,3}^T & - & \mathbf{p}_{n,1}^T \\ v_n \mathbf{p}_{n,3}^T & - & \mathbf{p}_{n,2}^T \end{pmatrix} \mathbf{X} = 0, \quad (7)$$

where  $u_i$  and  $v_i$  are the elements of  $\mathbf{x}_i$ .

However, it is not possible to reconstruct the world coordinates exactly from several views, but rather it is necessary to estimate them in an optimization function that minimizes some cost function, as the rays that can be backprojected from the image through the respective camera centers are skew and the system of Equations (7) is overdetermined.

As stated in [6] this kind of reconstructional problems can be treated as finding a least-squares solution to Equation (7), i.e., finding that  $\mathbf{X}$  that minimizes  $A\mathbf{X}$  subject to  $\|\mathbf{X}\| = 1$ .

## 4 Processing pipeline

The chosen approach can be partitioned into several smaller stages, which can be treated as individual modules as depicted in Figure 2. They are detailed in their respective subsections regarding a short overview of the overall application in Section 4.1, removing clutter (e.g. small globes, probe head) in Section 4.2, finding the line the needle is aligned to in section 4.3, detecting the tip along this line in Section 4.4, combining this data to find its representation in world coordinates in Section 4.5 and finally rejecting the wrongly detected needle with back-projection of the reconstructed data in Section 4.6.

### 4.1 Application

As already stated, this module is part of a more complex application supposed to assist the surgeon during biopsy of the prostate, i.e., exciting tissue for histological examination, the ablation of cancerous tissue or the implantation of low radiating beads that would destroy the former.

A precondition to this minimal invasive navigational surgery is knowledge of the exact location of the biopsy needles inside the prostate and the establishment of a mapping pathology of the cancer lesion in the single patient.

It is considered to take three images with a mobile C-arm, each with maximal mechanical offset (e.g.,  $-45^\circ$ ,  $0^\circ$  and  $+45^\circ$ ) but minimal femoral influence. Subsequently, line detection methods have to be adopted to detect the needle in these images. The needle images are projected back into 3D space, where the best fitting line will be computed. At the same time, a CT or MR based model of the

patient would be registered to the 3D space and the location of the penetrated tissue will be available.

For the purpose of reconstructing camera positions and orientations, a special target with small globes attached in a specific pattern visible to X-ray imaging has been designed and used for reconstruction of the camera orientation. Furthermore, the cameras are to be calibrated ahead of the biopsy.

Visualization and interaction methods should finally allow the surgeon to easily access the desired information. We chose Coin3D<sup>1</sup> for visualization of the resulting data.

### 4.2 Filtering background clutter

As shown in Figure 2, there is already some work done a priori, that can be used to ease the task of filtering background clutter:

- The geometry of the screen always stays the same. This implies, that we can use a constant mask to mask out the additional information that is written to the image, such as the name of the patient and others, so that we can focus on the actual image area, which is circular in our case.
- The cameras are already calibrated, i.e., the images are undistorted. This plays a very important role, as X-ray images are subject to a very strong pincushion distortion.
- The positions of the small globes are computed during an earlier step, so this information can be used to mask out the small beads.

After doing so, we want to mask out the probe head, too, as this is the only significant structure in the image besides the needle that we want to detect, that remains.

However, prior to processing the data, a Gaussian filter should be applied to the images, since X-ray image sources tend to be very noisy and we are going to compute second order derivatives later on, which are very sensitive to noisy images. As it will permit us to do scale space operations, too, we start by building the well known Gaussian pyramid [1] from the images. This is done the way described in [8], so we apply several filters at each octave of the scale.

To detect the probehead, it has turned out to be very efficient to threshold the image to a very low value at a higher scale and detect the biggest area remaining. After dilating this area several times and combining it with the already known information we have a good mask to apply to the image.

