

Reducing False Detections in Extracting 3D Anatomical Point Landmarks*

Sönke Frantz, Karl Rohr, and H. Siegfried Stiehl

Universität Hamburg, Fachbereich Informatik, Arbeitsbereich Kognitive Systeme
Vogt-Kölln-Str. 30, D-22527 Hamburg, Germany
{frantz,rohr,stiehl}@informatik.uni-hamburg.de

Abstract. Applying 3D differential operators to extract point landmarks from medical images generally suffers from false detections. A considerable number of these false detections is caused by neighboring structures that are included in the region-of-interest (ROI) specified by the observer. The main contributions of this paper are two different approaches to reducing false detections resulting from neighboring structures. First, we present a statistical differential approach to selecting a suitable ROI size automatically. Second, we propose a differential approach to incorporating prior knowledge of the intensity structure at a landmark. Also, to cope with anisotropic voxel sizes in estimating partial derivatives, we implemented a computationally efficient scheme based on cubic B-spline image interpolation. Experimental results based on 3D MR and CT images of the human head are presented.

Keywords: Point landmarks, differential operators, false detections

1 Introduction

Anatomical landmarks are useful features for a wide spectrum of applications in medical image analysis. More specifically, we are interested in 3D point landmarks of the human head, which can be used for point-based elastic image registration. Here, prominent landmarks are, e.g., the tips of the lateral ventricles, the tip of the external occipital protuberance, or the saddle points at the zygomatic bones. However, manually extracting such landmarks from images is generally very tedious and often prone to error. Consequently, automating landmark extraction is of central interest. To this end, computationally efficient 3D differential operators for point landmark detection were introduced [1],[2]. These operators employ only first order partial derivatives of the intensity function and hence are relatively robust w.r.t. noise (see [3] for a comparative performance analysis of various differential operators for landmark detection). We apply our operators within a semi-automatic procedure: (i) the observer determines the landmark position only roughly, (ii) a differential operator is applied within a region-of-interest (ROI) to detect landmark candidates, and (iii) the observer selects the best candidate. However, this procedure generally suffers from false

* This work has been supported by Philips Research Hamburg, Project IMAGINE (IMage- and Atlas-Guided Interventions in NEurosurgery).

detections caused by neighboring structures that are included in the ROI. The main contributions of this paper are two different approaches to reducing false detections resulting from neighboring structures. First, we address the problem of selecting a suitable ROI size that enables reliable landmark detection but excludes neighboring structures. We present a statistical differential approach to selecting a suitable ROI size automatically (Sec. 2). Second, we present a differential approach to incorporating prior knowledge of the intensity structure at a landmark. Detected candidates with an inconsistent intensity structure are rejected automatically (Sec. 3). Also, to cope with anisotropic voxel sizes in estimating partial derivatives, we implemented a computationally efficient scheme based on cubic B-spline image interpolation (Sec. 4). Experimental results based on 3D MR and CT images of the human head demonstrate the efficacy of our approaches in reducing false detections (Sec. 5).

2 Automatically selecting a suitable ROI size

The observer determines the position of the landmark only roughly. At this position, a ROI is placed where the differential operators are applied to detect landmarks. Usually, a fixed ROI size is used. Because of this, however, neighboring structures are often included in the ROI, which gives rise to additional detections. We propose an approach to automatically selecting a suitable ROI size, based on a 3D differential approach developed for refined landmark localization [4]. With this differential edge intersection approach, tangent planes are defined to locally approximate the surface at a landmark. The tangent planes are intersected to estimate the landmark position. This approach can be used for ROI size selection by taking the statistical uncertainty of the position estimate as a criterion for the consistency of the present intensity structure w.r.t. the expected surface geometry at the landmark (e.g., a tip). Preliminary results based on a 2D version of the approach were reported in [5].

Suppose we have placed a cubic ROI of width w , say, at the manually determined position. Let \mathbf{x}_w^* denote the position estimate from the 3D differential edge intersection approach [4]. The localization uncertainty of \mathbf{x}_w^* is given by the covariance matrix, $\mathbf{\Sigma}_w = \sigma_\varepsilon^2 (\sum_i \nabla g(\mathbf{x}_i) \nabla g(\mathbf{x}_i)^T)^{-1}$, where σ_ε^2 is a data-dependent noise term and $\nabla g(\mathbf{x}_i)$ denotes the intensity gradient at \mathbf{x}_i . The sum index i addresses all voxels within the ROI. A scalar measure for the localization uncertainty is the determinant of the covariance matrix, $U_w = \det(\mathbf{\Sigma}_w)$ (generalized variance). Exploiting the uncertainty U_w as a function of w , our strategy for selecting the optimal ROI size is as follows (see [5] for details): We start with a user-specified minimal ROI size (e.g., $w = 7$ voxels). If the ROI does not capture enough surface information to reliably estimate the landmark position, we expect that U_w is high. Hence, taking into account more image information by enlarging the ROI, we expect U_w to decrease. However, if neighboring structures finally begin to interfere, we expect a significant increase in U_w , which suggests that further enlarging the ROI is not useful. In our implementation, we detect such a signal change by requiring that (a) U_w increases and (b) the relative spatial variation of the position estimate exceeds a threshold. Finally, we select the ‘optimal’ ROI size based on minimal uncertainty.

