

Deconvolving Active Contours for Fluorescence Microscopy Images

Jo A. Helmuth and Ivo F. Sbalzarini*

Institute of Theoretical Computer Science and Swiss Institute of Bioinformatics,
ETH Zurich, Switzerland

Abstract. We extend active contours to constrained iterative deconvolution by replacing the external energy function with a model-based likelihood. This enables sub-pixel estimation of the outlines of diffraction-limited objects, such as intracellular structures, from fluorescence micrographs. We present an efficient algorithm for solving the resulting optimization problem and robustly estimate object outlines. We benchmark the algorithm on artificial images and assess its practical utility on fluorescence micrographs of the Golgi and endosomes in live cells.

1 Introduction

Active contours are among the most important frameworks for image segmentation. In the original formulation by Kass et al. [1], a contour is defined as a (closed or open) parametric curve in the image domain that minimizes an energy functional. Closed active contours can also be represented implicitly as level sets [2]. This is particularly beneficial when the number of objects to be segmented is not known a priori since it allows for topology changes during energy minimization. In both representations, the energy functional consists of two terms: (1) an *external energy* that depends on image data, and (2) an *internal energy* that solely depends on the geometry of the contour. While the former defines an unconstrained image segmentation problem, the latter provides regularization, helps overcoming local minima, and allows bridging regions with little information in the image. Many extensions of active contours have been proposed over the last two decades, including active masks [3,4], active surfaces [5], and stochastic active contours (STACS) [6,7]. For implicit active contours, very efficient algorithms have been developed based on narrow-band level sets [8] or graph cuts to minimize the energy functional [9].

Active contours are widely used in biological light-microscopy imaging. Their application has, however, so far been restricted to images of objects well above the resolution limit of the imaging equipment. If the size of the object becomes comparable to the width of the point spread function (PSF) of the microscope, the objects are under-resolved and active contour segmentations can no longer

* This work was supported by ETH Research Commission grant TH-1007-1 and the Swiss SystemsX.ch initiative, evaluated by the Swiss National Science Foundation.

be considered unbiased estimates of the object's geometry [10]. The only exception are objects of dimension 0 or 1, imaged with a symmetric PSF. This includes particles modeled as points [11] and filaments modeled as lines [12]. For extended objects of co-dimension zero, deconvolution can be used to estimate their geometry from an image blurred by the PSF. Deconvolution is, however, an inverse problem that is known to be ill posed [13,14]. Moreover, direct linear deconvolution is not feasible for sub-cellular structures [15]. Therefore, constrained iterative methods have to be used.

In this paper, we extend the active contour framework to constrained iterative deconvolution by including models of the objects and the imaging process. This involves replacing the external energy functional with a negative log-likelihood function and optimizing it under the constraint of the internal energy. This optimization involves simulating the forward imaging process. We propose an efficient minimization algorithm that uses a domain decomposition approach and exploits the linearity of the convolution operator. The resulting method is an iterative deconvolving active contour that is constrained by the object model and the imaging model. We demonstrate the accuracy and precision of deconvolving active contours on synthetic benchmark images of sub-resolution objects. We show that the present framework allows unbiasing the estimation of object geometries from fluorescence micrographs. We further demonstrate the applicability of the proposed algorithm to images of the Golgi complex and endosomes in live cells. The resulting estimated outlines allow biological observations that were not possible before.

2 Energy Functional

We seek a parametric description of a set S of outlines of objects, supported by an error-corrupted digital image $I_m = I + \varepsilon$. Given an imaging model $I(S)$, parameters Θ of the set of N objects $S = \{\Theta^k\}_{k=1}^N$ have to be found that best explain the measured image I_m . The imaging model $I(S)$ predicts the image I of the set S of objects in the absence of noise. This parameter estimation problem can be rephrased in the explicit active contour framework. Hereby, the outline of an object k is represented by a piece-wise linear spline $\Theta^k = [x_1^k, y_1^k, \dots, x_{n_k}^k, y_{n_k}^k]^T$. As described at the end of Sec. 3, the computational cost of the algorithm is not significantly influenced by the number of control points n_k used. We minimize the sum of internal and external energy

$$E(I_m, S) = \sum_{\Theta \in S} (E_b(\Theta) + E_s(\Theta)) + E_{\text{ext}}(I_m, I(S)), \quad (1)$$

