



Published in final edited form as:

Med Image Anal. 2016 August ; 32: 157–172. doi:10.1016/j.media.2016.03.007.

Automatic optimal filament segmentation with sub-pixel accuracy using generalized linear models and B-spline level-sets

Xun Xiao^{b,c,e}, Veikko F. Geyer^d, Hugo Bowne-Anderson^d, Jonathon Howard^d, and Ivo F. Sbalzarini^{a,b,c,*}

^aChair of Scientific Computing for Systems Biology, Faculty of Computer Science, TU Dresden, Dresden, Germany

^bMOSAIC Group, Center for Systems Biology Dresden (CSBD), Dresden, Germany

^cMax Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

^dDepartment of Molecular Biophysics and Biochemistry, Yale University, New Haven, CT, USA

^eNow at: European Research Center, Huawei Technologies, Munich, Germany

Abstract

Biological filaments, such as actin filaments, microtubules, and cilia, are often imaged using different light-microscopy techniques. Reconstructing the filament curve from the acquired images constitutes the filament segmentation problem. Since filaments have lower dimensionality than the image itself, there is an inherent trade-off between tracing the filament with sub-pixel accuracy and avoiding noise artifacts. Here, we present a globally optimal filament segmentation method based on B-spline vector level-sets and a generalized linear model for the pixel intensity statistics. We show that the resulting optimization problem is convex and can hence be solved with global optimality. We introduce a simple and efficient algorithm to compute such optimal filament segmentations, and provide an open-source implementation as an ImageJ/Fiji plugin. We further derive an information-theoretic lower bound on the filament segmentation error, quantifying how well an algorithm could possibly do given the information in the image. We show that our algorithm asymptotically reaches this bound in the spline coefficients. We validate our method in comprehensive benchmarks, compare with other methods, and show applications from fluorescence, phase-contrast, and dark-field microscopy.

Keywords

Filament segmentation; Level set; B-spline; Axoneme; Light microscopy; Microtubule

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

*Corresponding author. Tel.: +493512102525. ivos@mpi-cbg.de (I.F. Sbalzarini).

Author Contributions

Conceived the project: IFS and JH; developed the theory and the algorithm: XX and IFS; wrote the manuscript: IFS and XX; made the figures: XX; generated benchmark data: XX and VFG (FIESTA); preparation of biological samples, microscopy, and image acquisition: VFG; Edited and approved the final manuscript: all authors.

Introduction

Filamentous structures are ubiquitous in biology and are routinely imaged using different modalities. Examples range from cytoskeletal filaments, like microtubules (Ruhnow et al., 2011) and actin filaments (Gittes et al., 1993), to polymers (Graham et al., 2014), axonemes (Mukundan et al., 2014), sperm flagella (Rikmenspoel and Isles, 1985), nematodes (Geng et al., 2004; Ramot et al., 2008), and rodent whiskers (Clack et al., 2012). From the acquired images, one often aims to extract quantitative information about the filaments, such as their length (Ruhnow et al., 2011), curvature (Mukundan et al., 2014), bending (Gittes et al., 1993; Graham et al., 2014), or motion dynamics (Rikmenspoel and Isles, 1985). Doing so manually is prohibitive for large data volumes. Moreover, manual analysis introduces significant intra- and inter-rater variability and bias. Image segmentation techniques are available to automate the process. However, the problem of filament segmentation is not trivial because the thickness of the imaged filaments is often below the resolution limit of the microscope, and the signal-to-noise ratio (SNR) of the images is routinely low. While the former calls for filament localization with sub-pixel precision, the latter amplifies localization errors.

Addressing the problem of localization precision and error control in filament segmentation has been explored in various methods over the past decades. These methods can generally be classified into *region-segmentation* methods (e.g., Paul et al., 2013; Fuller et al., 2005), *curve-fitting* methods (e.g., Smith et al., 2010; Clack et al., 2012; Valdman et al., 2012; Xu et al., 2014), and combinations of the two (e.g., Ruhnow et al., 2011). Region-segmentation methods detect a connected region of pixels within which the filament lies, represented as either a *pixel mask* (Fuller et al., 2005; Goldstein et al., 2010; Paul et al., 2013), *particles* (Florin et al., 2005; Cardinale et al., 2012), or a *closed active contour* (Kass et al., 1988; Ronfard, 1994; Yezzi et al., 1997; Butenuth and Heipke, 2012; Zhang et al., 2012a; 2012b; Bernard et al., 2009). Usually, a closed active contour is implicitly represented as the zero-level set of a higher-dimensional function, which is called *level-sets methods* (Sethian, 1999).

