Published in final edited form as:

Signal Image Video Process. 2014 December 1; 8(1 Suppl): 41–48. doi:10.1007/s11760-014-0694-8.

Graph-regularized 3D shape reconstruction from highly anisotropic and noisy images

Christian Widmer,

Sloan Kettering Institute, 1275 York avenue, New York, NY, USA

Stephanie Heinrich,

Institute for Biochemistry, ETH Zurich, Otto-Stern-Weg 3, Zurich, Switzerland

Philipp Drewe,

Sloan Kettering Institute, 1275 York avenue, New York, NY, USA

Xinghua Lou,

Sloan Kettering Institute, 1275 York avenue, New York, NY, USA

Shefali Umrana, and

Sloan Kettering Institute, 1275 York avenue, New York, NY, USA

Gunnar Rätsch

Sloan Kettering Institute, 1275 York avenue, New York, NY, USA

Christian Widmer: ckwidmer@gmail.com; Stephanie Heinrich: stephanie.heinrich@bc.biol.ethz.ch; Philipp Drewe: drewe@cbio.mskcc.org; Xinghua Lou: loux@cbio.mskcc.org; Gunnar Rätsch: raetsch@cbio.mskcc.org

Abstract

Analysis of microscopy images can provide insight into many biological processes. One particularly challenging problem is cellular nuclear segmentation in highly anisotropic and noisy 3D image data. Manually localizing and segmenting each and every cellular nucleus is very time-consuming, which remains a bottleneck in large-scale biological experiments. In this work, we present a tool for automated segmentation of cellular nuclei from 3D fluorescent microscopic data. Our tool is based on state-of-the-art image processing and machine learning techniques and provides a user-friendly graphical user interface. We show that our tool is as accurate as manual annotation and greatly reduces the time for the registration.

Keywords

Cell nuclei detection; Shape reconstruction; 3D fluorescent microscopic data; Nuclear segmentation

1 Introduction

Imaging data, such as those from microscopic experiments, are a unique source of information in biology. Through fluorescent staining, they allow to investigate tissue composition, cell shapes and subcellular localization patterns. Consistent measurements of such data are often performed manually and thus time-consuming, which remains an obstacle in large-scale experiments. Methods that assist processing such complex, large data are therefore needed. These methods should not only speed up these measurements steps but also increase the reproducibility of the measurements. In this work, we will focus on the challenge of detecting cell nuclei from fluorescent microscopy images. For fluorescence microscopy, it is common practice for biologists to manually segment cells based on 3D visualization and then later quantify the signal within this segmentation (usually in a different channel). In particular, it is essential to address segmentation for anisotropic and highly noisy 3D microscopic images with possible staining defect—a very challenging problem that cannot be handled robustly by conventional computer vision methods such as blob detection [1–3], deformable model (e.g., level set [4–6] or combinational optimization (e.g., graph cut) [7,8], because staining defects normally lead to missing intensity within the body of the true nucleus (see supplement). A robust approach is highly desired.

Though microscopes are sometimes equipped with software to assist researchers on this task, more often than not the existing software only provides very rough polygon fits based on intensity values (see below or Figure 6 in the supplement). For instance, the microscope used in [9] was equipped with the 2D polygon finder software (DeltaVision SoftWorx software, Applied Precision/GE Healthcare). A major drawback of this software solution is that 3D information is not taken into account, which means segmentations are performed for each layer individually. Furthermore, each object in a bigger volume (with potentially hundreds of cells) has to be processed separately, making this step a major bottleneck. Finally, any prior knowledge about the structure of the objects of interest is ignored as fits are often nonparametric (e.g., [10,11]). This is suboptimal when segmenting cells or nuclei, as these objects have known structure that can be exploited. Due to these drawbacks, signal quantification is a time-consuming task, and large-scale quantification experiments become prohibitive. Besides, existing 3D segmentation programs (which often need a high signal-to-noise ratio) often perform tracking and segmentation in order to separate individual cellular structures from each other but do not extract quantitative information from the segmentation, for example total signal intensity in one fluorescence channel or a different channel from the same 3D stack [12–14]. For instance, [2] aids with the segmentations of cells in 3D images based on gradient flow tracking, a possible drawback being that the final segmentation is obtained using adaptive thresholding, which—in our experience—lead to problems in upper and lower planes of a 3D stack. Further, there is the widely used MATLAB package Level Set [15], which provides efficient implementations of level-set-based segmentation algorithms. In contrast to our approach, which is based on a parametric model that describes the biological structures that we are interested in, this library is quite general and makes it harder to encode special purpose prior knowledge. Similarly, [16] propose a contour-based minimum-model approach to cellular nuclei segmentation, which gives them more flexibility, possibly at the expense of recovering an accurate parametric description of

