

ADAPTIVE ACTIVE-MASK IMAGE SEGMENTATION FOR QUANTITATIVE CHARACTERIZATION OF MITOCHONDRIAL MORPHOLOGY

Kuan-Chieh Jackie Chen¹, Yiyi Yu¹, Ruiqin Li¹, Hao-Chih Lee¹, Ge Yang¹ and Jelena Kovačević^{1,2}

¹Dept. of Biomedical Eng. and Center for Bioimage Informatics, ²Dept. of Electrical & Computer Eng. Carnegie Mellon University, Pittsburgh, PA 15213, USA

ABSTRACT

We propose an automated algorithm for segmentation of mitochondria from widefield fluorescence microscopy images for quantitative morphology characterization. Mitochondria are membrane-bound organelles that are essential to cells of higher living organisms. Reliable and precise quantitative characterization of their shape is crucial to understanding related physiology and disease mechanisms. Building upon the active-mask framework developed for segmentation of confocal fluorescence microscope images, we propose a new adaptive region-based distributing function to effectively address the problem of halo artifacts that are common in widefield fluorescence images. Such artifacts prevent the segmentation of weak features of mitochondria using existing algorithms. We compare the algorithm to the original active-mask algorithm as well as the geodesic active contour algorithm based on hand-segmented ground truth, and find that it performs significantly better both qualitatively and quantitatively.

Index Terms— segmentation, active masks, mitochondria

1. INTRODUCTION

Mitochondria are membrane-bound organelles that are essential to eukaryotic cells of higher living organisms including humans. They serve critical functions such as energy conversion and distribution as well as intracellular signaling in many basic cellular processes [1]. To meet the changing needs of dynamic cellular processes, mitochondria frequently undergo controlled shape changes [2]. Reliable and precise quantitative characterization of mitochondrial shape is crucial to understanding related physiology and disease processes [3]. This, in turn, requires reliable and precise image segmentation to extract shapes of mitochondria from their microscopy images. The main goal of this study is to segment fluorescently tagged mitochondria, as depicted in Figure 1, from motor neurons of *Drosophila* 3rd instar larvae.

Segmentation of fluorescence microscopy images has attracted substantial attention in the past few years as it is a crucial first step in quantitative image data analysis. In particular, fluorescence images often have a *punctate* (dotted, unconnected) appearance, leading to the lack of discernible edges and precluding the use of edge-based algorithms. Two algorithms have been widely used: seeded watershed, designed as a general-purpose segmenter, and highly dependent on the choice of initial seeds (for example, [4]), and Voronoi segmentation (for example, [5]). Several attempts have also been



Fig. 1. GFP-labeled mitochondria (arrows), imaged in axons within segmental nerves of a dissected *Drosophila* 3rd instar larva.

made to use active-contour algorithms (for example, [6]), considered state-of-the-art in medical image segmentation. Active-contour methods offer flexibility, allowing for the choice and/or design of forces suited to a given application. In the segmentation work in our lab, we added a multiresolution version to active contours, achieving an order-of-magnitude increase in speed, as well as a multiscale version to smooth punctate patterns and *connect the dots* (that is, creating edges) [7]. These led to the *active-mask framework* [8] that is the basis for the work in this paper.

In widefield fluorescence microscopy images, an artifact known as *halo* (dim ring around a bright features appearing as a result of the acquisition method) is often present. As its intensity is in the range of the weak features of the mitochondria, it is a challenge to exclude it during segmentation. In this paper, we present an algorithm for automated segmentation of fluorescently-tagged mitochondria, which expands the active-mask framework to incorporate local statistical properties of the image and solves the issue of confusing halos with weak mitochondria features.

2. ACTIVE-MASK SEGMENTATION

Our previous efforts raised several questions: What is a contour in a digital image? Since updating the level-set function in active-contour algorithms can be very slow, how do we reconstruct it in its multiresolution version? These questions led us to formulate a set of requirements for a new algorithm that combines the flexibility offered by active-contour methods, speed offered by multiresolution methods, smoothing offered by multiscale methods, and statistical modeling offered by region-growing methods.