where the external energy is given by the similarity between the model image $I(S)$ and the real image I_m , quantifying the likelihood that the objects S have indeed created the observed image I_m . The internal energy comprises regularizations for bending and stretching of the contour as:

$$\begin{aligned} E_b(\Theta) &= \beta \sum_{i=1}^n \|\mathbf{x}_{i+1} - 2\mathbf{x}_i + \mathbf{x}_{i-1}\|^2 \text{ and} \\ E_s(\Theta) &= \alpha \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{x}_{i-1}\|^2 \end{aligned} \quad (2)$$

with $\mathbf{x}_i = (x_i, y_i)$ of the spline Θ . While the “bending stiffness” β limits undulations of the outlines, the “stretching stiffness” α constrains the shrinking of outlines to significant image energy gradients. We define $I(S)$ based on a function $O(i, j)$ over the image pixels $\{x_i\} \times \{y_j\}$. This *object intensity function* O represents the pre-imaging objects up to a multiplicative constant. In fluorescence microscopy, O is proportional to the concentration of fluorophores at each point in the focal plane. Formation of the model image (Fig. 2B) is done by convolving the object intensity function O with the (measured) PSF P of the microscope:

$$I = O * P, \quad (3)$$

where $*$ denotes the discrete convolution operator. In practice, we sample $O(i, j)$ at higher spatial resolution (two, three, or four-fold) than the measured image in order to include sub-pixel information. This requires down-sampling of I before comparison to I_m . The object intensity function $O(i, j)$ is defined from the set S of outlines by setting to non-zero values only the pixels (i, j) close to or enclosed by one of the outlines Θ^k :

$$O(i, j) = \begin{cases} c^k & \text{if } (x_i, y_j) \text{ enclosed by } \Theta^k \\ (1-d)c^k & \text{if } d = D((x_i, y_j), S) < 1 \\ 0 & \text{else} \end{cases}, \quad (4)$$

where D is the distance to the closest spline Θ^k and c^k the constant intensity of object k . Similar to the simplified Mumford-Shah functional [2], this object intensity function is piecewise constant, but with linearly decaying intensities at the boundaries. We favor this piecewise linear functional over more complex models as it increases the robustness of the estimator on noisy data. The external energy E_{ext} is given by:

$$E_{\text{ext}} = \sum_i \sum_j R(i, j) (I_m(i, j) - I(i, j))^2. \quad (5)$$

The weight matrix R allows including a model for the distribution of the imaging noise ε . In the absence of knowledge about ε , or for Gaussian white noise, R is the identity matrix.

Minimizing E over the Θ^k yields an estimate of $O(i, j)$. Direct estimation of O based on I_m amounts to direct linear deconvolution, a problem known to be ill posed. The present framework can thus be interpreted as a constrained, iterative deconvolution [13,14] with the constraints defined in Eqs. 2 and 4.

3 Minimization of the Energy Functional

We assume that initial estimates of the outlines are provided by a suitable pixel-based segmentation procedure and that they represent the correct topology of the objects. Further, the 2D PSF of the imaging device is assumed to be known. Minimizing Eq. 1 could then be done using any general purpose optimizer. A specialized procedure exploiting the structure of this high-dimensional problem

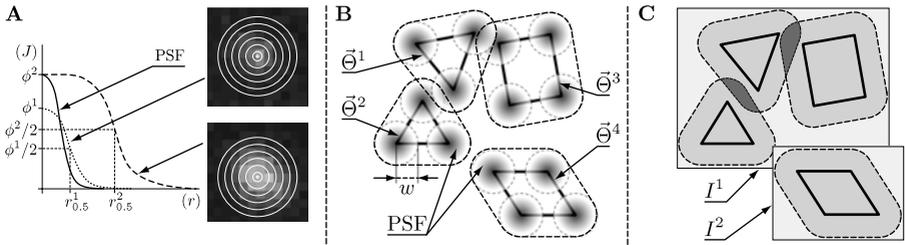


Fig. 1. (A) Radial intensity profile of a small (top image, dotted line) and a large (bottom image, dashed line) object. The intensity profile of the small object and the PSF (solid line) have a similar shape. (B) The width w of the PSF defines the influence regions (dashed lines) of the objects Θ^k (solid lines). (C) Objects with overlapping influence regions (dark gray areas) are grouped together.

is, however, favorable. Our algorithm consists of three steps: (1) estimating the object intensities c^k , (2) decomposing the image into smaller parts, and (3) estimating the precise outlines Θ^k of all objects.