ellipsoid-shaped nuclei in the presence of highly noisy data. Finally, in practice, many biologists use the software package Imaris (Bitplane/Oxford Instruments), a commercial software for 3D and 4D reconstruction, for which no source code is available (see below for a more detailed description of this software solution). We propose a new method that addresses some of these shortcomings. It can be applied to images containing multiple cells and is hand-tailored to nuclei of ellipsoid shape. The method adapts graph-regularized transfer learning [17] to the problem of parametric fitting in several layers in combination with a robust loss function, as used in support vector regression, to minimize the need for manual post-processing. Our proposed method thus provides biologists with a tool for high-throughput quantification experiments. In summary, our approach GRED (graph-regularized 3D shape reconstruction) brings together several ideas to come up with an optimized tool for a specific detection problem. In contrast to many existing software solutions, this may restrict flexibility in the type of biological objects we can track, but provides a highly accurate and efficient approach in the face of noisy data.

2 Methods

Our method performs the fitting in two steps: a preprocessing step to localize the nuclei that is based on multi-scale Hessian eigenvalue thresholding [18], followed by a parametric fitting procedure to compute the shape of each nucleus. These two steps will be explained in detail in the following.

2.1 Preprocessing

For preprocessing, we localize and extract individual nuclei from a larger volume. We apply Hessian eigenvalue thresholding introduced in [18], which finds sets of foreground pixels that cluster together. For this, we use Gaussian smoothing to aggregate mass from the neighborhood and emphasize the central regions of the blobs (local maxima). A Hessian representation is then used to find those local maxima, exploiting that they have negative Hessian eigenvalues while a stationary point (e.g., a saddle point) does not [18,19]. This is one for multiple resolutions in order to find clusters of all sizes. Therefore, simple thresholding of the eigenvalues can extract the foreground (cell). For a visualization of the individual steps of the method, see Fig. 1. We repeat this procedure at different scales (e.g., {2, 3, 4, 5, 6}) because a large Gaussian kernel strongly suppresses noise but yields merge errors (i.e., under-segmentation, because mass is aggregated within a larger neighborhood), while a small Gaussian kernel is sensitive to noise but better preserves the boundary. Results at different scales have characteristics that are complementary to each other and combining them produces less false-positives and merge errors [20].

2.2 Parametric fits

Our method is optimized for the detection of cellular nuclei, which are membrane-enclosed organelles in eukaryotic cells that contain most of the cells genetic material. The shape of these objects resembles a deformed ellipsoid. We argue that we can incorporate this prior knowledge about the shape by fitting parametric geometric objects such as an 3D ellipsoid or stack of 2D ellipses. An example how parametric fitting is beneficial in providing robust fits in the face of missing data points is shown in Fig. 2. While we limit ourselves to

discussing ellipsoid structures for the rest of this paper, our framework generalizes to other geometric objects such as splines.

2.2.1 Fitting circles—We start with the simplest parametric object we could use for this task: a circle (one for each plane of the 3D stack). This has obvious limitations, as many nuclei are not perfect circles, but rather correspond to ellipses in each layer. However, it may be easily derived and therefore constitutes a good starting point for the description of our method.

