

Closing Curves with Riemannian Dilation: Application to Segmentation in Automated Cervical Cancer Screening

Patrik Malm and Anders Brun

Centre for Image Analysis, Box 337, SE-751 05, Uppsala, Sweden
{patrik,anders}@cb.uu.se

Abstract. In this paper, we describe a nuclei segmentation algorithm for Pap smears that uses anisotropic dilation for curve closing. Edge detection methods often return broken edges that need to be closed to achieve a proper segmentation. Our method performs dilation using Riemannian distance maps that are derived from the local structure tensor field in the image. We show that our curve closing improves the segmentation along weak edges and significantly increases the overall performance of segmentation. This is validated in a thorough study on realistic synthetic cell images from our Pap smear simulator. The algorithm is also demonstrated on bright-field microscope images of real Pap smears from cervical cancer screening.

1 Introduction

Cervical cancer is the second most common type of cancer among women. During 2005 it caused a quarter of a million deaths worldwide, of which 80% occurred in developing countries [1]. Current screening programmes mainly use the staining and visual inspection method developed by Papanicolaou during the 1940s. Images from the Papanicolaou (Pap) test, commonly known as the Pap smear, are mostly analyzed manually by experts, which is a costly procedure [2]. In this paper we describe an automatic segmentation algorithm for cell nuclei in such images. In particular we introduce and study the use of anisotropic dilation for curve closing to achieve better segmentations.

When using edge detection algorithms, e.g., the Canny edge detector [3], it is common that the resulting edge map has gaps in the object borders. To achieve a segmentation based on the edge detection result, such flaws need to be corrected. The method described in this paper generates a non-Euclidean distance transform of the edge map, derived from local gradients in the image. The distance map is then used to perform anisotropic dilation, which we call Riemannian dilation since the distances are derived from geodesic distances in a Riemannian manifold [4]. Locally, the metric we use is related, but not identical, to the local structure tensor field [5,6,7] in the image. In this metric, the geodesic disc has an elongated appearance that follows edges and lines in the image. When it is used as structuring element in dilation, it repairs the gaps in the edge map. In

this paper, we show that by modeling the image as a Riemannian manifold in this manner, we achieve better performance for weak edges compared to using ordinary isotropic dilation.

1.1 Related Work

A thorough discussion regarding the concepts of spatially variant morphology can be found in [8]. In [9] a similar approach is described, where the structure tensor field is also used to perform spatially varying morphology in images.

2 Segmentation

The segmentation process can be thought of as consisting of two related processes, recognition and delineation. Recognition is the task of roughly deciding where in an image an object is, whereas delineation is the process of determining the precise spatial extent and point-by-point composition of the object [10].

The segmentation method described in this paper is mainly aimed at recognizing the location of nuclei in images. This is achieved by using the Canny edge detector followed by a series of morphological processes aimed at refining and closing detected edges and thereby producing recognizable objects (see Fig. 1).

The basic way to perform a closing operation (\bullet) is to do a binary morphological dilation (\oplus) followed by an erosion (\ominus),

$$A \bullet B = (A \oplus B) \ominus B, \quad (1)$$

where A represents the set and B the structuring element used [11]. If the structuring element B is disk shaped the same dilation can also be achieved by thresholding a distance transform generated using the binary object as a seed. However,




Fig. 1. The individual steps of the segmentation method when used to segment a nucleus (big roundish object in the center of the image): a) initial image, b) median filtering, c) Canny edge result, d) resulting distance map using the weighted measure, e) thresholded distance map, f) skeleton of threshold, g) final segmentation result achieved through filling of closed areas and opening, h) final segmentation overlaid on the original image

both methods are limited in that they dilate a binary object equally in all directions and everywhere in the image, without regard to the local image structure.

An alternative approach to curve closing is to perform an edge linking procedure in which matching pairs of endpoints are found and connected [12]. However, image content is still not taken into account. The two matching endpoints are connected using a straight line, making it possible to create boundaries that exclude part of the object or include background in the object segmentation.

Our method uses an image adaptive dilation, meaning that, from a given seed point, it propagates edges with regards to certain features in the image. This makes it possible to close edges in a way that better corresponds to the underlying image data.