While region-segmentation methods narrow down the localization of the filament, they do not segment a filament in the sense of a curve. Curve-fitting methods account for this by fitting a smooth open curve to the image in order to represent the filament (Wong et al., 1998; Sarry and Boire, 2001; Smith et al., 2010; Clack et al., 2012; Valdman et al., 2012; Xu et al., 2014), because characterizing a filament as a region of pixels is especially inappropriate if sub-pixel accuracy or smoothness are required. This is typically the case when quantifying filament polymerization kinetics by computing time derivatives of the filament length.

In both paradigms, the segmentation result can either be derived by solving an optimization problem (Paul et al., 2013; Fuller et al., 2005; Florin et al., 2005; Cardinale et al., 2009; Smith et al., 2010), or using filters (Saban et al., 2006; Danuser et al., 2000; Clack et al., 2012; Rigort et al., 2012). Formulating the task as an optimization problem, e.g. maximizing the Bayesian posterior of the segmentation to explain the image, provides principled ways of

including prior knowledge about the imaged filaments and the imaging modality in the form of the object and forward models.

Solving the resulting optimization problem, however, may be difficult. For discrete problems over pixel masks or filament pieces, dynamic programming (Clack et al., 2012) or graph-cuts (Kolmogorov and Zabini, 2004; Boykov and Kolmogorov, 2004; Boykov and Funkalea, 2006) can find the globally optimal solution. When using continuous filament representations, such as splines, active contours, or level-sets, however, optimization is mostly done locally, e.g., using gradient-descent (Zhang et al., 2012a; 2012b) or shape-gradient flow (Tsai et al., 2003; Law and Chung, 2009; Paul et al., 2013). For level-sets methods, variational approaches are popular, relaxing the requirement of previously having to know the number of the filaments to be segmented (Goldstein et al., 2010; Paul et al., 2013; Bernard et al., 2009). The main limitation of such local approaches is that image noise can cause the result to converge in a sub-optimal local minimum (Smith et al., 2010; Xu et al., 2014). In addition, curve-fitting methods usually require prior knowledge of the number of filaments present in the image (Wong et al., 1998; Sarry and Boire, 2001; Valdman et al., 2012).

Methods from both paradigms usually assume Gaussian image noise and the use of a certain type of microscopy, e.g., fluorescence microscopy. These assumptions may be inappropriate in some cases, for example when using phase-contrast microscopy, or in low-light conditions where the noise is Poissonian.

Despite these shortcomings, both paradigms also have unique advantages. Specifically, region-segmentation methods can localize multiple filaments automatically and in a computationally efficient way. When using a level-set segmentation method, the number of filaments present in an image does not need to be previously known or imposed. While curve-fitting methods require prior knowledge of the number of filaments to be fitted, they can properly characterize the geometry of filaments as open curves.

These observations motivate us to combine the two paradigms in a novel way. We propose a method that represents filaments as open curves, hence providing sub-pixel resolution in a geometric representation of the correct dimensionality, just as existing curve-fitting methods do. At the same time, however, rather than directly using a 1D curve representation, we use a level-set representation of these curves. This means that, unlike in curve-fitting methods, our method does not require the number of filaments to be previously known. It automatically handles multi-filament cases, as long as the filaments do not cross.

We show how to formulate the resulting optimization problem in a globally convex way (Paul et al., 2013) by measuring distances between segmentations and the image using the *Bregman divergence* (Banerjee et al., 2005; Bregman, 1967). We further derive an analytical expression for the gradient of the energy functional, enabling us to efficiently solve the convex problem by an easy-to-implement gradient-descent optimizer. The segmentation produced is the best possible one for the given model and image data. We also exploit the model-based framework to account for key physical properties of the microscope (i.e., image-formation model), which are rarely considered in other methods. This renders our

method fully automatic in the sense that it does not require prior manual segmentation or image pre-processing (e.g., inversion or smoothing).