The distance of a point $\mathbf{x} \in \mathbb{R}^2$, $i \in \{1, \dots, n\}$ to a circle with center $\mathbf{c} \in \mathbb{R}^2$ and radius $r \in \mathbb{R}$ is easily computed as

$$d(\mathbf{c}, r, \mathbf{x}) = |\|\mathbf{c} - \mathbf{x}\| - r|$$

Finding a circle parametrized by $\mathbf{c}$ and r that minimizes the sum of distances to points $\mathbf{x}_i$ corresponds to solving the following optimization problem:

$$\min_{\mathbf{c}, r} \sum_{i=1}^n L(d(\mathbf{c}, r, \mathbf{x}_i)),$$

where L is a loss function, such as the squared loss or the hinge loss. The choice of loss function L has important implications on the properties of the fit (e.g., robustness).

2.2.2 Fitting ellipses—A class of shapes that allows more flexibility for fitting nuclei in 2D are ellipses. An ellipse in 2D can be parametrized by a center point $\mathbf{c} = [c_x, c_y]^\top$ and two radii $\mathbf{r} = [r_x, r_y]^\top$ [21]. The points $[x, y]^\top$ on the ellipse centered at $\mathbf{c}$ are then given by the equation

$$\frac{(x - c_x)^2}{r_x^2} + \frac{(y - c_y)^2}{r_y^2} = 1. \quad (1)$$

An alternate parametrization (general conic) is given by

$$ax^2 + bxy + cy^2 + dx + ey + f = 0, \quad (2)$$

describing an ellipse if $b^2 - 4ac < 0$ [22]. Let

$$\mathbf{x} = [x^2, xy, y^2, x, y, 1]^\top, \quad (3)$$

$$\Theta = [a, b, c, d, e, f]^\top,$$

then points on the ellipse satisfy $\mathbf{x}^\top \Theta = 0$. The algebraic distance f of a point $\mathbf{x}$ to the ellipse parametrized by Θ is defined as:

$$f(\mathbf{x}, \Theta) = \mathbf{x}^\top \Theta. \quad (4)$$

The algebraic distance is an approximation of the Euclidean distance that has the advantage that it is much easier to compute.

Avoiding degenerate solutions: Additional constraints are necessary to avoid degenerate solutions. In order to avoid the trivial solution $\Theta = 0$ and recognizing that any multiple of a solution Θ represents the same conic, the parameter vector Θ is constrained in one way or the other [23]. Different algorithms for fitting ellipses often only differ in the way they constrain parameters. Many authors suggest $\|a\|^2 = 1$, others $a + c = 1$ or $f = 1$ [23].

Minimizing algebraic distance: For a general loss function, we arrive at the following formulation:

$$\min_{\Theta} \sum_{i=1}^N L(\Theta^\top \mathbf{x}_i) \quad (5)$$

s.t. solution nondegenerate

Depending on which combination of nondegeneracy constraints and loss function is used, different solvers are needed.

2.2.3 Robust loss function—It is well established that the squared loss is particularly prone to outliers [24,25], as distance is penalized quadratically. We therefore propose to use the ε -insensitive loss function for the problem at hand, where ε defines the width of the region, within which points inflict no error. The ε -insensitive loss has its background in the context of support vector regression [26–28]. It is also known as dead-zone penalty in other contexts [29] and is often used when a more robust error function is needed.

$$L_\varepsilon(r) = \begin{cases} |r| - \varepsilon, & \text{if } |r| > \varepsilon \\ 0 & \text{else.} \end{cases} \quad (6)$$