2.1 Method Overview

Figure 1 shows the individual steps of the described method. The initial image (a) is first median filtered to remove noise (b). Canny edge detection is then applied (c) and the result is used as a seed image for a tensor weighted distance transform, implemented¹ from [5], that will be further described in the following section (d). The distance map is thresholded (e) and the subsequent binary image cleaned up via skeletonization (f), filling of closed areas and finally an isotropic opening (binary erosion followed by dilation) in order to get the final segmentation result (g, h).

2.2 Tensor Weighted Distances

We model the image as a chart of a 2D Riemannian manifold. The distance transforms we compute are based on estimates of geodesic distance in this curved Riemannian manifold. Since we only approximate geodesic distances, there is a possibility that the triangle inequality is not fulfilled and that the corresponding space we create is not a metric. For the task at hand, however, this possibility is of less practical importance.

The concept of looking at an image as a landscape, where the gray-level values correspond to height, is well known. How distances can be calculated in such a landscape is also a previously studied concept [13,14]. A 2D image $f(x, y)$ is in this case represented by the embedding of a surface in \mathbb{R}^3 :

$$\chi = (x, y, \alpha f(x, y)), \quad (2)$$

where α is positive. From Eq. 2 we derive the pullback metric from χ in the image plane $(x, y) \in \mathbb{R}^2$:

$$\mathbf{B} = I + \alpha^2 \nabla f \nabla f^T = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \alpha^2 \begin{bmatrix} f_x f_x & f_x f_y \\ f_y f_x & f_y f_y \end{bmatrix}. \quad (3)$$

¹ Toolbox Fast Marching by Gabriel Peyré

<http://www.mathworks.com/matlabcentral/fileexchange/6110>

The scalar product of two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$, in a particular point in the image, is then $\langle \mathbf{u}, \mathbf{v} \rangle_B = \mathbf{u}^T \mathbf{B} \mathbf{v}$. With this metric, the distance between two neighboring pixels, separated by a vector \mathbf{u} (Fig. 2(a)), is approximately

$$\|\mathbf{u}\| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle_B} = \sqrt{\mathbf{u} \mathbf{u} + \alpha^2 (\nabla f \cdot \mathbf{u})^2}. \quad (4)$$

The effect of the α parameter on the resulting distance transform is illustrated in Fig. 2(b). When $\alpha = 0$ the distance is reduced to the Euclidean distance in the image plane, while for larger values of α , steps taken in the direction of the gradient are penalized. Equation 3 is known as the Beltrami framework [15] and for cell images it means that distances increase faster perpendicular to the nucleus border. As pointed out in [16], Eq. 3 may be generalized by replacing the outer product of the gradient with the structure tensor [17,18,19]

$$\mathbf{S} = I + \alpha^2 \nabla f \nabla f^T * G_\sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \alpha^2 \begin{bmatrix} f_x f_x & f_x f_y \\ f_y f_x & f_y f_y \end{bmatrix} * G_\sigma, \quad (5)$$

where $*G_\sigma$ is the convolution by a Gaussian with σ standard deviation. The use of the structure tensor extends the Beltrami framework to penalize both image gradients (edges) and ridges (lines). It is a generalization because when $\sigma \rightarrow 0$ then $\mathbf{S} \rightarrow \mathbf{B}$. It is also different from the pure structure tensor because of the inclusion of the unit metric and the α parameter.




Fig. 2. a) Illustration showing distance vector from a pixel with gradient ∇f . b) Effect of the α parameter on the distance transform.

3 Data

Our algorithms were applied on two types of data: realistic synthetic cell images and minimum intensity projections of focus stacks taken of cervical smears with a bright-field microscope.

3.1 Realistic Synthetic Cell Images

For validation and parameter tuning purposes we have created a program that generates synthetic cell images. This program is a part of ongoing research. Here we briefly describe the image formation in the Pap smear simulator.

The images used for this paper have dimensions 1000×1000 pixels. The dynamic range is set to $[0,255]$ in order to match the real data which is being mimicked. An overview over the different steps of the image generation can be seen in Fig. 3.