### 4.3 Detecting the line the needle lies on

The core needle which is used to extract the tissue samples is a hollow needle which contains the tissue sample after

<sup>1</sup>A free scene graph library under GPL - <http://www.coin3d.org>

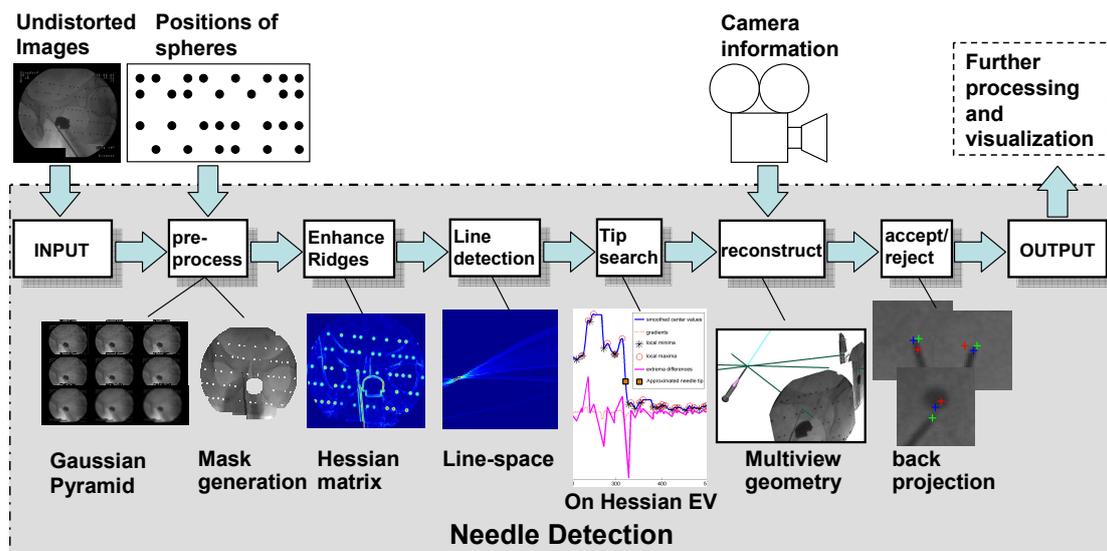


Figure 2: A schematic overview of the processing pipeline with its stages and its interaction with the entire application.

a successful shot. It can be considered an elongated rigid object of a small diameter of a few millimeters. When viewing the image it can be considered a thick line, so the problem of detecting the needle is reduced to detecting a line in the image. In [2] it is stated, that the Radon transform is strongly related to the Hough transform, which is commonly used to detect lines in images, as described in Section 3.1.

But applying the Radon transform on the preprocessed grey value image we have got so far will most probably not end up in a single point in parameter space, but rather in very unpredictable results, since the background will also contribute with a sine function for every point in the image. Therefore, the image will not be used directly for the Radon transform but it will be transformed to a different representation to enhance discrimination of lines against not entirely homogenous background utilizing the Hessian matrix described in Section 3.2.

As shown in Figure 3, the needle is very prominent feature in the image of the stronger eigenvalues - but so are the small globes and the probehead. However, these can be masked out to some degree as described in Section 4.2. What remains is the needle and some very low value ridges of the patient's bones. We have decided to threshold the stronger eigenvalue image at the half of the maximum value. This is used along with the conditions for ridge detection defined above as input to the Radon transform. The result (Figure 4) is then searched for local peaks after non-maximum suppression. The strongest of which is most likely to be the ridge of the needle.

#### 4.4 Detecting the needle's tip

After finding the line that is most probably aligned to the needle, we search for the position of the tip of the needle.

We first rotate the image containing the stronger eigenvalues, so that the needle is oriented with one of the image axes, which will ease the task, since we only have to slide a small window from one border of the image to the other and follow the ridge in a narrow band surrounding the line found.

During this process, the values beneath this sliding window area are stitched together, always centering the closest local maximum value. Of course, the masks defined earlier have to be taken into account, so that the algorithm would not accidentally snap to e.g. the contour of the probe head. This is simply done by repeating the last valid values as many times as needed and stick to the last detected needle center that was detected before entering the masked out area.

After this procedure, we have a one dimensional cross section of the needle ridge (see Figure 5). Since after the tip there should be a significant step among the values registered, the needle's tip recognition is based on these values.

