Segmentation of Tubular Structures in 3D Images Using a Combination of the Hough Transform and a Kalman Filter

Thorsten Behrens[1], Karl Rohr[2], and H. Siegfried Stiehl[1]

[1] Universität Hamburg, FB Informatik, AB KOGS, Vogt-Kölln-Straße 30, D-22527 Hamburg, {behrens,stiehl}@kogs.informatik.uni-hamburg.de

[2] International University in Germany, D-76646 Bruchsal, rohr@i-u.de

Abstract. In this paper, we present a new approach for coarse segmentation of tubular anatomical structures in 3D image data. Our approach can be used to initialise complex deformable models and is based on an extension of the randomized Hough transform (RHT), a robust method for low-dimensional parametric object detection. In combination with a discrete Kalman filter, the object is tracked through 3D space. Our extensions to the RHT feature adaptive selection of the sample size, expectation-dependent weighting of the input data, and a novel 3D parameterisation for straight elliptical cylinders. For initialisation, only little user interaction is necessary. Experimental results obtained for 3D synthetic as well as for 3D medical images demonstrate the robustness of our approach w.r.t. image noise. We present the successful segmentation of tubular anatomical structures such as the aortic arc or the spinal chord.

Keywords: 3D medical images, 3D tubular structure segmentation, minimal user interaction, randomized Hough transform (RHT), Kalman filter-based tracking

1 Introduction

Deformable models are often used to segment objects in complex 2D and 3D images (e.g. [10]). Since usually local optimisation methods are employed for deformable model fitting, model initialisation is generally required to be close to the real object to obtain reasonable fitting results. In particular, initialisation becomes a major problem for elongated and complicately shaped objects (such as long blood vessels), which are generally described by a large set of parameters. Typically, these model parameters have to be initialised manually. Thus, there is a clear need for automated methods, yielding an approximate segmentation of complex shaped tubular objects, while requiring *minimal* user interaction.

Only a few approaches (e.g. [14, 3]) consider the segmentation of tubular structures in 3D medical image data with minimal user interaction. However, their drawbacks are that either an ad-hoc slice-tracking procedure is used in

conjunction with threshold-based determination of the tube wall [14], or the method works only for a previously given scale of the tube diameter [3].

In this contribution, we introduce a new approach for segmentation of tubular structures in 3D image data. To detect objects in 3D image data, we extend in Sect. 2 the randomized Hough transform (RHT), a robust method for low-dimensional parametric object detection, while in Sect. 3 we describe how this method can be combined with a discrete Kalman filter to track the objects through 3D space. Segmentation is achieved by detecting elliptical cross sections or straight elliptical cylinder segments that are subsequently tracked through the 3D image. Our algorithm has been applied to both 3D synthetic and 3D medical images (Sect. 4).

2 Extensions of the Randomized Hough Transform

The Hough transform is a well-known method for parametric object detection, where detected pixels in the input image are mapped to a discrete parameter space, whose maxima represent object candidates. In order to detect ellipses or straight elliptical cylinders, which have a relatively large number of parameters (leading to *excessive* space/time complexity for the conventional Hough transform), we build upon the so-called randomized Hough transform (RHT) [16].

The RHT uses a randomly chosen subset of the input data, together with a so-called many-to-one sampling scheme that maps input point sets to zero-dimensional point sets in the parameter space. Thus, parameter spaces of higher dimensions remain tractable, if a dynamic accumulation scheme is used. Several drawbacks inherent to coarse-to-fine and parameter-space decomposition Hough transforms, such as increased noise sensitivity and projection artifacts [4, 9], can be alleviated by the RHT, because the parameter space can always have both full dimension and resolution. Our proposed extensions include a new object parameterisation for straight elliptical cylinders, new derivations for an adaptive sample size, and a novel input weighting scheme.