The **object intensities** c^k are separately **estimated** in order to use them in Eq. 4 and as a biologically relevant readout. For objects that are far larger than the width of the PSF, the object intensity is approximately equal to the intensity ϕ in the center of the image of the object (bottom image in Fig. 1A). Smaller objects have a central intensity that is reduced by a factor κ . In order to estimate κ^k for a given outline Θ^k , we analyze the radial intensity profile $J(r)$ of the object (Fig. 1A), found by averaging interpolated intensities along concentric circles around the intensity centroid. The same procedure is also used to measure the radially symmetric 2D PSF of the microscope from images of point-like sources such as fluorescent beads. The half width at half maximum (HWHM) $r_{0.5}$ of $J(r)$ serves as a size parameter (Fig. 1A). Since the dependence $\kappa = f(r_{0.5})$ is not explicitly known, we empirically calibrate it on synthetic images, generated by convolving circular objects of different sizes and known intensities c^i with the PSF. For each synthetic image I^i we measure $r_{0.5}^i$ and the central intensity ϕ^i , and then compute the calibration function as $\bar{\kappa}^i(r_{0.5}^i) = c^i/\phi^i$. Based on this function, the object intensity c^k of an experimentally observed object k can later be estimated as $c^k = \phi^k \bar{\kappa}(r_{0.5}^k)$ using (linear) interpolation. For point-like objects, $r_{0.5}$ converges to the HWHM of the PSF and does no longer vary with object size. The analysis of sizes and shapes is, therefore, restricted to objects with $r_{0.5}$ above an empirically determined threshold of 1.1–1.5 HWHM. Smaller objects are treated as circles with centers at the observed intensity centroid and object intensities c found by least squares regression on Eqs. 3 and 5.

In order to accelerate the computations and estimate the outlines of different objects on different computer processor cores in parallel, we **decompose the image** into smaller, independent parts. This is possible since the PSF P is of limited spatial extent and the influence of objects on distant pixels in the image can be neglected. We define the radius of the region of influence around each outline Θ^k as the width w where the PSF has decayed below 1% of its

peak (Fig. 1B). Objects with overlapping influence regions are grouped together and their influence regions are merged. This decomposes the image into (not necessarily disjoint) rectangular sub-images I^l containing the merged influence regions (Fig. 1C). Eq. 1 is then evaluated independently for each sub-image and we no longer consider pixels that are in no sub-image.

Estimating the outlines of the objects to sub-pixel resolution is done by minimizing the energy functional in Eq. 1 with respect to the Θ^k . This amounts to implicit deconvolution. The energies of different sub-images are minimized independently in parallel. We use the optimizer proposed by Kass et al. [1] with an explicit Euler method of adaptive step size. This requires approximating the partial derivatives of E with respect to the x and y positions of all spline points (x_i^k, y_i^k) , $i = 1, \dots, n_k$. We use a finite-difference approximation to the derivatives with respect to a point (x_i^k, y_i^k) on the outline Θ^k by displacing the point by a dynamically adapted Δx . This yields a deformed hypothetical outline Θ_*^k in a new set of outlines S_* . The derivative of the external energy E_{ext}^l in sub-image l is then approximated as:

$$\frac{\partial E_{\text{ext}}^l}{\partial x_i^k} \approx \frac{E_{\text{ext}}^l(I_m^l, I^l(S_*)) - E_{\text{ext}}^l(I_m^l, I^l(S))}{\Delta x}. \quad (6)$$