Since the values are still subject to noise and there are many small deviations of the data, it is smoothed with a moving average filter. Afterwards the local minima and maxima are detected and the differences between their values are computed. It is now supposed that the step, that marks the end of the needle provides a very steep slope with a very low value at the end. Since we would not want to threshold these values as the variations of the data are too high, we look for the steepest descent and its local minimum and the preceding maximum. Due to the Gaussian filtering of the image and the smoothing of the data along the ridge the needle tip will be between these two extrema. The following Langrange line interpolation between the two gradient values worked out very well, as the turning

Eigenvalue-decomposed - Stronger Eigenvalues @level 2      Eigenvalue-decomposed - Smaller Eigenvalues @level 2

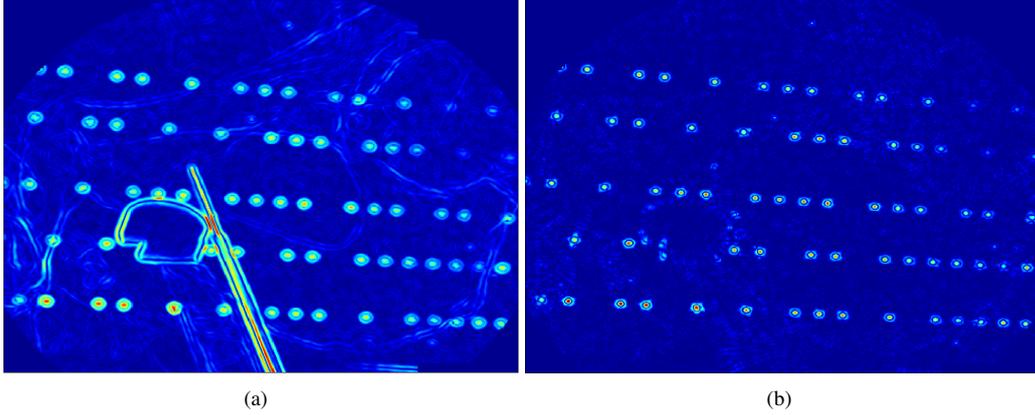


Figure 3: Eigenvalues of the  $|\lambda_1|$  3(a) and  $|\lambda_2|$  3(b) of the computed tensorfield

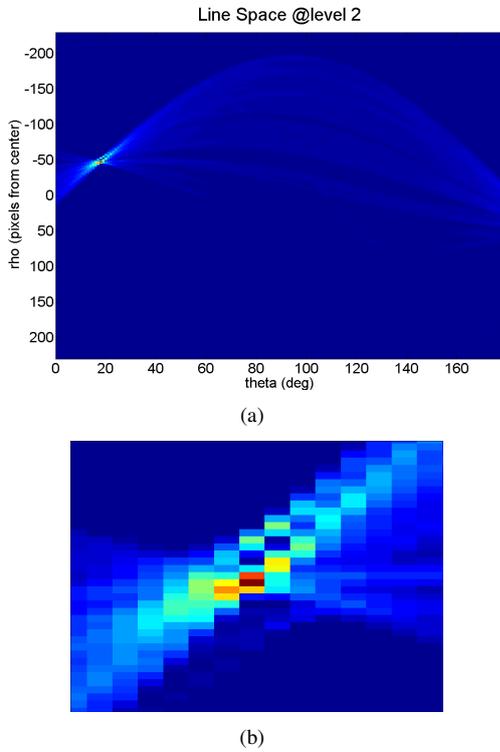


Figure 4: 4(a) Radon transform of the image, 4(b) detailing the maximum peak

point of the curve is a good indicator for the needle tip:

$$x_{\text{tip}} = \frac{x_i y_{i-1} - x_{i-1} y_i}{y_{i-1} - y_i} \quad (8)$$

where  $x_i$  and  $y_i$  refer to the values of the minimum of the differences and  $x_{i-1}$  and  $y_{i-1}$  refer to the maximum and  $x_{\text{tip}}$  is the position of the zero crossing along the needle.

## 4.5 3D-Reconstruction of the needle

After having performed the previous steps to the images it is now possible to reconstruct the needle (i.e., estimat-

ing the position and orientation of the needle in three dimensional space). As already mentioned earlier, this work is part of a framework and a lot of work is done and information is gathered before needle detection is started. The camera calibration matrix and the external camera matrix (i.e., the rotation matrix and the translation vector) are available and used for the reconstruction. Of course, this implies the projection matrices that can be computed out of the internal and external camera matrices are exact, but it is sufficient to consider the coordinates from the images to be erroneous.

### 4.5.1 Reconstruction of the needle tip in 3D-space

The solution to the problem is acquired applying the Direct Linear Transform (DLT), see Algorithm 4.1 for the way we reconstructed a single point. Please note that this is a modified form of the algorithm presented in [6], since we only have a single point in each image, we want to reconstruct.