3 Incorporating prior knowledge of the intensity structure at a landmark

We take advantage of prior knowledge of the intensity structure at a landmark to impose additional constraints for accepting a detected candidate. In an application, the user generally knows the landmark type (e.g., tip or saddle point) as well as the imaging modality. Here, we distinguish between tips and saddle points. Additionally, we distinguish between tips of dark and bright structures w.r.t. the background. To classify these structures, we exploit the curvature properties of the isointensity surface at a detected candidate.

Suppose we have detected a point \mathbf{x}_d on the surface of an anatomical structure. Let K denote the Gaussian curvature and H denote the mean curvature of the isointensity surface at \mathbf{x}_d (for details on computing differential measures of isointensity surfaces, we refer to, e.g., [6],[7]). Exploiting the sign of K , we distinguish between tips ($K > 0$) and saddle points ($K < 0$). Exploiting the sign of H , we further distinguish between tips of dark ($H < 0$) and bright ($H > 0$) structures w.r.t. the background. The candidate is rejected if the present classification is inconsistent with the expected intensity structure.

4 Coping with anisotropic voxel sizes in estimating partial derivatives

Our approaches described above require partial derivatives of the intensity function. Usually, the derivatives are calculated by applying sampled Gaussian derivative filters. However, problems generally arise when dealing with anisotropic data. Then, the filter sizes have to be adapted, which can result in instabilities. To cope with anisotropic voxel sizes, we implemented a computationally efficient scheme based on cubic B-spline image interpolation, following the work in [8]. With this approach, the partial derivatives are calculated based on the reconstructed continuous signal. Since the B-spline basis functions are separable, we only describe the 1D case and briefly point out the extension to 3D.

Let $\hat{g}(x)$ denote the cubic B-spline interpolated intensity function. Suppose the sampling distance is δ_x (in mm), i.e., the spatial variable in $\hat{g}(x)$ is given in δ_x units. To further smooth $\hat{g}(\cdot)$, we apply a Gaussian filter $G_\sigma(x) = 1/(\sqrt{2\pi}\sigma) \exp(-x^2/(2\sigma^2))$, which we also represent in terms of cubic B-splines as $\tilde{G}_\sigma(x)$. To take into account the different coordinate scalings, the Gaussian is rescaled by adapting the standard deviation σ (in mm) to $\sigma' = \sigma \cdot 1\text{mm}/\delta_x$. The whole computational chain for calculating the derivatives at the points of the discrete image can be expressed by convolving the B-spline coefficients of $\hat{g}(x)$ with a discrete filter. This filter is derived from (a) the B-spline coefficients of $\tilde{G}_{\sigma'}(x)$, (b) the finite difference operator, (c) the discrete B-spline, and (d) a normalization constant depending on δ_x . The extension to 3D consists of three 1D convolutions in x -, y -, and z -direction with such a discrete filter. Hence, calculating the derivatives this way has the same complexity as ‘conventional’ derivative calculation. The only additional computational burden lies in computing the B-spline coefficients of the image.

5 Experimental results

We applied our approaches to 3D synthetic images and 3D tomographic images. Here, we study the efficacy of our approaches in reducing false detections in extracting anatomical landmarks from medical images. Exemplarily, we consider the tip of the external occipital protuberance (MC5e), the tip of the left frontal horn of the ventricular system (MC6l), and the saddle point at the left zygomatic bone (MC15l). We report on experiments with a 3D T1-weighted MR image ($256 \times 256 \times 120$ voxels, voxel resolution $0.86 \times 0.86 \times 1.2 \text{mm}^3$) and a 3D CT image ($320 \times 320 \times 87$ voxels, voxel resolution $0.63 \times 0.63 \times 1 \text{mm}^3$). To detect landmarks, the differential operator $Op3 = \det(\mathbf{N})/tr(\mathbf{N})$ was used [2]. Here, \mathbf{N} denotes the averaged dyadic product of the intensity gradient, $\mathbf{N} = \overline{\nabla g \nabla g^T}$, and $\det(\cdot)$ and $tr(\cdot)$ denote the determinant and the trace of a matrix, resp. The scale of the Gaussian derivative filters was set to $\sigma = 1.0 \text{mm}$ (which, however, was adapted according to the voxel resolution). Averaging the gradient was done within a neighborhood of $5 \times 5 \times 5$ voxels. Landmark candidates were determined by searching for local maxima of the operator responses in $3 \times 3 \times 3$ neighborhoods. No thresholds were applied to the operator responses.

Reducing false detections by selecting a suitable ROI size First, we manually determined the landmark positions. Then, a ROI was placed at those positions. The ROI size was automatically selected as described above, starting with $w = 7$ voxels. The maximum ROI size was restricted to $w = 21$. The operator was applied within the automatically selected ROIs and the number of remaining detections is compared with the case when using a maximum ROI size of $w = 21$.