Framework. We termed this algorithm *active-mask segmentation* [8], as we moved from the concept of a *contour* (ill-defined for digital images) to that of a *mask* (well-defined). To preserve topology (for example, separate touching image features), multiple masks are introduced. Analogously to forces in active-contour methods, we defined *distributing functions*, acting on each mask as a whole. We initially designed two types of distributing functions: the *region-growing* ones, which depend on the statistical properties on the image, and the *voting-based* ones, which depend on the geometrical properties of the image. The region-based functions aim to separate statistically different regions (background/foreground), while the voting-based ones aim to separate foreground into multiple masks (for example, features such as mitochondria). We embed

We gratefully acknowledge support from the NIH through award EB009875, NSF through awards DBI-1052925, 1017278 and the CMU CIT Infrastructure Award, as well as discussions with Drs. Srinivasa and Fickus.

these into a multiscale block allowing the algorithm to find boundaries by going through multiple scales. A multiresolution block serves to speed up the computation by segmenting quickly a coarse version, then lifting up one resolution at a time to fine tune the result. The entire algorithm is illustrated in Figure 2.

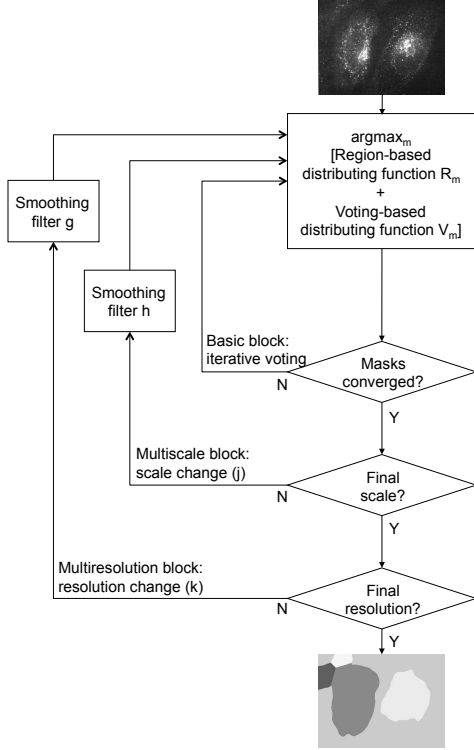


Fig. 2. Active-mask algorithm [8].

Given is a d -dimensional image f with $(M - 1)$ cells present. Similarly to [9], a collection of M masks will be a function ψ that assigns each pixel n a value $\psi_n \in \{1, \dots, M\}$, where n belongs to the m th mask if $\psi_n = m$. Each mask is a characteristic function of the appropriate region, $\mu_{m,n} = 1$ if $\psi_n = m$, and 0 otherwise.

Region-Based Distributing Functions. Let the first mask represent the background. The role of the region-based distributing function R_1 is to separate background from foreground by forcing those points with a low average intensity into the background mask μ_1 (we assume $R_i = 0$ for $i = 2, \dots, M - 1$),

$$R_{1,n} = \alpha \text{sig}\left(\beta((f \star h)_n - \gamma)\right), \quad (1)$$

where h is a *region-based lowpass filter*, in [8] chosen to be

$$h_n = e^{-|n|^2/a^2}, \quad (2)$$

and $a > 0$ is the *scale of the region-based lowpass filter*. The function sig is any sigmoid-type function, for example,

$$\text{sig}(x) = \text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2/2} dt,$$

so that it is ± 1 at $\pm\infty$, respectively. Function $f \star h$ denotes convolution adjusted for boundary issues, the *skewing factor* $\alpha \in (-1, 0)$ should be close to -1 , β determines the *harshness of the threshold*;

and *average border intensity* γ should be taken as the average intensity of those pixels which lie on the boundary between being inside every cell and being outside them all. Thus, for a pixel outside a cell, $(f \star h)_n < \gamma$, or $R_{1,n} \approx -\alpha \approx 1$. During voting described shortly, this will skew it so that the background is chosen.