The process starts with an empty image to which two levels of Perlin noise [20] is added to simulate general lighting inconsistencies and noise due to glass imperfections (Fig. 3a). The image area is then populated with cells.

Each generated cell is unique. The cytoplasm shape is created by selecting a field from a Voronoi diagram and then distorting the shape using Perlin noise. Texture generated from normally distributed noise and shadow created by adding a blurred layer are used to increase realism when drawing the finished shape in the image (Fig. 3b, c).

The nucleus shape, a randomly distorted circle, is created using a truncated Fourier series. Texturing of the nucleus includes two types of noise, a normally distributed noise to simulate chromatin texture and a Perlin noise to simulate staining imperfections (Fig. 3d, e).

Debris, e.g., blood cells, is generated as clusters containing three to ten small roundish objects. These are then added to the image as two layers with different levels of blurring (Fig. 3f, g).

The final step of creating the synthetic image is the addition of a small Gaussian blur ($\sigma = 0.8$) throughout the entire image. This ensures that all components are blended together in a way that simulates the point spread function of the microscope (Fig. 3h). The resulting images are hard to visually distinguish from real microscope images (Fig. 4).

3.2 Bright-Field Microscopy Images

The biological samples used in this study were prepared using the standard Papanicolaou protocol. We obtained focus stacks of the Pap smear samples using




Fig. 3. The key steps in the synthetic image creation: a) initial plane with noise, b) finished cytoplasm mask, c) cytoplasm subtracted from image, d) nucleus mask, e) nucleus subtracted from image, f) debris mask, g) debris subtracted from image, h) final image after blurring




Fig. 4. Synthetic image (left) with corresponding ground truth segmentation of nuclei (middle) compared to the minimum intensity projection of a focal stack taken of a Pap smear sample (right)

a MIRAX MIDI system (3DHISTECH Ltd., Budapest, Hungary) equipped with a $20\times / 0.8$ Plan-Apochromat objective. The images used have 1376×1032 pixels and a dynamic range in $[0, 255]$. For the application described in this paper, the focus stacks were reduced to a single image using minimum intensity value projection.

4 Evaluation

The focus of the evaluations performed in this paper is to obtain a high rate of nuclei recognition, rather than perfect delineation of every individual cell nuclei. For this reason we evaluate the segmentation result in terms of true positives (found nuclei) and false negatives (missed nuclei). An efficacy metric which compares the segmentation results with the ground truth for the synthetic images was implemented. The metric is based on theory described in [10], which compares two segmentation results of an object i , τ_{i_1} and τ_{i_2} , using the equation

$$r_i = \frac{|\tau_{i_1} \cap \tau_{i_2}|}{|\tau_{i_1} \cup \tau_{i_2}|}, r_I = \sum_{i \in I} r_i. \quad (6)$$

Here r_i represents the common part of the two segmentation results as a fraction of their combined area. If a ground truth is used as one of the segmentations in Eq. 6 then r_i represent the precision of the segmentation for each nucleus i . The global score r_I can be used to rate the performance over the entire image I .

We consider an object correctly segmented when r_i exceeds a threshold ρ . With the segmentation result given as a series of binary outcomes we are able to perform a statistical analysis on the results to compare two segmentation methods. We use a single sided McNemar test [21] that uses the values p_{TF} , the number of objects found by method 1 but not by method 2, and p_{FT} , the number of objects found by method 2 but not by method 1, to compare methods.

In the case of the synthetic data, the ground truth was known for all nuclei, making it possible to apply the efficacy metric described in Eq. 6 to evaluate segmentation results. The ratio threshold, ρ , was set to 0.7 for all evaluations. The results could then be statistically analyzed using McNemar's test.

Lacking a ground truth, the segmentation results for the Pap smear images were analyzed based on an expert's visual inspection. The McNemar test was then performed using the exact binomial distribution, instead of the χ^2 approximation, because of the small sample sizes in the real datasets.

5 Experiments

The experiments in this paper were focused at comparing the curve closing ability of the anisotropic dilation scheme to a standard isotropic one. The isotropic dilation method based on the distance transform, described in Section 2, was chosen for the comparison. It was implemented using the same framework as the anisotropic dilation only with the σ and α parameters set to 0. This means that dilation becomes fully dependent on the threshold, λ . The Canny edge parameters were heuristically selected prior to the evaluation process and then kept fixed throughout.

