

Learning a Cost Function for Microscope Image Segmentation*

Sharmin Nilufar¹ and Theodore J. Perkins²

Abstract—Quantitative analysis of microscopy images is increasingly important in clinical researchers’ efforts to unravel the cellular and molecular determinants of disease, and for pathological analysis of tissue samples. Yet, manual segmentation and measurement of cells or other features in images remains the norm in many fields. We report on a new system that aims for robust and accurate semi-automated analysis of microscope images. A user interactively outlines one or more examples of a target object in a training image. We then learn a cost function for detecting more objects of the same type, either in the same or different images. The cost function is incorporated into an active contour model, which can efficiently determine optimal boundaries by dynamic programming. We validate our approach and compare it to some standard alternatives on three different types of microscopic images: light microscopy of blood cells, light microscopy of muscle tissue sections, and electron microscopy cross-sections of axons and their myelin sheaths.

I. INTRODUCTION

Quantitative microscope image analysis is a new frontier in the efforts to understand cellular and molecular function and disease [1], [2]. Naturally, manual segmentation of microscope images is possible. However, given the huge quantities of data produced by modern imaging systems, manual image interpretation and information extraction is not only time consuming, but also costly and potentially inaccurate. Automated or semi-automated processes are the only viable approach for accurate analysis of these huge datasets. Segmentation is the first step of many microscope image analyses, and has been applied to the morphological analysis of cells [3] classification and clustering of cellular shape [4], leukocyte detection and tracking [5], neurite and filopodia tracing [6], [7], subcellular analysis [8], and so on. Still, automated segmentation is challenging, because images and image quality can vary based on technician, platform, staining, cut and other factors. Moreover, “automated” techniques typically required careful parameter tuning to obtain satisfactory results.

An increasing trend has been to use machine learning to tune or improve the performance of a computer vision system [9]. For instance, learning can combine different image features to improve edge and boundary detection [10], [11]. Learning is also an integral to segmentation by graphical model approaches [12], [13].

We explore the possibility of learning a cost function in an active contour (or snake) formulation based on dynamic

programming. Active contour models are widely used for segmentation, including for microscope images. Cost, or energy, functions based on the signed-magnitude of image gradients along the boundary [14], [15] or the variance of gradient magnitude can be successful [16]. However, for any given application, one must commonly adjust parameters of such functions and/or to preprocess the image so that the desired boundaries are detected. These ad hoc adjustments take time and negate some of the benefit one desires from an automated approach.

We have developed a learning-based system that a user trains by tracing the boundaries of just one or a few example objects. These are used to train a probabilistic classifier for recognizing object boundaries, which we then translate into a cost function. The user then interactively clicks on additional objects, and the system uses the learned cost function to identify its boundary automatically. The goal is to produce accurate segmentations of user-selected objects with minimal effort—while at the same time keeping the human in the loop for “sanity checking” and quality assurance. We find the approach works robustly and with a good accuracy on a variety of types of microscope images.

II. MATERIALS AND METHODS

A. Dynamic Programming for Segmentation

We use a dynamic programming formulation for segmentation, which is a variation on the approach we proposed in Nilufar & Perkins [17], and similar to that used in many previous studies [18], [16]. To segment a single object in image I , the formulation assumes we know a sourcepoint S within the object in the image—in our case, given by a user click. The software then seeks the best object boundary around the sourcepoint. The dynamic program is based on a radial mesh of grid points (Fig. 1), with R total radial lines and N points along each line. A boundary, or contour, is tgy equivalent to a vector $C \in \{1, \dots, N\}^R$, specifying the position of the contour along each radial line.

The formulation also assumes that we are given a cost function $J : \{1, \dots, R\} \times \{1, \dots, N\} \rightarrow \mathfrak{R}^+$, which specifies a non-negative cost for sending the contour through each possible point along each radial line. Intuitively, this cost represents how “boundary like” each point is, with higher costs meaning less boundary-like. However, it would be too simplistic to construct a boundary merely by choosing the least-cost point along each radial line. Such a contour could be highly non-smooth. Thus, we impose a constraint that the contour positions at adjacent radial lines are within δ positions: for a contour C , $|C_r - C_{r+1}| \leq \delta$ for all $r = 1, \dots, N - 1$, and also $|C_N - C_1| \leq \delta$.

*This work was not supported by any organization

¹S. Nilufar is with Ottawa Hospital Research Institute, Ottawa, Ontario, Canada K1H8L6 snilufar@gmail.com

²T. Perkins is with the University of Ottawa and Ottawa Hospital Research Institute, Ottawa, Ontario, Canada K1H8L6 tperkins@ohri.ca, Web: www.perkinslab.ca




Fig. 1: Example of a dynamic programming mesh centered on an object, and an optimized boundary.