In Fig. 1, the operator responses at the detected positions are drawn as a function of the distance to the manual position. Those candidates that lie within the automatically selected ROI are indicated by bold bars. Other candidates that would additionally be obtained by using the maximum ROI size are indicated by narrow bars. Moreover, the total number of detections as well as the number of remaining detections within the automatically selected ROI are given. In the case of MC5e and MC6l in MR, we obtain optimal ROIs that yield a unique candidate. In the case of MC5e in CT, two candidates remain, both of which are located at this structure. In the case of MC6l in CT and MC15l (both modalities), the automatically selected ROI size corresponds to the maximum ROI size. Further experiments with other data and other landmarks showed similar results, i.e., in a number of cases, adapting the ROI size avoids additional false detections from neighboring structures. Also, we tested the effect of varying the manual positions (e.g., due to different observers) on selecting the ROI size. Our approach appeared to be robust w.r.t. varying manual positions.

Reducing false detections by incorporating prior intensity knowledge A ROI of size $21 \times 21 \times 21$ voxels was placed at the manually determined landmark positions. In the case of MC6l (CT), a ROI of size $21 \times 21 \times 14$ was used since this landmark is located at the border of the image. The operator was applied within the ROIs and at each detected point the present intensity structure was classified.




Fig. 1. Reducing false detections by automatically selecting a suitable ROI size. The operator responses are drawn as a function of the distance to the ROI center. Bold bars indicate the remaining detections within the automatically selected ROI. The numbers in the captions mean *total number of detections/remaining number of detections*.

In Fig. 2, the operator responses at the detected positions are drawn as a function of the distance to the manual position. Those candidates where the present intensity structure corresponds to the expected one are indicated by bold bars. Other candidates that would additionally be obtained without this distinction are indicated by narrow bars. Moreover, the total number of detections as well as the number of remaining detections after classification are given. We see that a large number of detections that partly show significant operator responses are rejected due to an inconsistent intensity structure. However, in either case, those candidates that best detect the landmark, i.e., having minimal distance to the manual position, are correctly classified. Thus, the localization performance has not been affected. In sum, classifying the detections as described above has proven to be very effective in reducing false detections.

Both approaches to reducing false detections, i.e., automatically selecting the ROI size and incorporating prior intensity knowledge, have so far been studied separately. However, combining both approaches can further improve the results. For example, in the case of MC5e in CT, then only one candidate remains.

6 Conclusion

The performance of a semi-automatic procedure for extracting anatomical landmarks from medical images substantially depends on the number of detected landmark candidates. We have proposed two different approaches to reducing false detections. In experiments with 3D tomographic images, automatically selecting the ROI size and incorporating prior knowledge of the intensity structure at a landmark have proven to be effective in reducing false detections.

References

1. K. Rohr, H.S. Stiehl, R. Sprengel, W. Beil, T.M. Buzug, J. Weese, and M.H. Kuhn. Point-Based Elastic Registration of Medical Image Data Using Approximating Thin-Plate Splines. In K.H. Höhne and R. Kikinis, eds., *Proc. VBC'96*, LNCS 1131, pp. 297–306. Springer, 1996.
2. K. Rohr. On 3D differential operators for detecting point landmarks. *Image and Vision Computing*, 15(3):219–233, 1997.




Fig. 2. Reducing false detections by incorporating prior knowledge of the intensity structure at a landmark. The operator responses are drawn as a function of the distance to the ROI center. Bold bars indicate the remaining detections after classification. The numbers in the captions mean *total number of detections/remaining number of detections*.

3. T. Hartkens, K. Rohr, and H.S. Stiehl. Evaluierung der Detektionsleistung von 3D-Operatoren zur Ermittlung anatomischer Landmarken in tomographischen Bildern. In T. Lehmann, V. Metzler, K. Spitzer, and T. Tolxdorff, eds., *Proc. 2. Workshop Bildverarbeitung für die Medizin*, Informatik aktuell, pp. 93–97. Springer, 1998.
4. S. Frantz, K. Rohr, and H.S. Stiehl. Refined Localization of Three-Dimensional Anatomical Point Landmarks Using Multi-Step Differential Approaches. In K.M. Hanson, ed., *Proc. SPIE's Medical Imaging: Image Processing*, vol. 3338, pp. 28–38. SPIE, 1998.
5. S. Frantz, K. Rohr, and H.S. Stiehl. Multi-step Procedures for the Localization of 2D and 3D Point Landmarks and Automatic ROI Size Selection. In H. Burkhardt and B. Neumann, eds., *Proc. ECCV'98*, LNCS 1406, pp. 687–703. Springer, 1998.
6. L.M.J. Florack, B.M. ter Romeny, J.J. Koenderink, and M.A. Viergever. General Intensity Transformations and Differential Invariants. *Journ. of Mathematical Imaging and Vision*, 4:171–187, 1994.
7. J.-P. Thirion. Extremal Points: Definition and Application to 3D Image Registration. In *Proc. CVPR'94*, pp. 587–592. IEEE Computer Society Press, 1994.
8. M. Unser, A. Aldroubi, and M. Eden. B-Spline Signal Processing: Part I—Theory. *IEEE Trans. on Signal Processing*, 41(2):821–833, 1993.