Algorithm development was done in the MATLABTM programming environment (The MathWorks Inc., Natick, MA) using the DIPimage toolbox version 2.0.1 [22].

5.1 Synthetic Data

A quantitative evaluation was performed using synthetic data sets generated as described in Section 3.1. A total of 160 images were used, each containing approximately 60 nuclei (not counting those touching the image border). The total number of fully segmentable nuclei was 9750.

Prior to evaluation the parameters of the dilation schemes were optimized based on the criterion described in Section 4 using a brute force approach. The parameter tuning was performed by running the segmentation algorithm on three synthetic images with a range of parameters. The dilation schemes were rated based the image score of Eq. 6. The optimal parameters were then chosen to be $\alpha = 8$, $\sigma = 6$ and $\lambda = 210$ for the anisotropic dilation and $\lambda = 6.5$ for isotropic dilation. In Fig. 5 maximum value projections illustrating the relationships between the three parameters σ , α and λ for the anisotropic parameter optimization are shown.

Using the optimized parameters all 160 images were then analyzed. The results are illustrated in Table 1(a) in the form of a contingency matrix. The results were analyzed using a single sided McNemar test, mentioned in Section 4: $H_0 : p_{TF} = p_{FT}$ and the alternative hypothesis $H_1 : p_{TF} > p_{FT}$. The test showed that the anisotropic method was significantly better, $p \ll 10^{-6}$.

5.2 Bright-Field Microscope Images

Because of the limited number of Pap smears available, the experiments using real data were focused towards a qualitative analysis. We used two Pap smear images, one where the cells were less densely packed and with a smaller amount




Fig. 5. A illustration of the evaluated parameter space for the anisotropic dilation scheme. The figure shows the maximum value projections along the λ (left), α (middle) and σ axis (right). The 'X' marks the global optimum.

Table 1. Tables showing the resulting contingency matrices for the experiments performed on the synthetic (a) and the real (b, c) datasets. The results corresponding to the two real images have been divided with respect to the type of cellular distribution. The matrices show the distribution of correct (cor.) and incorrect (in.) segmentations.

		(a)		(b)		(c)	
		Synthetic data		Real data (dense)		Real data (sparse)	
		Aniso.		Aniso.		Aniso.	
		cor.	in.	cor.	in.	cor.	in.
Iso.	cor.	6444	419	26	0	37	0
	in.	1487	1400	2	6	0	5




Fig. 6. The difference in curve closing ability between an anisotropic dilation and a standard isotropic dilation

of debris and one we consider to have a more normal composition. In total the two images contained 76 nuclei. No ground truth existed for the segmentation so parameter tuning was performed manually, starting from optimal values found for synthetic images (described in Section 5.1). We used the previously found optimal values for α and σ , but changed λ slightly.

We experienced that both dilation schemes found the same nuclei in the less cluttered image but that the anisotropic approach was able to segment two more nuclei in the cluttered image. This behavior is demonstrated in Fig. 6. The contingency matrices for the experiments are shown in Table 1(b) and 1(c). The results are encouraging but not significant due to the small sample size: The single side McNemar test did not reject the zero hypothesis, $p \leq 0.25$, for the results in the cluttered image.

6 Conclusions

We can conclude that Riemannian dilation performs better than isotropic dilation in a realistic biomedical setting where it is used for curve closing. Even though there is a need for a more complete validation on real data our extensive study on realistic simulated data shows the benefits of the anisotropic framework.

Through the use of our Pap smear simulator we were able to tune the parameters optimally by an exhaustive search in the relevant part of the parameter space. In the future, we expect to reduce the number of parameters in the segmentation by exploiting the almost linear dependency shown in Fig. 5 (right). When more Pap smear images become available in our project, we will have the option to auto-tune the algorithm by including a cell classification step, combined with a performance criteria, in a closed loop. We also expect the Pap smear simulator to be helpful in the future, to test image processing methods, and we plan to develop it further.