Technically, this is made possible by representing the filaments as a vector level-set with level functions represented as B-spline surfaces (Bernard et al., 2009). The evolution of those surfaces is driven toward the optimal segmentation by a convex Bregman energy (Paul et al., 2013), which derives from a generalized linear model (GLM) (McCullagh, 1984; Nelder and Wedderburn, 1972) for the noise in the image (see (Paul et al., 2013) for details). This allows for any pixel-wise noise distribution from the exponential family, which includes the most common cases in microscopy, such as Gaussian, Poisson, and Bernoulli noise (Paul et al., 2013). This relaxes the assumption of Gaussian pixel noise, rendering the model more flexible.

In order to confirm the optimality of our solution, we derive an information-theoretic lower bound for the segmentation error in filament segmentation. This lower bound defines how well any unbiased algorithm can possibly perform given the information in the image. We find that the algorithm presented here asymptotically reaches the bound in the spline coefficients, showing that the bound is attainable.

In summary, our method can (1) automatically detect and optimally segment a previously unknown number of filaments as smooth open curves with sub-pixel resolution, and (2) be adapted to different types of image data from different microscopy modalities. The method is easy to implement, since it amounts to standard gradient descent over a convex function.

The paper is organized as follows: in Section 2, we present the optimization framework, mathematically formalizing the filament segmentation problem. The algorithm design and the theoretical lower bound for the segmentation accuracy are introduced in Sections 3 and 4, respectively. Results of comprehensive benchmarks and experiments are shown in Section 5, and we conclude this work in Section 6.

2. Problem formulation

Before providing the details of our method, we mathematically formulate the filament segmentation problem. The first thing we need is a mathematical representation of the filaments.

We represent filaments as vector level-sets, because by using a single level function the level set is necessarily a closed curve. In order to represent open curves, we use a vector level-set with two components. The zero level set of the first component ϕ describes a closed curve that contains the filament (black circle in Fig. 1a). The part of that closed curve where the second component ψ is positive then is the actual open filament (black arc in Fig. 1b). The vector level function is represented as a B-spline surface. This geometric model can represent an arbitrary number of non-crossing and non-overlapping filaments. Filaments are overlapping as soon as they are too close to each other for the microscope optics to resolve them as two.

Based on this filament representation, it is possible to predict the image that one *expects* to see when imaging those filaments with a particular microscope. This *forward model* hence simulates the generative process of image formation. The so-predicted model image is then compared with the actually observed image data. The difference between the two (measured in a suitable metric) then drives the evolution of the filament representation so as to minimize this difference. A regularizer is used to avoid over-fitting. The data-fitting term and the regularizer together form the *energy functional* that is to be minimized in order to find the best possible segmentation. The variables of the optimization problem are the coefficients of the B-spline vector level-set surfaces. They evolve such that the represented filaments lead to an expected image that is as close as possible to the actually observed image.

2.1. Image model

The forward model assumes that each filament has a uniform intensity along its length, and that the background is uniformly homogeneous. If these assumptions are not met, the algorithm will still work, but the result is no longer guaranteed to be globally optimal, as we show below. The (unknown) intensities of the filament and background are thus represented by a 2-vector $\boldsymbol{\beta} = [\beta_1, \beta_2]^T$. Let $H_f(\mathbf{x})$ be an indicator function such that $H_f(\mathbf{x}) = 1$ means that the point \mathbf{x} lies on a filament, $H_f(\mathbf{x}) = 0$ is background. Hence, the intensity distribution in the model sample *before* imaging is

$$\mu(\mathbf{x}, \boldsymbol{\beta}) = \beta_1 H_f(\mathbf{x}) + \beta_2 (1 - H_f(\mathbf{x})). \quad (1)$$

The indicator function $H_f(\mathbf{x})$ is represented by the vector level set. In order to represent non-intersecting filaments, we use two level functions $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$. Open-curve filaments are then represented by

$$\Gamma(\mathbf{x}) = \{\mathbf{x} : \phi(\mathbf{x}) = 0 \cap \psi(\mathbf{x}) > 0, \mathbf{x} \in \Omega_{\mathcal{I}}\}, \quad (2)$$

where $\Omega_{\mathcal{I}}$ is the image domain. The level set $\phi(\cdot) = 0$ describes a closed contour, while the condition $\psi(\cdot) > 0$ cuts the closed contour to an open curve since only a subset of \mathbf{x} on $\phi(\mathbf{x}) = 0$ can yield $\psi(\mathbf{x}) > 0$. The above indicator function $H_f(\mathbf{x})$ is then given by

$$H_f(\mathbf{x}) = \delta(\phi(\mathbf{x})) H(\psi(\mathbf{x})), \quad (3)$$

where $\delta(\cdot)$ is a Dirac and $H(\cdot)$ a Heaviside distribution. This is equivalent to condition (2) because if $\phi(\cdot) = 0$ and $\psi(\cdot) > 0$, $H_f(\cdot) = 1$, meaning that \mathbf{x} belongs to a filament. Otherwise $H_f(\cdot) = 0$, meaning the point belongs to the background.