It has two important properties that make it appealing. The first is that it does not penalize points that are within a rim of the stacked ellipsoid. This captures the intuition that the nuclear membrane has a certain thickness, and we therefore do not want to penalize points that are within the membrane. Second, the loss is affine (linear minus some offset that depends on ε). This means that outliers are not penalized as severely as with the squared loss, yielding a more robust error function. In the following, we show that although non-differentiable, the ε -insensitive loss may be expressed in the form of a constrained optimization problem. As a first step, we note that L_ε may be written as

$$L_\varepsilon(r) = \max(|r| - \varepsilon, 0). \quad (7)$$

Plugging the above loss into Eq. (5) and neglecting the nondegeneracy constraints for now yield the optimization problem:

$$\min_{\Theta} \sum_{i=1}^N \max(|\mathbf{x}_i^\top \Theta| - \varepsilon, 0). \quad (8)$$

We now make use of the fact that $\max(a, b)$ can be expressed to the smallest upper bound of a and b [29], i.e.,

$$\begin{aligned} \max(a, b) &= \min_c \\ \text{s.t. } a &\leq c, \\ b &\leq c. \end{aligned} \quad (9)$$

Furthermore, we exploit that the absolute value may be expressed as the maximum of two linear functions $|r| = \max(r, -r)$. We use the latter to move $|\mathbf{x}_i^\top \Theta|$ from the objective to the constraints using newly introduced slack variables s_i . This gives rise to:

$$\begin{aligned} \min_{\Theta, s_i} \sum_{i=1}^N \max(s_i - \varepsilon, 0) \\ \text{s.t. } \mathbf{x}_i^\top \Theta \leq s_i, \\ -\mathbf{x}_i^\top \Theta \leq s_i \end{aligned} \quad (10)$$

Using the same scheme, the other max is moved to the constraints using Eq. (9), introducing variables t_i . We arrive at:

$$\begin{aligned} \min_{\Theta, s_i, t_i} \sum_{i=1}^N t_i \\ \text{s.t. } \mathbf{x}_i^\top \Theta \leq s_i, \\ -\mathbf{x}_i^\top \Theta \leq s_i, \\ s_i - \varepsilon \leq t_i, \\ 0 \leq t_i. \end{aligned} \quad (11)$$

2.2.4 Graph regularization—So far, we have derived a method for robust parametric fits in $2D$. In principle, we could extend this approach to $3D$ parametric objects. However, the data we consider are highly anisotropic in the sense that the z -axis has lower resolution and should therefore be treated differently. This motivates our alternative approach. To share information of the ellipse fitting across the z -layers, we propose to jointly fit ellipses in all layers and penalize differences between parameter vectors of neighboring layers by means of regularization term R :

$$R(\Theta_1, \dots, \Theta_N) = \sum_{i=1}^{N-1} \|\Theta_i - \Theta_{i+1}\|_p, \quad (12)$$

where $\|a\|_p$ is the p -norm. The effect of this smoothing is shown in Figure 7 in the Supplement. This smoothness regularizer is a special case of a general graph regularizer, which is often used in the context of multitask learning [17,30], as edges only exist between neighboring layers. Note that in the above formulation, we have not settled on a particular norm; however, in the following, we will instantiate to the ℓ_1 -norm.

2.2.5 Linear program (LP) formulation—We now start to combine all pieces to obtain the final optimization problem. Starting from Eq. (11), we add the graph regularizer from Eq. (12) to the mix.

Note that to avoid the trivial solution, we add additional constraints as discussed in Sect. 2.2.2. Here, we use $\Theta_{i,1} + \Theta_{i,3} = 1$ (i.e., $a + c = 1$).

$$\begin{aligned} \min_{\Theta_i, s_i, t_i} \quad & \sum_{i=1}^M \sum_{t=1}^{N_i} t_{t,i} + \sum_{i=1}^{M-1} |\Theta_i - \Theta_{i+1}| \\ \text{s.t.} \quad & \mathbf{x}_i^\top \Theta \leq s_i, \\ & -\mathbf{x}_i^\top \Theta \leq s_i, \\ & s_i - \varepsilon \leq t_i, \\ & 0 \leq t_i, \\ & \Theta_{i,1} + \Theta_{i,3} = 1. \end{aligned}$$