Acknowledgments

Ewert Bengtsson, Gunilla Borgfors, Cris Luengo, Robin Strand and Carolina Wählby are all acknowledged for their valuable input and support. The project "3D analysis of chromatin texture in cell nuclei to improve cervical cancer screening" is funded by the Swedish Research Council (2008-2738).

References

1. WHO: Comprehensive cervical cancer control: A guide to essential practice. WHO Press (2006)
2. Grohs, H., Husain, O. (eds.): Automated Cervical Cancer Screening. IGAKU-SHOIN Medical Publishers, Inc. (1994)
3. Canny, J.: A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence* 8, 679–698 (1986)
4. Do Carmo, M.: *Riemannian Geometry*. Birkhäuser, Boston (1992)
5. Prados, E., Lenglet, C., et al.: Control theory and fast marching techniques for brain connectivity mapping. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, New York, NY, vol. 1, pp. 1076–1083. IEEE, Los Alamitos (2006)

6. Jeong, W.K., Fletcher, P., et al.: Interactive visualization of volumetric white matter connectivity in dt-mri using a parallel-hardware hamilton-jacobi solver. *IEEE Transactions on Visualization and Computer Graphics* 13, 1480–1487 (2007)
7. Tsai, R., Zhao, H., Osher, S.: Fast sweeping algorithms for a class of hamilton-jacobi equations. *SIAM journal on numerical analysis* 41, 673–694 (2004)
8. Bouaynaya, N., Charif-Chefchaoui, M., Schonfeld, D.: Theoretical foundations of spatially-variant mathematical morphology part 1: Binary images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, 823–836 (2008)
9. Breuss, M., Burgeth, B., Weickert, J.: Anisotropic continuous-scale morphology. In: *Pattern Recognition and Image Analysis, Pt 2, Proc.*, Berlin, Germany, vol. 2, pp. 515–522 (2007)
10. Udupa, J., LeBlanc, V., et al.: A framework for evaluating image segmentation algorithms. *Computerized Medical Imaging and Graphics* 30, 75–87 (2006)
11. Gonzalez, R., Woods, E.: 9. In: *Digital Image Processing*, 3rd edn., pp. 633–635. Pearson Education, London (2008)
12. Ghita, O., Whelan, P.: Computational approach for edge linking. *Journal of Electronic Imaging* 11, 479–485 (2002)
13. Ikonen, L.: Pixel queue algorithm for geodesic distance transforms. In: *Andrès, É., Damiand, G., Lienhardt, P. (eds.) DGCI 2005. LNCS*, vol. 3429, pp. 228–239. Springer, Heidelberg (2005)
14. Verbeek, P., Verwer, B.: Shading from shape, the eikonal equation solved by grey-weighted distance transform. *Pattern Recognition Letters* 11, 681–690 (1990)
15. Sochen, N., Kimmel, R., Malladi, R.: A general framework for low level vision. *IEEE Transactions on Image Processing* 7, 310–318 (1998)
16. San Jose Estepar, R.: *Local Structure Tensor for Multidimensional Signal Processing. Applications to Medical Image Analysis*. PhD thesis, University of Valladolid, Spain (2005)
17. Knutsson, H.: A tensor representation of 3-D structures. Poster presentation of 5th IEEE-ASSP and EURASIP Workshop on Multidimensional Signal Processing, Noordwijkerhout, The Netherlands (1987)
18. Bigün, J., Granlund, G.H.: Optimal orientation detection of linear symmetry. In: *IEEE First International Conference on Computer Vision*, London, Great Britain, pp. 433–438 (1987)
19. Förstner, W., Gülch, E.: A fast operator for detection and precise location of distinct points, corners and centres of circular features. In: *Proc. ISPRS Intercommission Workshop*, pp. 281–304 (1987)
20. Perlin, K.: An image synthesizer. In: *SIGGRAPH 1985: Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, New York, USA, vol. 19, pp. 287–296 (1985)
21. Dietterich, T.: Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation* 10, 1895–1923 (1998)
22. Luengo Hendriks, C.L., van Vliet, L.J., Rieger, B., van Ginkel, M.: *DIPimage: a scientific image processing toolbox for MATLAB*. Computer Program (1999)