The best object boundary is the one with minimal sum of costs, subject to the δ -smoothness constraint. To find this boundary, we use standard finite-horizon dynamic programming [18] to find the least-cost path from each possible start position $C_1 \in \{1, \dots, N\}$ on the first radial line, through each successive radial line $r = 2, \dots, N - 1$, and finally to each possible end position $C_N \in \{1, \dots, N\}$ on the N^{th} radial line. Naturally, each step of these paths is constrained to be within δ positions radially of the previous one. Finally, we find the best path by searching for the combination of C_1 and C_N that minimizes the sum of costs, subject to $|C_N - C_1| \leq \delta$.

For all examples in this paper, we used a grid with $R = 250$ radial lines and $N = 95$ positions along each radial line, starting at 3 pixels out from the sourcepoint, and separated by 2 pixels. Based on preliminary testing, we chose a smoothness parameter of $\delta = 2$; a value of $\delta = 1$ did not allow contours to change quickly enough to follow object boundaries, whereas larger values resulted in jumpy boundaries. In principle, δ might be learnt from sample contours, as we will describe in the next section for the cost function.

B. Learning the Cost Function

For a particular image segmentation problem, we use a probabilistic classification approach to learn a cost function based on one or more user-specified contours in one or more sample images. For each user-provided contour, we compute the centroid and take it to be the sourcepoint of that object. Based on that sourcepoint, we create a radial mesh, just as described in the previous section. Along each radial line, we identify the position closest to the user-provided contour. These positions are considered to be positive examples in a binary classification task, and all other positions along the radial lines are considered to be negative examples.

To help us discriminate the positive examples from the negative, we need a set of features. In all our examples,

we use a common set of features containing information on changes in intensity, color and texture at each position along each radial line. Specifically, we use the (numerically estimated) derivatives along the radial line of the following six properties: (1) grayscale intensity, (2–4) hue, saturation and value, as defined in the HSV colorspace, (5) local image entropy (as computed by MATLAB), and (6) the output of a multiscale ridge detector [19].

Having specified the features at each position along each radial line, and having specified that class of each position (negative or positive), we have a fully-defined binary classification problem. We employ a standard logistic classifier, with weights fit by the `mnrfit` function in MATLAB. The resulting classifier ϕ outputs the probability of the positive class, given a vector of feature values. We turn this into a cost function for the dynamic program by taking the negative logarithm, $J = -\log \phi$. As such, a minimum-cost contour is equivalent to the contour with maximum joint probabilities of the positive class, subject to the smoothness constraints.

C. Datasets

We applied the proposed method on two types of light microscopy images and one electron microscopy image.

Muscle fiber image: This is an image from Mouse Tibialis Anterior (TA) muscle. The TA muscle received a cardiotoxin (CTX) injection followed by siRNA/Lipofectamine RNAiMax injections 6h, 48h and 96h later. Seven days after the CTX injection, mice were sacrificed, TA muscles were isolated, fixed in paraformaldehyde and processed in paraffin blocks. Four micron thick sections were stained with hematoxylin-eosin. The image was acquired with a Carl Zeiss Axio Imager. Figure 2(a) shows a crop from the original, larger image, which we use for testing. The goal is to determine the cross sectional area distribution of all myofibers in the experimental sample; the larger goal of the project is to study muscle regeneration and degenerative diseases.

Blood cell image: This is an image of cultured human red blood cells that were differentiated ex vivo from adult hematopoietic stem/progenitors cells as previously described [20]. Cells were concentrated by Cytospin, fixed in methanol for 2 min. and stained with May-Grünwald Giemsa. Figure 2(e) shows a crop from a larger image which we use for testing. The larger goal of the project is to perturbations in blood cell differentiation in leukemia.

Nerve image: This is an image of the cross section of mouse tibial nerves. Mice were anaesthetized with Avertin and perfused transcardially with Karnovsky's fixative (4%PFA, 2%glutaraldehyde, 0.1 M sodiumcacodylate, pH7.4). The optic nerve was then removed and postfixed in Karnovsky's fixative at 4C. Fixed optic nerves were cut into ultrathin sections, stained with uranyl acetate and lead citrate, and analyzed by electron microscopy. The goal of this staining is to calculate the axon and myelin diameters. Number of myelinated fibers relative to total fibers can then be determined and subdivided into groups by axon diameter. The G ratio can be calculated by dividing the axon diameter




Fig. 2: Panels (a), (e) and (i) show original image of muscle, blood and nerve, (b), (f) and (j) show detected contours with proposed snake (c), (g) and (k) detected contours with GICOV, (d), (h) and (l) show detected contours with MaxGrad method respectively.

by the axon plus myelin diameter. The image is shown in Figure 2(i). The larger goal of the project is to understand the demyelinating diseases.