2.1 Object Parameterisations

For the object parameterisation of an *ellipse*, we apply an approach proposed for the Hough transform in [8, p. 151], which determines a unique ellipse (if there is any) that passes through a set of five coplanar points. The advantages over other methods are that no additional information (such as local edge direction) is needed, and because the problem is reduced to linear equations, standard numerical methods can be applied. As discussed in Sect. 3, the plane from which the samples are taken should lie orthogonal to

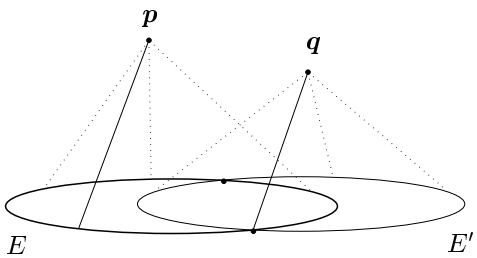

Fig. 1: Determining the cylinder axis

the tube axis. The *straight elliptical cylinder* is a degenerate quadric, so that a direct approach in the above sense leads to a rather complex non-linear algebraic problem. For the sake of speed, we therefore devised the following two-step algorithm for straight elliptical cylinder parameter calculation, using the fact that it is a ruled surface, but still avoiding the use of local edge direction (which is hard to estimate in a robust way with the required accuracy):

1. Select randomly five coplanar points and determine the ellipse E through them (if there is any).
2. Select randomly two distinct points $\boldsymbol{p}$ and $\boldsymbol{q}$ outside the plane containing E. Sweep a line through $\boldsymbol{p}$ and points along ellipse E, and calculate the ellipse E' in the same plane generated by sweeping simultaneously a parallel line through $\boldsymbol{q}$ (see Fig. 1). The intersection points of these two ellipses (it can be shown that there are at maximum two) determine the possible cylinder axes: A line through one of those intersections and $\boldsymbol{q}$ (solid line in Fig. 1) is parallel to the cylinder axis, which of course intersects the two ellipses' plane at the center of E. Having determined the axis, the remaining cylinder parameters can easily be calculated.

2.2 Adaptive Sample Size

When using random sampling schemes for the Hough transform, it is normally no longer necessary to consider *all* points of the input data set. Usually, a subset will suffice, whose size depends on the input data quality. This subset is typically generated by drawing a sample of a given size from the input data points. Instead of the formulas given in [15], which are only hints to sample size selection, we suggest a different way to derive an estimation of an appropriate (data-dependent) sample size. This enables us to state an upper bound for the probability of false detections. Although the derivation is strictly correct only for an unlimited sample size, in practice the given error limit is almost always valid.

Because the counts in the accumulator cells are binomially distributed [15], we can approximate an individual cell's count for large sample sizes by the normal distribution. Thus, if we know *a priori* the probability p_s for sampling points from a significant object and the probability p_{ns} for sampling points from a non-significant object, it can be shown that if the sample size satisfies

$$n \geq \frac{z^2 \left(\sqrt{p_s \left(1 - p_s\right)} + \sqrt{p_{ns} \left(1 - p_{ns}\right)} \right)^2}{\left(p_s - p_{ns}\right)^2} \quad , \tag{1}$$

then with probability of at least $1 - \mathrm{erf}(z/2)$ the counts of significant cells are larger than those of non-significant ones. This actually prevents non-significant objects from being erroneously detected, because for the RHT, the detected object coincides with the maximal accumulator count. The value z thus parameterises the sample size and the remaining error probability.

We can refine this a priori sample size estimate by calculating the probability p'_s, which denotes the *actual* probability for sampling points from an object

corresponding to a given cell, using the theorem of the iterated logarithm by Khintchine (e.g. [2, pp. 204]):

$$p'_s \geq \frac{nk + n\log(\log(n)) - \sqrt{n\log(\log(n))\,(2nk - 2k^2 + n\log(\log(n)))}}{n^2 + 2n\log(\log(n))} \quad , \qquad (2)$$

where n denotes the current sample size and k the current cell count. As it is only required for the most significant object's cell count to be larger than non-significant cells, we can substitute for p_s in (1) the estimate p'_s calculated from the maximum accumulator cell count to determine a data-dependent sample size.

2.3 Windowing the Hough Transform

In images from cluttered environments, non-relevant image structures often distract the RHT from the relevant features, especially when the objects to be detected are tiny compared to the image size. Therefore, masking out all but small areas around the features has proven to be advantageous [6, 7], where usually rectangular, binary-valued windows are used. We extend masking towards arbitrary discrete functions that are adapted to the problem at hand: As the input pixels are selected randomly for the RHT, we have to induce an appropriate discrete probability density function onto the input image, with high values at the most possible object locations. These arbitrary discrete probability functions can be realised by employing the so-called alias method described in [13].