The derivatives with respect to other points and in the y -direction are found similarly. The hypothetical sub-image $I^l(S_*)$ needs to be computed once per iteration of the minimization procedure. Since convolution is a linear operation, the image $I^l(S_*)$ can be expressed as the sum of $I^l(S)$ and a change ΔI^l caused by the deformation of Θ^k . ΔI^l is found by computing $O(i, j)_*^k$ from Θ_*^k , subtracting it from $O(i, j)^k$, and convolving this difference $\Delta O(i, j)^k$ with P . This drastically reduces the number of compute operations and the run-time. The computational cost of the algorithm is dominated by the cost of the convolutions; the cost of the active contour updates is insignificant. The total cost of all convolutions is proportional to the total number of non-zero entries in *all* $\Delta O(i, j)^k$, which scales linearly with the length of the outline. Reducing the number of control points n_k would thus not lead to significant computational savings since the total length of the outline (and hence the computational cost of the convolutions) remains the same. After computing all partial derivatives of the image error, one Euler step is performed. Minimization starts from the initial outline estimates and ends when the L_∞ -norm of the change of outlines between two iterations, $\|\Theta_{i+1} - \Theta_i\|_\infty$, is below a user-defined threshold. This implies that the gradient-descent minimizer has converged to a (local) minimum. Local minima can possibly be escaped by increasing α .

4 Benchmarks

We quantify the accuracy and precision of the presented algorithm on synthetic benchmark images of diffraction-limited objects. Images are generated as illustrated in Fig. 2A to C using a real, measured PSF with full width at half maximum FWHM = 322 nm. We consider images of a circular object with diameter

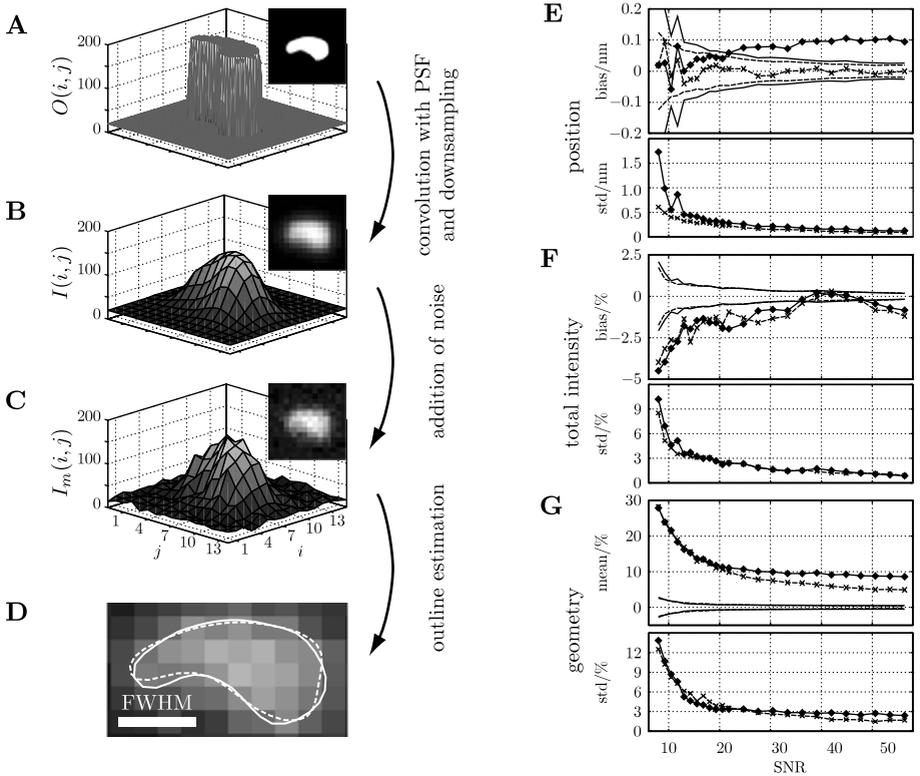


Fig. 2. (A–C) Generation of the synthetic benchmark images. (D) True (solid line) and estimated (dashed line) outline for SNR = 12.5 (geometry error = 14%). (E–G) Benchmark results for a pear-shaped (solid lines with diamonds) and circular (dashed lines with crosses) object. Lines without markers delimit the ± 1 standard error interval.

1.5 FWHM and of a 0.6 to 1.2 FWHM wide and 2.8 FWHM long pear-shaped object. The ground-truth $O(i, j)$ is generated from the true outlines according to Eq. 4. The object intensity c is set to 200 and a background level of $b = 20$ is added (Fig. 2A). Different amounts of noise yield signal-to-noise ratios $\text{SNR} = (c - b) / \sigma_c$ in the range of 7 to 56, where σ_c is the noise level in the center of the objects. For each object and SNR, 250 test images are generated, object outlines are estimated, and compared to ground truth.