Voting-Based Distributing Functions. These functions basically perform majority voting based on statistical properties of local regions. This is done by computing a convolution (adjusted for boundary issues) with a *voting-based lowpass filter* g , for example,

$$g_n = 1 - \text{erf}(|n|^2/b^2 - 1),$$

where b is the *scale of the voting-based lowpass filter*. At iteration step i , the *voting-based distributing function* $V_{m,i}$ is

$$V_{m,i,n} = (\mu_m \star g)_n.$$

Basically, $V_{m,i,n}$ tells us how much the pixel n wants to move to mask m , given its local neighborhood.

Basic Block. Combining these functions, the basic iteration is

$$\psi_{i+1,n} = \underset{m=1,\dots,M}{\text{argmax}} [V_{m,i,n} + R_{m,n}].$$

Thus, if $\alpha \approx -1$ and a pixel is inside a cell, R_1 will prejudice the voting so that any mask but background one is chosen. Similarly, if a pixel is outside a cell, R_1 will sway the voting towards the background. The algorithm can be initialized randomly. The first few iterations will typically produce a rough separation of background/foreground, followed by a fine separation of individual masks.

Multiscale Block. Changing the scale a of the region-based lowpass filter h in (2) can help overcome the problem of splits. Choosing a large a allows the algorithm to converge quickly, though the resulting masks may be larger than the cells they are segmenting. However, such a result can be used as a starting point for a new iteration with a smaller a .

Multiresolution Block. To speed up the algorithm we first segment the coarse version and then gradually lift up the resolution and segment again. The lifting is done by upsampling and filtering, for example using a simple Haar filter [8].

3. ADAPTIVE ACTIVE-MASK SEGMENTATION

Given the flexibility offered by the active-mask segmentation, we extend the framework to capture the mitochondrial morphology from the fluorescence image data as shown in Figure 1. Unlike the confocal microscopy images in [8], the images of mitochondria were obtained from widefield epifluorescence microscopy. The segmentation challenge posed by these images is that the level of fluorescence intensity varies significantly among mitochondria. Consequently, the foreground signal is far from uniform. The signal-to-noise ratios of certain mitochondria are so low that it is difficult to visually determine their boundaries (Figure 3). More importantly, widefield fluorescence images also contain halos as we explained earlier (see Figure 3 for an example). Whereas the halos are generally dimmer than the area inside the mitochondrion, some halos are as bright as or even brighter than the dimmer mitochondria. This effect could be attributed to a combination of out-of-focus fluorescence and scattered background fluorescence. When we attempted to segment the image using the active-mask algorithm (and several other segmentation algorithms such as global thresholding, active contours), this became a problem. If we determined a value of γ in (1) to exclude the halo,



Fig. 3. (a) Two mitochondria (arrows), a very bright one on the left and a very dim one on the right, from Figure 1 (contrast-enhanced for visibility). (b) The hand-segmented mask of the brighter mitochondrion (black area) is overlaid on the same image. A halo surrounding the brighter mitochondrion has intensity values similar to those of the dimmer one.

we would be excluding the dimmer object as well. Conversely, if we captured the dimmer objects, we would be capturing the halo as well. The solution is to extend the active-mask framework to precisely segment the mitochondria while excluding the halos. While there exist studies addressing similar problems in different imaging modalities [10, 11], we prefer a simple and computationally efficient version of the region-based distributing function.