Again, using the fact that $|a| = \max(a, -a)$, we push the graph regularizer to the constraints:

$$\begin{aligned} \min_{\Theta_i, s_i, t_i} \quad & \sum_{i=1}^M \sum_{j=1}^{N_i} t_{i,j} + \sum_{i=1}^{M-1} \sum_{j=1}^D u_{i,j} \\ \text{s.t.} \quad & \Theta_{i,j} - \Theta_{i+1,j} \leq u_{i,j} \\ & -\Theta_{i,j} + \Theta_{i+1,j} \leq u_{i,j} \\ & \forall i \in [1, M-1], j \in [1, D], \\ & \mathbf{x}_i^\top \Theta \leq s_i, \\ & -\mathbf{x}_i^\top \Theta \leq s_i, \\ & s_i - \varepsilon \leq t_i, \\ & 0 \leq t_i, \\ & \Theta_{i,1} + \Theta_{i,3} = 1. \end{aligned}$$

The above problem consists of a linear objective and linear constraints and can therefore be solved with a linear program solver. We used the freely available GNU Linear Programming Toolkit to solve the above optimization problem (<http://www.gnu.org/software/glpk/glpk.html>).

3 Experiments

Robustness analysis

To evaluate the proposed robust loss in the context of parametric fits, we set up an experiment using synthetic data. We compare to the standard squared loss and to RANSAC [31], which is the standard approach for robust parametric fits. The idea of RANSAC is to iteratively fit a standard parametric model (in our case using the ℓ_2 -loss), whereas at each iteration a subset of points is chosen, such that outliers are discarded.

To investigate the robustness of the three mentioned approaches, we varied two sources of noise. First, we sampled $n = 200$ points $\mathbf{x}_1, \dots, \mathbf{x}_n$ from an ellipse parametrized by Θ_0 , where we added normally distributed noise to the x and y coordinates according to a zero-mean Gaussian with varying standard deviations. Second, we added uniformly distributed points to the training set (sampled in the interval $[-3, 3]$) to simulate contaminations, whereas the number of contaminating points C constitutes the second parameter controlling the noise level. For the noise parameters, we used the values $\sigma = \{0.0, 0.1, 0.2, \dots, 2.0\}$ and $C = \{0, 10, 20, \dots, 200\}$. Both sources of noise were jointly increased in 10 steps. We assessed the error by comparing the recovered parameter vectors Θ_{fit} to the true parameter Θ_0 (i.e., $error = \|\Theta_{fit} - \Theta_0\|$). Examples of the synthetic data at different noise levels, and the corresponding fits are shown in Fig. 4a-c. A comparison of the three approaches is shown in Fig. 3, where the error is shown as a function of the noise level.

We observe that, as to be expected, all methods recover the correct parameter vector if no noise is present. Further, as we increase the noise level, we observe that the robust loss consistently produces more accurate results than the ℓ_2 -loss. RANSAC oftentimes manages to obtain good solutions as well (even better solutions than for robust loss for a few cases), but has a higher variance, which reflects the fact that it is not guaranteed to obtain an optimal solution, but may get stuck in a local optimum. Furthermore, the RANSAC solution requires many repeated fits of the parametric model, leading to an increased computational burden compared to the other two methods that each solve exactly one convex optimization problem.

We conclude that for the setting we investigated, our method is comparable to (or better than) RANSAC, the standard approach for robust parametric fits. However, our method has lower variation between noise levels and a lower computational cost, due to its convexity.