### 4.5.2 Reconstruction of the needle orientation in 3D-space

In Section 4.3, we have already come up with a fairly robust solution to detecting the line that supports the needle. In this Section we will show, how the orientation of the needle can be recovered.

For the reconstruction of a line we need to know two points on that line in every image. If these point pairs happen to correspond to every other pair of the other images, we can easily recover the line by reconstruction of each point with the method described above and compute the line vector with basic vector algebra.

We already have obtained one point in every image and its reconstructed 3D-point so far, but finding a second point could turn out difficult. One could say, that the guidance of the needle is thicker than the needle and detecting the obviously characteristic positions of where the needle

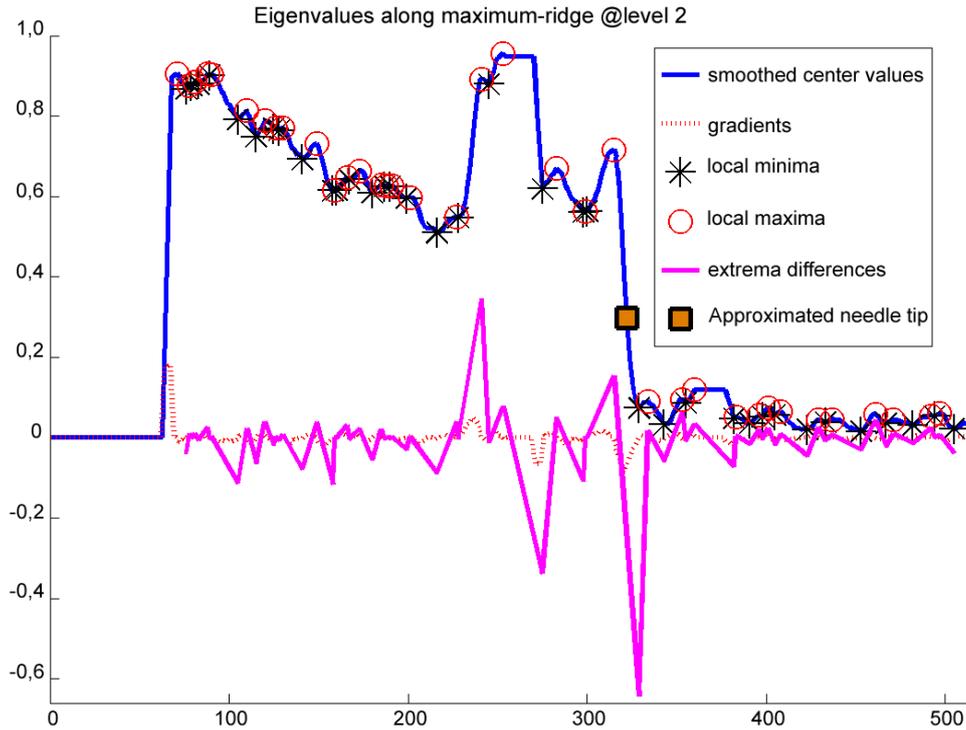


Figure 5: Assembled values beneath the needle ridge with marked extrema and backward-differences The square indicates the estimated needle tip.

---

**Algorithm 4.1** Direct Linear Transform triangulation for one point

---

**Require:**

image coordinates  $\{\mathbf{x}_1 = (u_1 v_1)^T, \dots, \mathbf{x}_n = (u_n v_n)^T\}$ ,  
 projection matrices  $\{P_1, \dots, P_n\}$

**for**  $i = 1, \dots, n$  **do**

- compute the point normalizing transform matrices

$$T_i = \begin{pmatrix} 1 & 0 & -u_i \\ 0 & 1 & -v_i \\ 0 & 0 & 1 \end{pmatrix}$$

- transform and normalize the projection matrices

$$\tilde{P}_i = T_i P_i, \hat{P}_i = \frac{\tilde{P}_i}{\sqrt{\text{tr}(\tilde{P}_i^T \tilde{P}_i)}}$$

**end for**

- stack each normalized projection matrix' first and second row vector in a  $2n \times 4$ -dimensional matrix

$$A = -(\hat{p}_{1,1} \hat{p}_{1,2} \dots \hat{p}_{n,1} \hat{p}_{n,2})^T$$

- compute the singular value decomposition of A

$$A = UDV^T$$

**return** with  $D = \text{diag}(d_1 \dots d_n)$  and  $d_i \geq d_{i+1}$

$$X = V(0 \dots 1)^T$$


---

leaves the guidance leads to several other point correspondences that would be sufficient for the task.