We quantify accuracy and precision of the estimated position, total intensity, and geometry (Fig. 2E to G). The position error is defined as the difference between the true and estimated x -position of the intensity centroid¹. The total intensity error is given by the difference in total intensity (sum over all $O(i, j)$ enclosed by the outline) between the estimated object and ground truth, divided by ground truth. The geometry error is defined as the sum of non-overlapping

¹ We prefer this over the Euclidean distance in (x, y) since it enables correlating shape asymmetries with the position errors in the different directions.

areas of the true and estimated outlines, normalized by the area enclosed by the true outline. For both shapes, precision and accuracy of the estimated position are in the range of a few nanometers (Fig. 2E). The errors in the y -direction are comparable for the pear-shaped object and identical for circular shape (data not shown). The position bias for the circular shape is always within an interval of \pm the standard error, and hence not significant. As expected, a small systematic position bias can be observed for the pear-shaped object. Since this shape is not symmetric, the underestimation of high curvatures due to E_b causes a small shift in the position estimate toward the less curved side. For both shapes the standard deviation of the relative total intensity error drops below 5% for SNRs larger than 10 (Fig. 2F). The bias is larger than the standard error, almost always negative, and converges to about -1%. This is due to E_s favoring shorter outlines and thus decreasing the enclosed object intensity. The means and standard deviations of the geometry errors of both shapes converge to values of less than 10% and 3%, respectively (Fig. 2G). The bending and stretching energies E_b and E_s prevent the mean geometry error from converging to zero, and the underestimation of high curvatures causes an additional bias for the more complex pear-shaped object. Nevertheless, we observe that the estimated outlines visually reproduce well the essential characteristics of the true outlines, even at SNRs below 15 and, therefore, error levels $>10\%$ (Fig. 2D).

Adjusting the bending stiffness β (Eq. 2) allows trading the accuracy of the outline estimation against its robustness by limiting undulations of the contour. Low values lead to a higher noise sensitivity (less regularization), but allow better estimation of high curvatures. In order to qualitatively assess this trade-off, we apply the algorithm to synthetic images of a triangle (Fig. 3), generated as described above. As expected, we observe that low SNRs favor high values of β , and vice versa. The algorithm is robust over two orders of magnitude of β . Only the most extreme case ($\beta = 0.02$, SNR = 5) exhibits significant shape instabilities. The stretching stiffness α has much less influence on the final contour. Higher values lead to faster convergence of the algorithm and better escape from local minima. At the same time, however, they bias the outlines to shorter, more contracted contours. We find a value of $\alpha = 0.005$ to be sufficient to overcome local minima and speed up convergence.

5 Application to Real Data

We demonstrate the utility of the present algorithm on fluorescence microscopy images of different intracellular structures. First, we apply it to an image of the Golgi complex in HeLa cells labelled by fluorescent giantin antibodies (Fig. 4A and B). The Golgi is a complex-shaped intracellular organelle composed of membrane stacks of about $5\ \mu\text{m}$ size. The same image was also used to demonstrate active mask segmentation [3]. We show how such a coarse, pixel-level segmentation can be refined by the present implicit deconvolution method. We start from a rough manual segmentation obtained from Fig. 11d of Ref. [3]. Since no information about the PSF was available, we model it by a Gaussian with

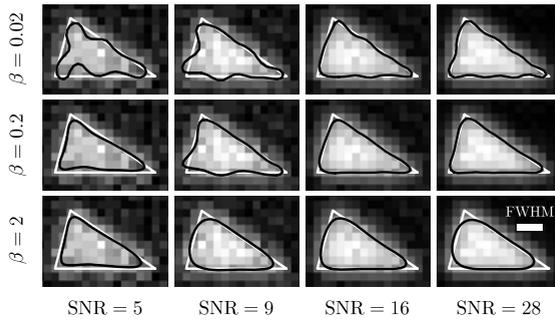


Fig. 3. Estimation of a highly curved outline at different SNRs and values of β ($\alpha = 0.005$ in all cases). White lines: true outlines; black lines: estimated outlines.