III. RESULTS

We applied our method to all three images, providing a single user-traced training contour for each one. For the muscle image, we sought to detect the muscle fibers (purple objects), and for the blood images we sought to detect the cell boundaries (rather than the pink nuclear boundaries, which are visually stronger). For the nerve image we learned two cost functions, one for detecting the inner whitish axon, and one for detecting the dark outer myelin sheath. We then clicked on each object in the test images, and used dynamic programming with our learned cost functions to identify the object boundaries. Figure 2(b,f,j) shows the results. Close inspection shows some small segmentation errors, but most objects are correctly segmented.

We compared our results to two other popular dynamic programming frameworks namely the MaxGrad snake (which maximizes the sum of intensity-image gradient magnitudes) and the GICOV snake (which minimizes the variance of the gradient magnitude along the path) [16]. Results are shown in Figure 2(c,d,e,h,k,l). For the muscle fiber task, the MaxGrad and GICOV snake boundaries more often overlap, due in part to influence of the dark blue/black nuclei. In the blood cell images, the MaxGrad snake is, unsurprisingly, attracted by the stronger nuclear boundary in many cases. The GICOV snake performs similarly. In the nerve image, MaxGrad and GICOV snakes can be set to detect inner and outer boundaries by specifying the desired gradient direction (light-to-dark or dark-to-light). They achieve better success than in the blood images, though the boundaries are prone to skipping to a neighboring axon.

As a more formal evaluation, we manually segmented each object in each image, and we compared the three approaches quantitatively using Dice coefficients [21]. For a given object corresponding to a set of pixels T (for True pixels, based on our manually-performed segmentation) and with

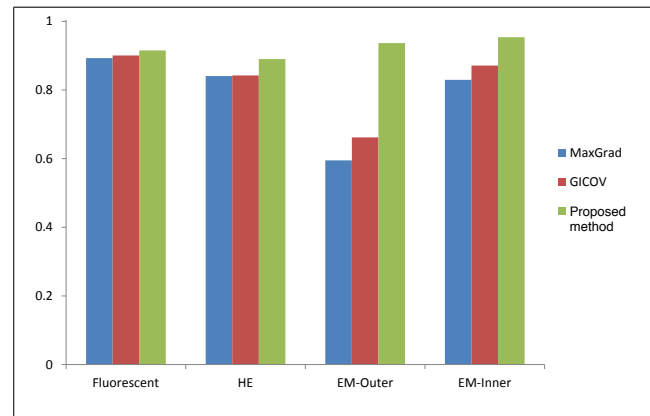


Fig. 3: Dice coefficients for the proposed supervised snake, GICOV snake and MaxGrad snake on three different types of microscope images.

estimated segmentation E (from one of the three automatic approaches), the Dice coefficient is $2 \times |T \cap E| / (|T| + |E|)$. The Dice coefficient is a maximum of one, when T and E are in perfect agreement, and as low as zero if they have no overlap at all. For each image and for each automatic segmentation approach, we computed the average Dice coefficient across all objects. The results show that our learning-based approach had superior performance on all four segmentation tasks (Figure 3).

IV. CONCLUSIONS

We have proposed a general procedure by which just one or a few training contours, combined with standard probabilistic classification, can learn to automatically identify the boundaries of desired objects in microscopy images. The identical approach, with the same dynamic programming mesh, features, and logistic classifier, was successful on all three types of images: light microscopy images of blood cells and of a muscle tissue section, as well as electron microscopy. Our method solves some practical difficulties with standard snake models—it identifies a good cost function and avoids hand tuning snake parameters. Our choice of the dynamic programming framework also avoids difficulty with finding the optimal contour, which are present in parametric-curve and level-set formulations of snakes. In these respects, our approach has similar advantages to GraphCut [13] and its many variations—efficient optimization and learning of a cost function. We tried a version of GraphCut on our test images and performance was poor. However, the software we obtained uses different features and training information; we are working on a more fair comparison. We are currently conducting a more thorough empirical evaluation of our system, on a larger set of microscope images, looking at the importance of the number and quality of training contours, and the effects of different features, classifier, and smoothness and grid parameters. A drawback to our approach is that the dynamic programming mesh cannot capture objects strong concavities. We are working on a more powerful dynamic programming approach that would overcome these difficulties. Finally, we have assumed that a sourcepoint for each object is given. An alternative would be to learn to detect sourcepoints. How to combine this with cost function learning remains to be seen.