However, this approach seems far too vulnerable to mis-detections, and there exists an easier way, which is only subject to wrong output, if the previous stages fail. This again utilizes the SVD and computes the optimal orientation of the needle in a least-squares sense:

- Extend the calibration matrix to a homogeneous  $4 \times 4$ -matrix and recompute the projection matrix - which will now be invertible - for each image.
- Compute the camera center for each image.
- Compute the 3D-coordinates for the two points of the needle's line crossing the border of their respective image with the inverted projection matrix. An arbitrary  $z$ -value can be chosen.
- Compute the normals of the planes that are spanned by these two vectors and the camera projection center.
- Stack the normals into a new matrix and compute the SVD of this matrix.
- The column of  $V$  that corresponds to the smallest eigenvalue of the matrix will be the best fit to the normals.

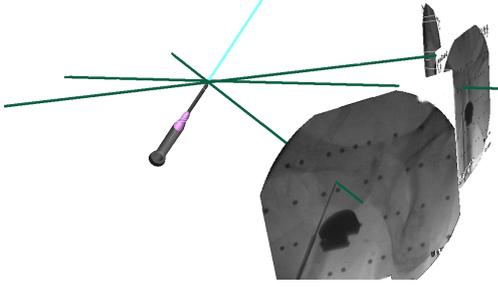


Figure 6: Final output visualized

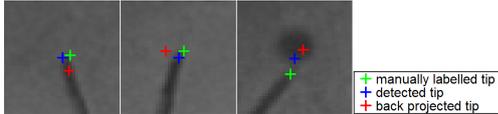


Figure 7: Back-projection of 3D-point of the computed needle tip

#### 4.6 Improving robustness

After reconstructing the three dimensional point of the needle tip, it is easy to project this point back onto the images, so we are able to compute the distance between the initially detected needle tip in an image and the back projected needle tip. If we only let for images with this distance below a previously specified upper limit and reject groups of images, where this value is outside on at least one image, this makes the procedure very robust towards outliers. If needle detection fails at any stage, it does not seem very likely that it fails on all three images in a very similar way, i.e., detecting an entirely different correspondence. Furthermore, this did not happen during the experiments at any time, so this assumption is considered safe.

### 5 Experimental results

Following the development of the pipeline described in the previous section, it was implemented in MATLAB. After obtaining the results from the detection and reconstruction steps, an intermediate visualization was made with Coin3D (Figure 6) to get an impression thereof. A very approximate model of the biopsy needle is placed with its tip at the point and orientation reconstructed. As shown in Figure 6, the values obtained seem very promising at a first glance. Designed as a proof-of-concept application, it was not supposed to perform fast. Nevertheless, the processing time of about ten seconds per image triplet is already considered sufficient for surgery assistance purposes.

The data available to test consisted of 87 images of three patients with shots from three different views. To get a more meaningful evaluation the tip of the needle has been marked by hand on every image. The three available coordinates are shown in Figure 7. For evaluating the accuracy

	$d(\mathbf{x}, \hat{\mathbf{x}})$	$d(\mathbf{x}, \hat{\hat{\mathbf{x}}})$	$d(\hat{\mathbf{x}}, \hat{\hat{\mathbf{x}}})$
minimum	0	0.846	0.582
maximum	10.438	12.416	8.464
mean	3.509	4.227	3.638
std deviation	2.506	2.449	2.146

Table 1: Distances between the hand labelled needle tip the initial detection in the first column and the back projection of the reconstructed needle tip in the second. The third column shows the difference between first detection and backprojection. All values are measured in pixels.