As an illustrating example explaining the vector level-set representation, in Fig. 1a, a closed contour (shown as a black circle) is represented by the first level set $\phi(\mathbf{x}) = 0$. This closed contour $\phi(\mathbf{x}) = 0$ is cut so as to satisfy $\psi(\mathbf{x}) > 0$, which directly interprets condition (2). In Fig. 1b, the resulting indicator function $H_f(\mathbf{x})$ in Eq. (3) is shown in the 2D image domain.

The image of $H_f(\mathbf{x}) = 1$ (red curve) presents an open curve, while $H_f(\mathbf{x}) = 0$ is the dark background. This forms a solution space for filament segmentation, since different indicator functions $H_f(\cdot)$, represented by the two level functions, correspond to different segmentation results.

The resulting expected model image is then formed by running the hypothetical scenery μ through the image-formation forward model. This predicts what the image would look like in expectation when seen through that particular microscope.

2.2. Representation of level functions

The two level functions $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$ can be represented in a number of ways. A common choice is to use signed-distance functions. This, however, requires re-initialization (or penalization) during contour evolution in order to ensure that the signed-distance property is not lost (Sethian, 1999). In our application, we do not require distance information in the filament domain. Following earlier works, we hence simplify the representation by using B-spline functions (Maeland, 1988; Farin, 1996; Bernard et al., 2009) for both level-set functions. Representing the level functions as B-spline surfaces moreover provides an analytical form (polynomial) of each segmented filament, which provides direct access to geometric features such as gradients, normals, and curvature (Bernard et al., 2009). Using $\phi(\mathbf{x})$ as an example ($\psi(\mathbf{x})$ is analogous), the function is

$$\phi(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}_+^2} c_\phi[\mathbf{k}] b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right), \quad (4)$$

where \mathbb{Z}_+^2 is a 2D mesh grid of positive integers that has the same size as the image domain

$\Omega_{\mathcal{Q}}$, \mathbf{k} is an integer 2-index, $c_\phi[\mathbf{k}]$ is the spline coefficient set, and $b_n(\frac{\mathbf{x}}{h+1} - \mathbf{k})$ is an n th-order basis function with the size h controlling the spacing between two consecutive nodes.

The dimensionality of the resulting optimization problem is given by the number of spline coefficients used to represent the level functions. This number is controlled by h . We illustrate B-spline surfaces of different sizes h in Fig. 2. For $h = 0$, the B-spline function has one coefficient (node) per image pixel. For $h = 1$, one pixel is left empty between any two consecutive nodes, thus quartering the number of coefficients in 2D. Thus, increasing h reduces the resolution of the spline, but also reduces the dimensionality of the resulting optimization problem.

On the one hand, smaller h allow the segmented filaments to trace smaller undulations. On the other hand, smaller h increases the computational cost and renders the segmentation more sensitive to noise. We investigate the trade-off of the scaling size h in simulation results in Section 5.3, and provide good parameter choices.

2.3. Energy functional

The energy functional is a criterion that evaluates the quality of a given segmentation with respect to the image data. If the model image shows the correct filaments, the energy should be minimal. Following the standard approach, we use an energy functional composed of two parts: a *data-fitting term* that measures the distance between the model image and the data image, and a *regularization term* that penalizes over-fitting.

2.3.1. Data-fitting term—The data-fitting term measures how closely the image one expects to see under the current segmentation hypothesis matches the observed data image. This includes a simulation of the image-formation process in the given microscope, and a pixel noise model. We model image formation as a convolution with a kernel K , which is known from optics or calibration measurements. The expected model image hence is $K * \mu(\mathbf{x}, \boldsymbol{\beta})$. This simulates how the imaging system maps the scene to an image. In fluorescence microscopy, K is the point-spread function (PSF). For phase-contrast microscopy, K is given by a non-linear imaging model that involves Bessel functions, as previously described (Yin et al., 2010).