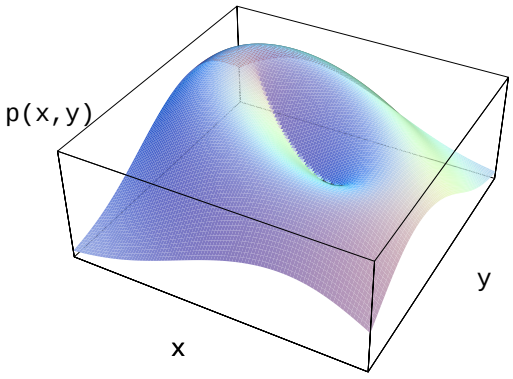

Fig. 2: Example of a continuous window function for ellipse detection

As the Kalman filter (explained below) yields estimates of the expected position, size, and shape of the input object in the next slice, the image can be windowed appropriately. Assuming a Gaussian distribution of the ellipse's main axes errors and exact estimates of the center position and the rotation angle, the window function in Fig. 2 shows the incorporation of Kalman prediction into the RHT.

3 Kalman Filter Approach

In order to track the detected objects (either ellipses or straight elliptical cylinders) through 3D image data, we apply a discrete Kalman filter [5]. Apart from the useful prediction information which can be exploited as mentioned in Sect. 2.3, this approach allows to include an explicit model of the tubular structure's axis curve. For modeling this axis, we used both a linear and a quadratic Taylor approximation.

If $\boldsymbol{x}_k$ is a point on the tube axis at time $k\Delta t$, an estimate at instant $(k+1)\Delta t$ can be calculated as follows (quadratic model):

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \Delta t \boldsymbol{x}'_k + \frac{1}{2}\Delta t^2 \boldsymbol{x}''_k, \quad \boldsymbol{x}'_{k+1} = \boldsymbol{x}'_k + \Delta t \boldsymbol{x}''_k, \quad \boldsymbol{x}''_{k+1} = \boldsymbol{x}''_k \ , \qquad (3)$$

where Δt denotes the time interval and $\boldsymbol{x}'_k$, $\boldsymbol{x}''_k$ the first and second derivative w.r.t. time, resp.

Therefore, taking together the results from the previous sections, we were able to implement a tube tracker that adapts itself to the input data quality and restricts random sampling effectively to the Kalman-predicted area. Furthermore, the Kalman prediction of the axis curve offers an additional advantage: If the intersection plane is not orthogonal to the cylinder axis, even tubular structures with elliptical cross sections generally do not yield ellipses when intersected by a plane (see e.g. [12] or [11] for a more in-depth discussion). Thus, when employing the detection of elliptical cross sections, it is essential to re-slice the 3D input data locally, orthogonal to the Kalman prediction of the axis curve. Otherwise, when tracking highly curved structures, the plane intersection curve will differ significantly from an ellipse after a few steps. This can lead to inaccuracies due to model mismatch.

4 Experimental Results

We performed about 20,000 experiments using 3D synthetic binary data with varying object sizes and noise levels to assess the performance and robustness for all four possible combinations of the described object parameterisations (ellipse and straight elliptical cylinder) and the axis approximation (linear and quadratic). It turned out that the detection of elliptical cross sections and the tracking with the quadratic model in (3) performed best. For this variant we show in Fig. 3 one of the binary input images used (slices of a torus), which was degraded with 10% shot noise, and the rendered segmentation result. The plot on the right hand side of Fig. 3 depicts the mean fraction of successfully segmented 3D input images (solid line). This means that the whole arch was tracked and the detected axis was never farther away from the real axis than twice the radius. The mean of that distance from the true arch axis is depicted by the dashed line, whereas the vertical error bars for both plots denote the standard deviation from the corresponding mean. The results are promising, especially regarding noise insensitivity. Up to a noise level of 15%, the test structures were always fully tracked. The small systematic error w.r.t. the true torus axis (≈ 0.5 pixel) is due to discretisation artifacts in the RHT accumulator. In comparison, the other variants were more sensitive to noise. However, for small noise levels, tracking elliptical cylinder segments was significantly more accurate than tracking elliptical cross-sections.