Adaptive Region-Based Distributing Function. Inspired by [12], we propose a new distributing function R_1 ; we call it *adaptive region-based distributing function*,

$$R_{1,n} = \alpha \text{sig}\left(\beta((f \star h)_n - ((f \star p)_n + \gamma_a))\right), \quad (3)$$

for some region-based lowpass filter h and *adaptive threshold lowpass filter* p . We choose h as in (2) and p as

$$p_n = e^{-|n|^2/c^2},$$

where $c > 0$ is called the *scale of the adaptive threshold lowpass filter*. In the above, $f \star p$ denotes convolution designed for local weighted average intensity and *threshold adjustment* γ_a indicates the adjustment of adaptive threshold from local weighted average which is determined as $\gamma_a = (\mu_f - \mu_b)/2$, where μ_f is the average intensity of the dim object, and μ_b is that of the background. The idea here is to compare the intensity value of a given pixel with the weighted average of the intensities in some neighborhood (a disk area with the radius of three times the scale c). For such an R_1 , for a typical pixel n inside a mitochondrion, $(f \star h)_n > ((f \star p)_n + \gamma_a)$, and so $R_{1,n} \approx \alpha \approx -1$. During the voting described earlier, $R_{1,n}$ then skews the voting so that any mask but the background one is chosen. On the other hand, for a typical pixel n in the area of the halo, the local weighted average intensity is raised depending on the distance between the pixel and the closest mitochondrion. Therefore, similarly to those pixels in the background, $(f \star h)_n < ((f \star p)_n + \gamma_a)$, and so $R_{1,n} \approx -\alpha \approx 1$ (see Figure 4). During the voting, $R_{1,n}$ then skews the voting so that the background is chosen.

4. RESULTS AND DISCUSSION

Data Set. Mitochondria, fluorescently tagged with GFP, from motor neurons of dissected *Drosophila* 3rd instar larvae were imaged. The data set we used in this study consists of 45 images selected from 15 movie sequences. Each movie sequence consists of 181 frames (varied sizes, about 1300×300) containing 10–25 mitochondria each. Three frames (#20, #100, and #180) from the beginning, middle, and end of each sequence were used for testing so that potential impact of photobleaching can be examined. The effective pixel size is 64.5 nm. To generate ground truth, we used images hand-segmented by outlining splines enclosing each mitochondrion.

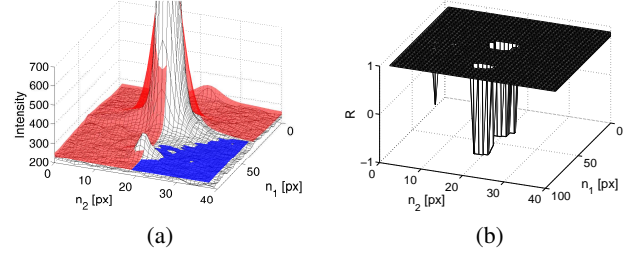


Fig. 4. (a) Illustration of the effect of the adaptive threshold plane $((f \star p)_n + \gamma_a)$ (red) versus the global threshold plane γ (blue) overlaid on the smoothed image $(f \star h)_n$ in 3D. (b) The resulting soft-thresholded image in 3D as in (3).

Experimental Setup. We compare the performance of the proposed algorithm to our original active-mask algorithm and a geodesic active contour (GAC) algorithm from [13]. GAC begins with an initial level set, given by the region of the interest, and evolves the contour according to the speed image, created by a sigmoid function of the soft-threshold image, $S = \tanh((f \star h) - \gamma_\ell)$, where $f \star h$ is defined previously and γ_ℓ is determined by the statistical properties of the images. The output level set generated by the GAC is then binary thresholded to produce the segmentation.

All algorithms are implemented in MATLAB except the GAC level-set evolving function, for which we use ITK through MATITK [14]. The runtime for the adaptive active-mask algorithm remains roughly the same as the original active-mask algorithm, which is about 2 minutes for the finest resolution, while the one for the GAC level-set algorithm is about 30 seconds. All computations were performed on a desktop workstation (2x Intel Xeon E5503 2.00 GHz and 16GB RAM).