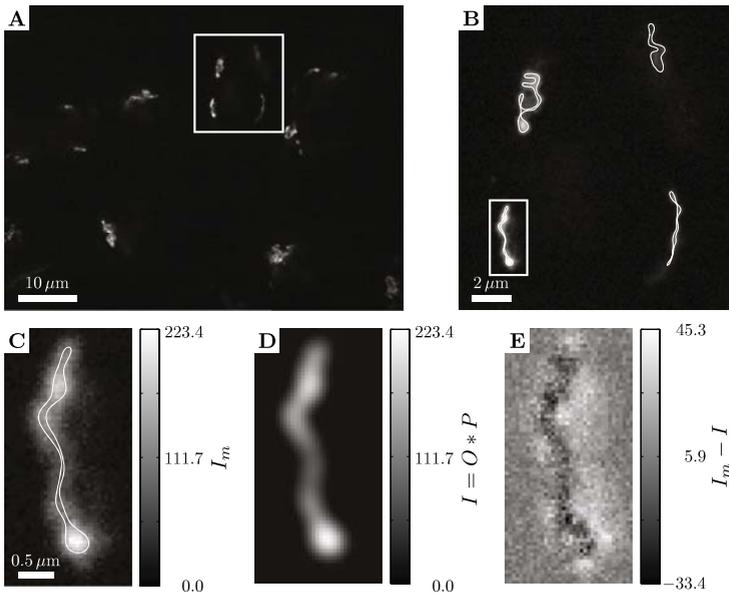


Fig. 4. Outline estimation of the Golgi complex in HeLa cells [16]. (A) Original image containing multiple cells. Magnification of a single cell (B) and a single Golgi segment (C) with estimated outlines (white line); model image (D) and residual error (E).

$\sigma = 150$ nm. This is a conservative choice for the imaging set-up used (spinning disk confocal, $NA = 1.4$, oil immersion). The final deconvolving active contours (Fig. 4B and C) capture well the morphological characteristics of the Golgi and the model image I (Fig. 4D) is remarkably close to the real image (Fig. 4C). The estimated outline shows no obvious signs of over-fitting. The residual error $I_m - I$ (Fig. 4E) shows that the model image trends to be too bright in the center of the object and under-represents the blur around it. As shown below, this is likely due to a too narrow model PSF.

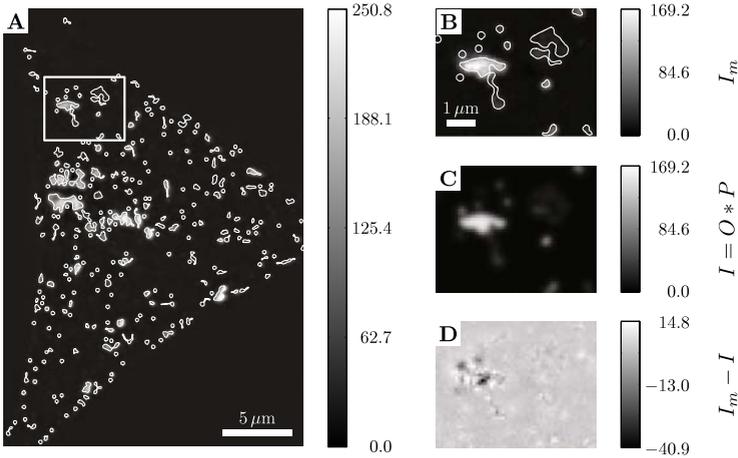


Fig. 5. Outline estimation (white lines) of endosomes in a live HER 911 cell (A). (B) Magnification, (C) model image, and (D) residual error in the magnified region.

The second application considers live HER 911 cells expressing EGFP-tagged Rab5, a protein marker for endosomes. With diameters of about 500 nm, endosomes are much smaller than the Golgi and they appear more compact. Initialized with watershed segmentation, the present algorithm estimates the outlines as shown in Fig. 5A and B. The PSF (FWHM = 322 nm) of the microscope (spinning disk confocal, NA = 1.35, oil immersion) was measured from images of sub-diffraction objects as described in Sec. 3. Fig. 5B shows complex-shaped endosome outlines in close vicinity. The outlines follow well the subjective contours in the images, even for very dim objects. The correspondence between the real and the model image (Fig. 5C) is remarkable. Also, unlike in the Golgi case, there is no clear trend in the residual error (Fig. 5D), highlighting the advantage of using the true, measured PSF. Except for slight over-estimation of the central intensity of the large object on the left, the residual error is dominated by detector noise.