We also conducted experiments using 3D Magnetic Resonance (MR) data, applying the same approach as above (detection of elliptical cross sections and tracking with the quadratic model). To generate the binary data sets necessary

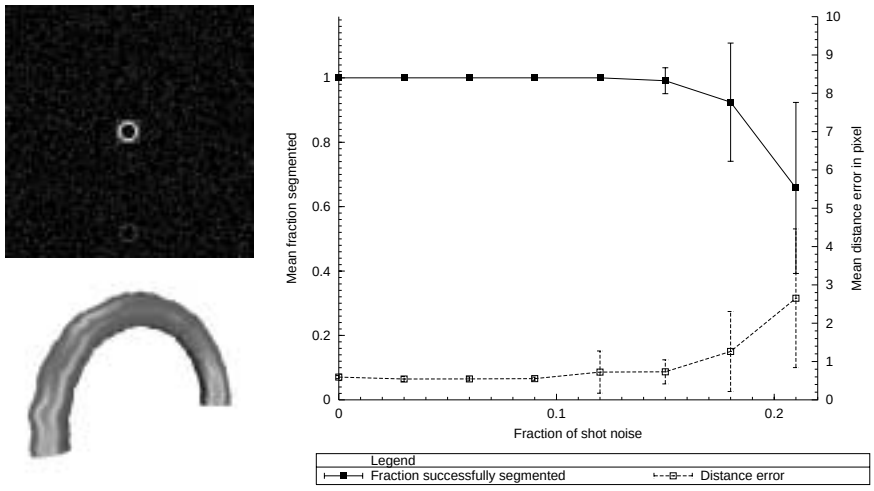

Fig. 3. Experiments using 3D synthetic data. The two images on the left hand side show at the top a slice of a 3D input data set (torus, degraded with 10% shot noise) and below the segmentation result. On the right hand side, a plot of two statistics for the experiments with the method using detection of elliptical cross sections and the quadratic axis model is given, where every point is the mean value of 20 to 40 separate experiments.

for the algorithm, we used a 3D extension of the edge detection method described in [1]. Figure 4 shows the segmentation result of the human aorta in a thorax MR angiography after 3D edge detection. The only initial parameters required for the algorithm were the start point of the aorta axis, the initial radius, and an approximate initial direction. The algorithm then continued segmentation through 3D space until the bottom of the image data volume was reached.

Using the same approach, Fig. 5 shows the result for a 3D MR image of the human head, where the spinal chord was segmented, starting from the bottom of the image up to the medulla oblongata.

5 Summary

We have proposed a new segmentation method for tubular structures in 3D image data, which is based on an extension of the randomized Hough Transform and a discrete Kalman filter. Experimental results obtained for both 3D synthetic and 3D MR image data showed promising results, especially regarding the robustness w.r.t. noise. Out of the four variants we investigated, the one using a quadratic axis model and the detection of elliptical cross sections performed best. Because of its robustness and the small number of required initial values, we characterize our novel algorithm to be well suited for coarse segmentation as well as for initialisation of complex deformable model fitting.

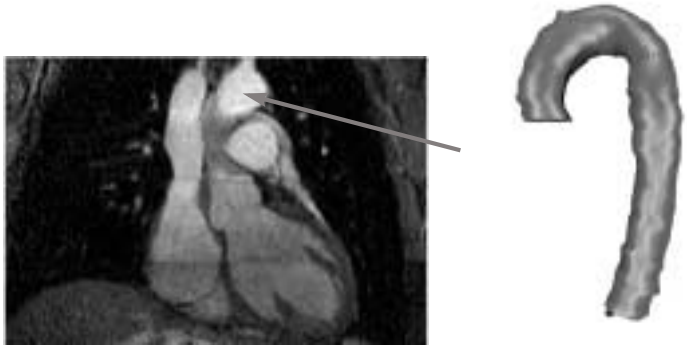

(a) Coronal slice of the thorax (b) 3D segmentation of the aorta

Fig. 4. Experiments using real 3D MR-angiography data of the thorax. CPU time (without prior 3D edge detection) for segmentation of the aorta was 347.29 sec on a 300 MHz Sun Ultra 2 workstation. The arrow marks the position of the aorta when crossing the image plane.