We used similar parameters as in [8] for both original active-mask and adaptive active-mask algorithms with the initial number of masks $M = 256$, initial resolution level $K = 3$, final resolution level $K_0 = 0$, scale parameters $a = 3.5$, $b = 8$, $c = 5$, and skewing factor $\alpha = -0.9$. For the original active-mask algorithm, as in [8] we determined β and γ based on one image from each movie sequence and used the same numbers for all the other images from that movie, and similarly for γ_a for the adaptive active-mask algorithm. For GAC, the following parameters were used: $\omega_{advection} = 1$, $\omega_{curvature} = 1$, and $\omega_{propagation} = 10$ (see [15] for more details).

Qualitative Performance Evaluation. We compare the results of the adaptive active-mask, the original active-mask, and the level-set algorithms against the hand-segmented masks, shown in Figure 5. As previously noted, the original active-mask algorithm performs poorly on the image which contains both dim objects and halos. In Figure 5(c) and (d), we can see that even with different levels of threshold-based tuning, it cannot produce precise segmentations. It either loses dim objects or includes halos to create a larger area than that marked by hand segmentation for brighter objects. The GAC level-set algorithm also shares similar problems. On the other hand, the adaptive active-mask algorithm produces segmentations matching the ground truth better than the other two algorithms.

Quantitative Performance Evaluation. We used two standard performance measures: the recall (R)/precision (P) and area similarity (AS) [16] by comparing the algorithm-produced segmentations (SG) to the ground truth (GT) for each mitochondrion in each image. The recall/precision rates are calculated as follows:

$$R = \frac{|TP|}{|GT|}, \quad P = \frac{|TP|}{|SG|},$$

where a true positive (TP) refers to a mitochondrion correctly segmented; a false positive (FP) to background area false segmented; and a false negative (FN) to a mitochondrion not segmented.

The AS normalizes twice the area of the intersection of the SG and the GT by their total area. This measure will penalize the algorithm that produces larger or smaller regions. According to [16], performance of the AS $\geq 70\%$ generally implies a good agreement of the algorithm's result with the ground truth.

Results. A summary of the results based on these methods is given in Table 1. All measures have been averaged over all mitochondria and all images. In terms of recall and precision, which roughly measure the percentage of usable segmentation, the adaptive active-mask algorithm outperforms the others by a fair margin. Note that recall and precision only present coarse results, as they do not take into account the fit of the segmented area to that marked by hand segmentation. From the AS results for each algorithm, as expected, the adaptive active-mask algorithm maintains good agreements with the ground truth (AS = 70.54%) with higher recall and precision rates. This is because the adaptive threshold enables the region-based function to exclude halos without losing the dimmer objects, while the other two algorithm can only trade off one for the other.

5. CONCLUSIONS

We present an automated algorithm for segmentation of mitochondria in widefield fluorescence images. The algorithm extends our previous work on active-mask segmentation, originally developed for confocal fluorescence microscope images. It adaptively extends active-mask segmentation by introducing a region-based distributing function that allows for dim artifacts such as halos to be discarded while still preserving dim objects of interest. The algorithm performs well when benchmarked using hand-segmented ground truth and significantly outperforms the original active-mask algorithm and the GAC algorithm both qualitatively and quantitatively.

6. REFERENCES

- [1] H. M. McBride, M. Neuspiel, and S. Wasiak, "Mitochondria: more than just a powerhouse," *Current Biology*, vol. 16, pp. R551–R560, 2006.
- [2] R. E. Jensen, "Control of mitochondrial shape," *Current Opinion in Cell Biology*, vol. 17, pp. 384–388, 2005.
- [3] D. C. Chan, "Mitochondria: dynamic organelles in disease, aging, and development," *Cell*, vol. 125, pp. 1241–1252, 2006.
- [4] C. Wählby, *Algorithms for Applied Digital Image Cytometry*, Ph.D. thesis, Uppsala Univ., 2003.
- [5] T. R. Jones, A. E. Carpenter, and P. Golland, "Voronoi-based segmentation of cells on image manifolds," *Lecture Notes in Computer Science*, pp. 535–543, 2005.
- [6] A. Sarti, C. O. Solorzano, S. Lockett, and R. Malladi, "A geometric model for 3-D confocal image analysis," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 12, pp. 1600–1609, 2000.
- [7] G. Srinivasa, M. C. Fickus, and J. Kovačević, "Multiscale active contour transformations for the segmentation of fluorescence microscope images," in *Proc. SPIE Conf. Wavelet Appl.*