The second component is the noise process, providing a likelihood to compare the model image to the image data. Assume that the intensity value $u(\mathbf{x})$ at location \mathbf{x} is a realization from a stochastic process with probability density function (p.d.f.) $p(u(\mathbf{x}) | \boldsymbol{\theta})$ and unknown parameters $\boldsymbol{\theta}$. In our problem, the unknown parameters are the mean intensities β_1 and β_2 of the fore.p.f.f and background, respectively. The likelihood of a segmentation given in the image data is $\ell(\mu(\mathbf{x}, \boldsymbol{\beta}) | u(\mathbf{x}))$. This likelihood increases as the model (i.e., segmentation) $\mu(\cdot, \cdot)$ converges to the correct solution.

To estimate the unknown parameters β_1 and β_2 , we use a *linear predictor* for the parameters of any p.d.f. from the exponential family (EF) of distributions. A non-linear *link function* correlates the linear predictor to the unknown parameters. This is the standard GLM framework in statistics.

We give a brief example to illustrate the GLM idea. We are given a set of pixels sampled from a probability distribution with unknown mean. We can estimate the mean as the average of the samples. In this case, the linear predictor of the GLM is an averaging operation, and the link function is the identity function because we let the mean directly equal the average of the samples. In the general case, we can specify an arbitrary link function. If we for example let the log of the parameters equal the sample average, we enforce positivity of estimated intensity values. Hence, a GLM generalizes the classical estimation of the mean by the average, allowing additional flexibility.

In our problem, the linear predictor is the model $\mu(\mathbf{x}, \boldsymbol{\beta})$. We further use the link function $g(\mathbb{E}(u(\mathbf{x}))) = K * \mu(\mathbf{x}, \boldsymbol{\beta})$, or $\mathbb{E}(u(\mathbf{x})) = g^{-1}(K * \mu(\mathbf{x}, \boldsymbol{\beta}))$ where $g^{-1}(\cdot)$ is the *inverse link function*.

¹ The likelihood of the model can then be written as

¹If $g^{-1}(\cdot)$ is the identity function, i.e., $I(\cdot)$, this reduces to the classical method of estimating the mean as the average of the samples.

$$\ell(\mu(\mathbf{x}, \boldsymbol{\beta})|u(\mathbf{x})) = p(u(\mathbf{x})|g^{-1}(K*\mu(\mathbf{x}, \boldsymbol{\beta}))). \quad (5)$$

Though it would be possible to directly maximize the likelihood in Eq. (5), likelihoods of some p.d.f.'s from the EF are complicated. For example, the p.d.f. of a Poisson distribution contains factorials and is a discrete function. More importantly, some p.d.f.s are not globally convex and, as a result, global optimality cannot be guaranteed. Hence, a better data-fitting term is needed.

We address these problems by replacing the likelihood with the Bregman divergence (Bregman, 1967) induced by that specific p.d.f. from the EF. The Bregman divergence generalizes the familiar squared Euclidean distance to a class of distances that all share similar properties. The most important property for our purpose is that a Bregman divergence is globally convex. This guarantees global optimality of the solution.

The form of the Bregman divergence varies according to the p.d.f. chosen as the noise model. For example, assuming $p(\cdot)$ to be Gaussian, the Bregman divergence is the squared Euclidean distance $\|u(\mathbf{x}) - K*\mu(\mathbf{x}, \boldsymbol{\beta})\|_2^2$. If $p(\cdot)$ is a Poisson distribution, the corresponding

Bregman divergence is the relative entropy $u(\mathbf{x}) \log \frac{u(\mathbf{x})}{K*\mu(\mathbf{x}, \boldsymbol{\beta})} + K*\mu(\mathbf{x}, \boldsymbol{\beta}) - u(\mathbf{x})$. The Bregman divergences associated with any p.d.f. from the EF are given in Table 1 of Paul et al. (2013). We denote by $\mathbf{B}_{p^*}(\cdot|\cdot)$ the Bregman divergence induced by the noise model p.d.f. p^* . Integrating over the image domain $\Omega_{\mathcal{I}}$, the final data-fitting energy term then reads

$$\mathcal{E}_d(\mathbf{x}, \boldsymbol{\beta}) = \int_{\Omega_{\mathcal{I}}} \mathbf{B}_{p^*}(u(\mathbf{x})|g^{-1}(K*\mu(\mathbf{x}, \boldsymbol{\beta}))) \, d\mathbf{x}. \quad (6)$$