Evaluation in practice

To evaluate the quality of our fits on real data, we compared them to manually curated segmentations obtained using the software of the microscope manufacturer. The 2D polygon finder program (SoftWorx, Applied Precision/GE Healthcare) was used to semi-automatically segment yeast nuclei labeled with Cut11-mCherry in 3D as previously described [9]. Images are 3D stacks of fission yeast nuclei expressing Cut11-mCherry, which localizes to the nuclear rim, and spindle assembly checkpoint proteins marked with GFP. Imaging and image analysis were performed as described in [9] with $n > 200$ cells. In brief, live cell imaging was conducted on a DeltaVision Core system (Applied Precision) equipped with a climate chamber (EMBL) using a 60 $\times$ /1.4 Plan Apo oil objective

(Olympus), and data were recorded with a CoolSnap CCD camera (Roper Scientific). Z-stacks of 4 μ m thickness were acquired for both mCherry and GFP fluorescence, with single planes spaced by 0.2 μ m. The imaged area spanned 256 $\times$ 256 pixels with 2 $\times$ 2 binning. Uneven illumination of the imaged area was corrected by flatfielding. All images were deconvolved using SoftWorx software. Individual nuclei were cut out and segmented using the graph-regularized fitting program or by manual segmentation setting a threshold value of 1.5 $\times$ average cytoplasmic signal. The resulting segmentation of the nucleus was projected to the GFP channel, and the sum of GFP intensity per sum of area was calculated. The results are shown in Fig. 5. We observe that our approach has an almost perfect correlation to the manually curated ground truth, while the existing microscope software shows a considerable deviation. Further, to compare the graph-regularized fitting software with other 3D segmentation software, the widely used Imaris 3D reconstruction software (Bitplane/Oxford Instruments) was chosen to segment nuclei. We compared nuclei showing heterogeneous signal intensity, which is one of the central challenges for segmentation tools but a fundamental property in cellular systems [32,33]. In brief, 3D tracking with Imaris was performed using the surpass mode and an outer shell was build with the surface option and the surface generation wizard. Smoothing was enabled, using a grain size of 0.430 μ m. To set an initial threshold, the background subtraction (local contrast) method was chosen and the diameter of the largest sphere that fits into the volume was set to 4 μ m. The threshold was then manually adjusted if the initially created surface only partially fitted the nuclear signal. Intensity and volume measurements were automatically performed by the software. Imaris performed less accurate (see Fig. 5). The correlation of the determined intensities between manual tracking and graph-regularized fitting was $r = 0.9662$, whereas Imaris performed substantially worse with $r = 0.8912$. This was due to: (a) setting only a single threshold value for the complete volume determination. Employing a single value cannot account for global intensity variations within the 3D stack, whereas the graph-regularized fitting tool sets a threshold for individual planes and is therefore independent of intensity fluctuations between individual planes, and (b) problems with tracking heterogeneous samples. Both the Imaris software and the 2D polygon finder (SoftWorx) detect cellular structures with homogeneous signals well but perform weakly if such structures show unequal signal distribution. For example, both programs cannot discriminate between an ‘inner’ and ‘outer’ region of cellular nuclei marked only at the nuclear rim. Therefore, regions within cellular structures are either falsely excluded or included, or incompletely tracked (see supplement). In contrast, the graph-regularized fitting software can account for heterogeneous signal intensities as it integrates prior knowledge about the structure of interest into the 3D fitting process.

4 Conclusion

We have presented GRED, a tool for parametric fitting of cell-like objects in fluorescence microscopy images. We have shown that in the context of unequal signal distribution, we can substantially improve upon existing software solutions. Combined with an easy-to-use user interface, our tool enables biologists to perform high-throughput quantitative experiments in fluorescence microscopy [9], if predefining the structure of the to-be-tracked cellular structure is advantageous.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We gratefully acknowledge core funding from the Sloan Kettering Institute (to G.R.), from the Ernst Schering foundation (to S.H.) and from the Max Planck Society (to G.R. and S.H.). Part of this work was done while C.W., S.H., P.D. and G.R. were at the Friedrich Miescher Laboratory of the Max Planck Society and while C.W. was at the Machine Learning Group at TU-Berlin.