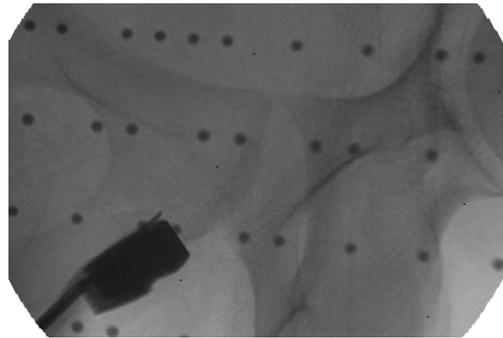


Figure 8: Needle occluded by the ultrasonic device

it is interesting to look at the distances between these coordinates in the images, as the ground truth is not available in 3D-coordinates. Following the initial needle tip detection, the distance can be computed between the detected coordinate  $\hat{\mathbf{x}}$  and the labelled coordinate  $\mathbf{x}$ . To evaluate the results of the reconstructional step backprojection of the reconstructed needle tip from 3D-space coordinate  $\mathbf{X}$  to the respective image coordinates  $\hat{\hat{\mathbf{x}}}$  is performed and the distance to the labelled coordinate is computed again. As we have already mentioned, the distance between  $\hat{\mathbf{x}}$  and  $\hat{\hat{\mathbf{x}}}$  is crucial to the robustness of the algorithm. It is used as an estimate of the confidence in the reconstructed data.

In Table 5 the underlying values are already restricted to results that passed outlier rejection. Although the maximum deviation of the reconstructed tip is very high, Approximately 68% of the reconstructed data lies within five pixels of the needle tip, which is roughly the thickness of the needle, which is most probably accurate enough for many applications.

As shown in Figure 8, it is possible, that the ultra sonic device overlaps the needle in a way that makes needle detection and thus reconstruction impossible. At the moment this case is not handled separately and simply rejected.

### 6 Conclusions

This paper presents the design and implementation of the processing pipeline of a needle detection and reconstruction system intended to be used as one component in a

framework for computer aided biopsy. It takes several established techniques to accomplish this task, such as the Radon transform to detect line features in the image and the application of the Hessian matrix that enhances the input to the former with respect to ridge detection, making it more robust. The parts involved to reconstruct the needle's tip and orientation are well known nowadays and a lot of research has already been put into it, so it can be said that the methods are very reliable.

Although experimental results show very good accuracy in reconstruction and robustness to noise and clutter there are still many problems to solve, involving making the detection of the needle's line and the probe head more independent of scale and further improve the detection of the needle's tip, so that the reconstruction can be done with higher accuracy. However, the system is supposed to be running in a closed loop, i.e., the results are available very shortly after taking the X-rays, so it would be possible to take another set of X-rays if the results are not convincing due to a lack of accuracy because of e.g. the needle tip being occluded by a bead in one of the images.

Clearly, the most considerable result would be the exact and standardized extraction of prostate tissue samples and the precise re-localization of that tissue during surgery.

Future work is planned to deal in more detail with robustness against occlusion and blurred images, which at the moment is handled by simply rejecting the entire imageset that caused the backprojection to differ from the detected needle tips too much.

Further refinement of the needle-tip-detection algorithm is needed, as it is necessary to detect it within sub-millimeter accuracy, as well as the incorporation of additional data from the ultrasound device that would be already available during biopsy.

## Acknowledgements

We would like to thank Markus Grabner for providing most valuable advice and being an excellent mentor during this project. Additionally, we want to thank Bernhard Kainz for supplying a lot of material considering this task from the medical point of view and Matthias R  ther for his help with solving multiview geometry reconstruction problems.

## References

- [1] P. J. Burt and E. H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, COM-31,4:532–540, 1983.
- [2] S.R. Deans. *The Radon Transform and Some of its Applications*. Wiley, 1983.
- [3] R. Ebrahimi, S. Okazawa, R. Rohling, and S.E. Salcudean. Hand-held steerable needle device. *Medical*

*Image Computing and Computer-Assisted Intervention (MICCAI)*, 2879/2003:223–230, 2003.

- [4] D. Glozman and M. Shoham. Image-guided robotic flexible needle steering. *IEEE Transactions on Robotics*, 23(3):459–467, June 2007.
- [5] T. Gross. GPU-based 2D/3D registration of X-Ray and CT data. Master's thesis, Graz University of Technology, Institute for Computer Graphics and Vision, Graz, Austria, January 2008.
- [6] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [7] A. M. L  pez, F. Lumbreras, J. Serrat, and J. J. Villanueva. Evaluation of methods for ridge and valley detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(4):327–335, 1999.
- [8] D. Lowe. Object recognition from local scale-invariant features. In *Proc. Seventh IEEE International Conference on Computer Vision The*, volume 2, pages 1150–1157, 20–27 Sept. 1999.
- [9] J. B. Antoine Maintz and Max A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [10] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *International Journal of Computer Vision*, 65(1–2):43–72, 2005.