in Signal and Image Proc., San Diego, CA, Aug. 2007, vol. 6701:18, pp. 1–15.

- [8] G. Srinivasa, M. C. Fickus, Y. Guo, A. D. Linstedt, and J. Kovačević, "Active mask segmentation of fluorescence microscope images," *IEEE Trans. Image Proc.*, vol. 18, no. 8, pp. 1817–1829, Aug. 2009.
- [9] B. Merriman, J. K. Bence, and S. J. Osher, "Motion of multiple junctions: A level set approach," *Journ. Comp. Phys.*, vol. 112, pp. 334–363, 1994.
- [10] L. Bradbury and J. W. L. Wan, "A spectral k-means approach to bright-field cell image segmentation," in *Proc. IEEE Int. Conf. EMBS Soc.*, Sept. 2010, pp. 4748–4751.
- [11] Z. Yin, K. Li, T. Kanade, and M. Chen, "Understanding the optics to aid microscopy image segmentation," in *Proc. Int. Conf. Med. Image Computing and Computer-Assisted Intervention*, Berlin, Germany, 2010, pp. 209–217, Springer-Verlag.
- [12] P. D. Wellner, "Adaptive thresholding for the DigitalDesk," Tech. Rep., Rank Xerox Research Centre, 1993.
- [13] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. Journ. Comp. Vis.*, vol. 22, pp. 61–79, Feb. 1997.
- [14] V. Chu and G. Hamarneh, "MATLAB-ITK interface for medical image filtering, segmentation, and registration," in *Proc. SPIE Conf. Medical Imaging*, 2006, vol. 6144, pp. 1–8.
- [15] L. Ibanez, W. Schroeder, L. Ng, and J. Cates, *The ITK Software Guide*, Kitware, Inc., second edition, 2005.
- [16] A. P. Zijdenbos, B. M. Dawant, R. A. Margolin, and A. C. Palmer, "Morphometric analysis of white matter lesions in MR images: Method and validation," *IEEE Trans. Med. Imag.*, vol. 13, pp. 716 – 724, Dec. 1994.

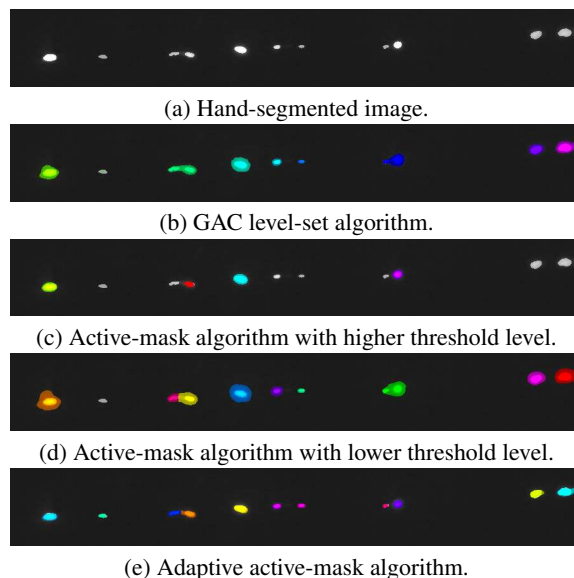


Fig. 5. Segmentation results. Each colored region corresponds to a different mitochondrion, overlaid on the hand-segmentation masks in white for comparison. Colors are selected randomly.