References

1. Santella A, Du Z, Nowotschin S, Hadjantonakis A, Bao Z. A hybrid blob-slice model for accurate and efficient detection of fluorescence labeled nuclei in 3D. *BMC Bioinformatics*. 2010; 11(1):580–2010. [PubMed: 21114815]
2. Li G, Liu T, Tarokh A, Nie J, Guo L, Mara A, Holley S, Wong S. 3D cell nuclei segmentation based on gradient flow tracking. *BMC Cell Biol*. 2007; 8(1):40. [PubMed: 17784958]
3. Coelho LP, Shariff A, Murphy RF. Nuclear segmentation in microscope cell images: a hand-segmented dataset and comparison of algorithms. *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2009
4. Cremers D, Rousson M, Deriche R. A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *Int J Comput Vis*. 2007; 72(2):195–215.
5. Dufour A, Shinin V, Tajbakhsh S, et al. Segmenting and tracking fluorescent cells in dynamic 3-D microscopy with coupled active surfaces. *IEEE Trans Image Process*. 2005; 14(9):13961410.
6. Srinivasa G, Fickus MC, Guo Y, et al. Segmenting and tracking fluorescent cells in dynamic 3-D microscopy with coupled active surfaces. *IEEE Trans Image Process*. 2009; 18(8):1817–1829. [PubMed: 19380268]
7. Boykov Y, Funka-Lea G. Graph cuts and efficient ND image segmentation. *Int J Comput Vis*. 2006; 70(2):109–131.
8. Al-Kofahi Y, Lassoued W, Lee W, et al. Improved automatic detection and segmentation of cell nuclei in histopathology images. *IEEE Trans Biomed Eng*. 2010; 57(4):841–852. [PubMed: 19884070]
9. Heinrich S, Geissen E, Kamenz J, Trautmann S, Widmer C, Drewe P, Knop M, Radde N, Hasenauer J, Hauf S. Determinants of robustness in spindle assembly checkpoint signalling. *Nat Cell Biol*. 2013; 15(11):1328–1339. [PubMed: 24161933]
10. Pécot T, Singh S, Caserta E, Huang K, Machiraju R, Leone G. Non parametric cell nuclei segmentation based on a tracking over depth from 3D fluorescence confocal images. *9th IEEE International Symposium on Biomedical Imaging (ISBI)*. 2012:170–173.
11. Uzunbas MG, Soldea O, Unay D, Cetin M, Unal G, Ercil A, Ekin A. Coupled nonparametric shape and moment-based intershape pose priors for multiple basal ganglia structure segmentation. *IEEE Trans Med Imaging*. 2010; 29(12):1959–1978. [PubMed: 21118755]
12. Stegmaier J, Otte JC, Kobitski A, Bartschat A, Garcia A, Nienhaus GU, Strähle U, Mikut R. Fast segmentation of stained nuclei in terabyte-scale, time resolved 3D microscopy image stacks. *PLoS One*. 2014; 9(2):e90036. [PubMed: 24587204]
13. Svoboda D, Kozubek M, Stejskal S. Generation of digital phantoms of cell nuclei and simulation of image formation in 3D image cytometry. *Cytometry A*. 2009; 75(6):494–5099. [PubMed: 19291805]
14. Lou X, Kang M, Xenopoulos P, Muñoz-Descalzo S, Hadjantonakis AK. A rapid and efficient 2D/3D nuclear segmentation method for analysis of early mouse embryo and stem cell image data. *Stem Cell Reports*. 2014; 2(3):382–397. [PubMed: 24672759]
15. Mitchell IM. A toolbox of level set methods. *UBC Department of Computer Science Technical Report, TR-2007-11*. 2007