6 Conclusions and Discussion

We have introduced deconvolving active contours, extending explicit active contours to iterative constrained deconvolution by replacing the image-based external energy with a model-based likelihood function that includes prior knowledge about the imaging process. The algorithm iteratively refines an initial image segmentation using regularized optimization. Optimizing the likelihood function is computationally more involved than optimizing a classical pixel-based energy. We have thus introduced a special-purpose algorithm that uses domain decomposition parallelism and exploits the linearity of the convolution operator. Run-times on a desktop computer are on the order of seconds for individual objects

(e.g. 1.2s for the object in Fig. 2D). The presented algorithm enables estimating the outlines of diffraction-limited, asymmetric objects to sub-pixel accuracy. The benchmarks demonstrated localization precision in the nanometer range (better than 0.01 FWHM) and fluorescence intensity estimation to within a few %. We have further demonstrated the practical utility of deconvolving active contours on images of the Golgi complex and endosomes in live cells.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *Int. J. Computer Vision*, 321–331 (1988)
2. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Trans. Image Process.* 10, 266–277 (2001)
3. Srinivasa, G., Fickus, M.C., Guo, Y., Linstedt, A.D., Kovačević, J.: Active mask segmentation of fluorescence microscope images. *IEEE Trans. Image Process.* 18, 1817–1829 (2009)
4. Srinivasa, G., Fickus, M.C., Gonzalez-Rivero, M.N., Hsieh, S.Y., Guo, Y., Linstedt, A.D., Kovačević, J.: Active mask segmentation for the cell-volume computation and Golgi-body segmentation of HeLa cell images. In: *Proc. IEEE Int. Symp. Biomed. Imaging*, pp. 348–351 (2008)
5. Dufour, A., Shinin, V., Tajbakhsh, S., Guillen-Aghion, N., Olivo-Marin, J.C., Zimmer, C.: Segmenting and tracking fluorescent cells in dynamic 3-D microscopy with coupled active surfaces. *IEEE Trans. Image Process.* 14, 1396–1410 (2005)
6. Pluempitwiriyawej, C., Moura, J.M.F., Wu, Y.J.L., Ho, C.: STACS: New active contour scheme for cardiac MR image segmentation. *IEEE Trans. Med. Imaging* 24, 593–603 (2005)
7. Coulot, L., Kirschner, H., Chebira, A., Moura, J.M.F., Kovačević, J., Osuna, E.G., Murphy, R.F.: Topology preserving STACS segmentation of protein subcellular location images. In: *Proc. IEEE Int. Symp. Biomed. Imaging*, pp. 566–569 (2006)
8. Shi, Y., Karl, W.: Real-time tracking using level sets. In: *Proc. IEEE Conf. CVPR*, vol. 2, pp. 34–41 (2005)
9. El-Zehiry, N., Elmaghraby, A.: An active surface model for volumetric image segmentation. In: *Proc. IEEE Int. Symp. Biomed. Imaging*, pp. 1358–1361 (2009)
10. Streekstra, G.J., van Pelt, J.: Analysis of tubular structures in three-dimensional confocal images. *Network: Comput. Neural Syst.* 13, 381–395 (2002)
11. Sbalzarini, I.F., Koumoutsakos, P.: Feature point tracking and trajectory analysis for video imaging in cell biology. *J. Struct. Biol.* 151, 182–195 (2005)
12. Li, H., Shen, T., Smith, M.B., Fujiwara, I., Vavylonis, D., Huang, X.: Automated actin filament segmentation, tracking and tip elongation measurements based on open active contours. In: *Proc. IEEE Int. Symp. Biomed. Imaging*, pp. 1302–1305 (2009)
13. Sibarita, J.B.: Deconvolution microscopy. *Adv. Biochem. Eng. Biot.* 95, 201–243 (2005)
14. Sarder, P., Nehorai, A.: Deconvolution methods for 3-D fluorescence microscopy images. *IEEE Signal Process. Mag.* 23, 32–45 (2006)
15. Meijering, E., Smal, I., Dzyubachyk, O., Olivo-Marin, J.C.: Time Lapse Imaging. In: *Microscope Image Processing*, ch. 15. Academic Press, London (2008)
16. Guo, Y., Linstedt, A.D.: COPII-Golgi protein interactions regulate COPII coat assembly and Golgi size. *J. Cell Biol.* 174, 53–63 (2006)