ACKNOWLEDGMENT

For providing us with sample microscope images as well as expertise in assessing segmentations, we thank our collaborators: Dr. Alex Blais (muscle sections), Drs. Marjorie Brand and Carmen G. Paliı (blood cells), and Drs. Rashmi Kothary and Yves De Repentigny (nerve sections). This work was supported in part by a Government of Ontario Ministry of Research and Innovation (MRI) grant, and by a grant from the National Sciences and Engineering Research Council of Canada.

REFERENCES

[1] Hanchuan Peng. Bioimage informatics: a new area of engineering biology. *Bioinformatics*, 24(17):1827–1836, 2008.

[2] Jason R Swedlow, Ilya G Goldberg, Kevin W Eliceiri, OME Consortium, et al. Bioimage informatics for experimental biology. *Annual review of biophysics*, 38:327, 2009.

[3] Yoo-Jin Kim, Thomas Brox, Wolfgang Feiden, and Joachim Weickert. Fully automated segmentation and morphometrical analysis of muscle fiber images. *Cytometry Part A*, 71(1):8–15, 2007.

[4] Yung-Fu Chen, Po-Chi Huang, Ker-Cheng Lin, Hsuan-Hung Lin, Li-En Wang, Chung-Chuan Cheng, Tsung-Po Chen, Yung-Kuan Chan, and John Y Chiang. Semi-automatic segmentation and classification of pap smear cells. *IEEE J Biomed Health Inform*, 18(1):94–108, 2014.

[5] Gang Dong, Nilanjan Ray, and Scott T. Acton. Intravital leukocyte detection using the gradient inverse coefficient of variation. *IEEE Trans. Med. Imaging*, 24(7):910–924, 2005.

[6] Khalid A. Al-kofahi, Sharie Lasek, Donald H. Szarowski, Christopher J. Pace, George Nagy, Senior Member, James N. Turner, and Badrinath Roysam. Rapid automated three-dimensional tracing of neurons from confocal image stacks. *IEEE Transactions on Information Technology in Biomedicine*, 6:171–187, 2002.

[7] Sharmin Nilufar, Anne Morrow, Jonathan Lee, and Theodore J. Perkins. Filodetect: automatic detection of filopodia from fluorescence microscopy images. *BMC Systems Biology*, 7:66, 2013.

[8] Yanhua Hu, Elvira Osuna-Highley, Juchang Hua, Theodore Scott Nowicki, Robert Stolz, Camille McKayle, and Robert F Murphy. Automated analysis of protein subcellular location in time series images. *Bioinformatics*, 26(13):1630–1636, 2010.

[9] Viren Jain, H Sebastian Seung, and Srinivas C Turaga. Machines that learn to segment images: a crucial technology for connectomics. *Current opinion in neurobiology*, 20(5):653–666, 2010.

[10] Piotr Dollar, Zhuowen Tu, and Serge Belongie. Supervised learning of edges and object boundaries. In *CVPR*, volume 2, pages 1964–1971. IEEE, 2006.

[11] D.R. Martin, C.C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(5):530–549, 2004.

[12] Stan Z Li. *Markov random field modeling in computer vision*. Springer-Verlag New York, Inc., 1995.

[13] Yuri Y Boykov and M-P Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *Proceedings of the Eighth IEEE International Conference on Computer Vision, 2001*, volume 1, pages 105–112. IEEE, 2001.

[14] Chenyang Xu and Jerry L Prince. Snakes, shapes, and gradient vector flow. *Image Processing, IEEE Transactions on*, 7(3):359–369, 1998.

[15] HyunWook Park, Todd Schoepflin, and Yongmin Kim. Active contour model with gradient directional information: directional snake. *Circuits and Systems for Video Technology, IEEE Transactions on*, 11(2):252–256, 2001.

[16] Nilanjan Ray, Scott T. Acton, and Hong Zhang. Seeing through clutter: Snake computation with dynamic programming for particle segmentation. In *ICPR*, pages 801–804. IEEE, 2012.

[17] Sharmin Nilufar and Theodore J. Perkins. Learning to detect contours with dynamic programming snakes. In *ICPR 2014*, to appear.

[18] A.A. Amini, T.E. Weymouth, and R.C. Jain. Using dynamic programming for solving variational problems in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(9):855–867, 1990.

[19] Tony Lindeberg. Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, pages 224–270, 1994.

[20] Carmen G Paliı, Carolina Perez-Iratxeta, Zizhen Yao, Yi Cao, Fengtao Dai, Jerry Davison, Harold Atkins, David Allan, F Jeffrey Dilworth, Robert Gentleman, et al. Differential genomic targeting of the transcription factor tall1 in alternate haematopoietic lineages. *The EMBO journal*, 30(3):494–509, 2011.

[21] L. R. Dice. Measures of the Amount of Ecologic Association Between Species. *Ecology*, 26(3):297–302, July 1945.