16. Wienert S, Heim D, Saeger K, Stenzinger A, Beil M, Hufnagl P, Dietel M, Denkert C, Klauschen F. Detection and segmentation of cell nuclei in virtual microscopy images: a minimum-model approach. *Sci Rep*. 2012; 2:503. [PubMed: 22787560]
17. Evgeniou T, Micchelli CA, Pontil M. Learning multiple tasks with kernel methods. *J Mach Learn Res*. 2005; 6(1):615–637.
18. Lou X, Koethe U, Wittbrodt J, Hamprecht FA. Learning to segment dense cell nuclei with shape prior. *CVPR*. 2012
19. Gonzalez, RC.; Woods, RE. *Digital Image Processing*. Prentice Hall; Upper Saddle River: 2008.
20. Lou X, Koethe U, Wittbrodt J, Hamprecht FA. Improved automatic detection and segmentation of cell nuclei in histopathology images. *IEEE on Computer Vision and Pattern Recognition (CVPR)*. 2012
21. Weisstein, EW. Ellipse. From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/Ellipse.html>
22. Rosin PL. Assessing error of fit functions for ellipses. *Graph Models Image Process*. 1996; 58(5): 494–502.
23. Fitzgibbon A, Pilu M, Fisher RB. Direct least square fitting of ellipses. *IEEE Trans Pattern Anal Mach Intell*. 1999; 21(5):476–480.
24. Newcomb S. A generalized theory of the combination of observations so as to obtain the best result. *Am J Math*. 1886; 8:343–366.
25. Huber PJ. The 1972 wald lecture robust statistics: a review. *Ann Math Stat*. 1972; 43(4):1041–1067.
26. Vapnik, V. Volume 8 of *Statistics for Engineering and Information Science*. Springer; Berlin: 1995. *The Nature of Statistical Learning Theory*.
27. Smola, AJ.; Schölkopf, B. Volume 64 of *Adaptive Computation and Machine Learning*. MIT Press; Cambridge: 2001. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*.
28. Smola AJ, Schölkopf B. A tutorial on support vector regression. *Stat Comput*. 2004; 14(3):199–222.
29. Boyd, S.; Vandenberghe, L. *Convex Optimization*. Cambridge University Press; Cambridge: 2004.
30. Widmer, C.; Kloft, M.; Görnitz, N.; Rätsch, G. *ECML 2012*. Springer; Berlin: 2012. Efficient training of graph-regularized multitask SVMs; p. 633-647.
31. Fischler MA, Bolles RC. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun ACM*. 1981; 24(6):381–395.
32. Elsasser WM. Outline of a theory of cellular heterogeneity. *Proceedings of the National Academy of Sciences*. 1984:5126–5129.
33. Li G, Liu T, Tarokh A, et al. 3D cell nuclei segmentation. *BMC Cell Biol*. 2007; 8(1):40. [PubMed: 17784958]

Fig. 1. Visualization of the preprocessing procedure, with individual processing steps as described in the main text. The images shown contain nuclei as studied and described in [9]. **a** Raw image, **b** Gaussian smoothing, **c** eigenvalues, **d** seeds, **e** labeled seeds, **f** boxes

Fig. 2. a Gives an example of a parametric fit for a volume. **b, c.** Two slices through the volume in **a** for which a good parametric fit is obtained in the face of missing data points

Fig. 3.

Comparison of the error of the squared and robust loss for different noise levels. The noise parameter controls both the Gaussian error from sampling points from the ground truth and the number of points from a uniform distribution. The error is measured as the squared norm of the difference between the parameter vector of the ground truth and the recovered parameter vector

Fig. 4.

Points sampled from the ground truth with Gaussian noise are shown in *black*, and the ones sampled from a uniform noise distribution are shown in *cyan*. The squared loss fit is shown in *red*, the ransac-fit is shown in *blue*, and the epsilon-insensitive robust fit is shown in *green*. **a** Low noise, **b** medium noise, **c** high noise (color figure online)

Fig. 5.

The *top row* shows results on real data comparing the 2D polygon fit by the microscope manufacturer to our approach. Taking into account manual post-processing, the total time taken for an experiment is shown in the *last column* (r denotes the Pearson correlation coefficient). The *bottom row* shows results on heterogeneous samples, which poses the biggest challenges to tracking software, comparing our tracking approach GRED to Imaris, a state-of-the-art software package for 3D-tracking