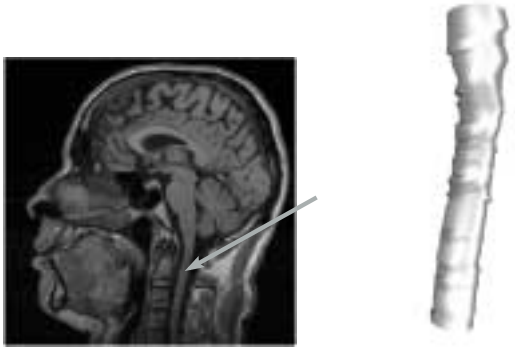

(a) Sagittal slice of the human head (b) 3D segmentation of the spinal chord

Fig. 5. Experiments using real 3D MR data of the human head. CPU time (without prior 3D edge detection) for segmentation of the spinal chord (see arrow) was 108.46 sec on a 300 MHz Sun Ultra 2 workstation.

Acknowledgement

The 3D MR and MRA image data has been kindly provided by Dr. T. Schäffter from Philips Research Laboratories Hamburg. Many helpful comments by Sönke Frantz are gratefully acknowledged.

References

1. J. F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):679–697, 1986.
2. W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, Inc., New York, 3rd edition, 1968.
3. M. Hernández-Hoyos, A. Anwander, M. Orkisz, J.-P. Roux, P. Douek and I. E. Magnin. A deformable vessel model with single point initialization for segmentation, quantification and visualization of blood vessels in 3D MRA. In *Proc. MICCAI 2000*, LNCS 1935, pp. 735–745. Springer-Verlag, 2000.
4. J. Illingworth and J. Kittler. A survey of the Hough transform. *Computer Vision, Graphics and Image Processing*, 44(1):87–116, 1988.
5. R. E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME, Journal of Basic Engineering*, (82):35–45, 1960.
6. H. Kälviäinen, P. Hirvonen, L. Xu and E. Oja. Comparisons of probabilistic and non-probabilistic Hough transforms. In *Proc. 3rd ECCV*, LNCS 800, pp. 351–360. Springer-Verlag, 1994.
7. V. Kyrki and H. Kälviäinen. Combination of local and global line extraction. *Real-Time Imaging*, 6:79–91, 2000.
8. V. F. Leavers. *Shape Detection in Computer Vision Using the Hough Transform*. Springer-Verlag, London, 1992.
9. V. F. Leavers. Survey: Which Hough transform? *Computer Vision, Graphics and Image Processing: Image Understanding*, 58(2):250–264, 1993.
10. T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996.
11. T. O' Donnell, A. Gupta, and T. Boult. A new model for the recovery of cylindrical structures from medical image data. In *Proc. CVRMed-MRCAS'97*, LNCS 1205, pp. 223–235. Springer-Verlag, 1997.
12. B. I. Soroka, R. L. Andersson, and R. K. Bajcsy. Generalised cylinders from local aggregation of sections. *Pattern Recognition*, 13(5):353–363, 1981.
13. A. J. Walker. An efficient method for generating discrete random variables with general distributions. *ACM Trans. on Mathematical Software*, 3(3):253–256, 1977.
14. O. Wink, W. J. Niessen, and M. A. Viergever. Fast quantification of abdominal aortic aneurysms from CTA volumes. In *Proc. MICCAI 1998*, LNCS 1496, pp. 138–145. Springer-Verlag, 1998.
15. L. Xu and E. Oja. Randomized Hough transform (RHT): Basic mechanisms, algorithms, and computational complexities. *Computer Vision, Graphics and Image Processing: Image Understanding*, 57(2):131–154, 1993.
16. L. Xu, E. Oja, and P. Kultanen. A new curve detection method: Randomized Hough transform (RHT). *Pattern Recognition Letters*, 11(5):331–338, 1990.