2.3.2. Regularization term—The regularizer is required to prevent over-fitting the noise in the image. It is a necessary component of any segmentation energy function and plays the role of a Bayesian prior in the estimation problem. We here use the popular total variation (TV) regularizer:

$$\mathcal{E}_r(\mathbf{x}, \boldsymbol{\beta}) = \int_{\Omega_{\mathcal{I}}} \|\nabla \mu(\mathbf{x}, \boldsymbol{\beta})\|_2 \, d\mathbf{x}, \quad (7)$$

where $\nabla \mu$ is the gradient of the model image.

The total energy functional is then given by

$$\mathcal{E}_T(\mathbf{x}, \boldsymbol{\beta}) = \mathcal{E}_d(\mathbf{x}, \boldsymbol{\beta}) + \lambda \mathcal{E}_r(\mathbf{x}, \boldsymbol{\beta}), \quad (8)$$

where the scalar regularization coefficient $\lambda > 0$ weights the regularizer term. This function is convex, since both \mathcal{E}_d and \mathcal{E}_r are convex, multiplying a convex function by a positive constant leaves it convex, and the sum of two convex functions is convex. Increasing λ increases the penalization of small undulations of the segmented filaments. Both λ and the scaling size h of the B-spline functions affect the robustness against noise and the segmentation precision of the method.

Now we can state the filament segmentation problem more formally: The goal is to design an algorithm that finds the optimal coefficients $c_\phi^*[\mathbf{k}]$ and $c_\psi^*[\mathbf{k}]$ such that the energy in Eq. (8) is minimized, hence

$$(c_\phi^*[\mathbf{k}], c_\psi^*[\mathbf{k}]) = \arg \min_{c_\phi[\mathbf{k}], c_\psi[\mathbf{k}]} \mathcal{E}_T(\mathbf{x}, \boldsymbol{\beta}). \quad (9)$$

Since the spline coefficients are normalized to $[-1, 1]$ (see Section 3.2.2), the domain of the optimization variables is the hypercube $[-1, 1]^{|\mathbf{k}|}$, which is a convex set. Minimizing a convex function over a convex domain constitutes a globally convex optimization problem.

The result will hence be a globally optimal filament segmentation $H_f^*(x)$.

3. Algorithm design

We provide an algorithm that globally optimizes the coefficients $c_\phi[\mathbf{k}]$ and $c_\psi[\mathbf{k}]$ of the two B-spline level-set functions and estimates the mean intensities $\boldsymbol{\beta}$ of the fore- and background regions. We use alternating minimization (Gunawardana and Byrne, 2005) over the energy functional in Eq. (8). We first introduce the method to estimate the intensities $\boldsymbol{\beta}$, followed by the procedure used to optimize the level-set coefficients. Since both the level-set optimization problem and the intensity estimation problem are convex, the alternating minimization scheme is guaranteed to converge (Gunawardana and Byrne, 2005). Note that the two problems are likely not jointly convex, but separately convex, which is enough for convergence (Gunawardana and Byrne, 2005). If needed, a jointly convex formulation could possibly be derived by extending recent works (Brown et al., 2012). A post-processing step then extracts parametric polynomial curves for all filaments.

3.1. Photometric estimation

Photometric estimation aims to estimate the mean intensity vector $\boldsymbol{\beta}$ for fixed $H_f(x)$. For a perfect imaging process (i.e., if $K = \mathbf{1}$), estimating $\boldsymbol{\beta}$ is trivial, as we can directly use the average intensity of a region as its mean. With an imaging-distortion kernel K acting on $\mu(\cdot)$, however, the optimal value for $\boldsymbol{\beta}$ is not a simple average (Paul et al., 2013).

In this case, we minimize the data-fitting term to find the most likely value of $\boldsymbol{\beta}$. This convex problem is efficiently solved using the *Fisher scoring algorithm* (McCullagh, 1984; Nelder and Wedderburn, 1972). We use the freely available solver from Paul et al. (2013), and refer to it as the *region statistics solver*. This solver is called multiple times in the alternating minimization framework. After each call, $H_f(x)$ is re-optimized using the new photometries,

iterating until convergence (which is guaranteed for a separately convex problem). For this geometric optimization step, the spline coefficients are optimized using gradient descent.

3.2. Geometric optimization

The gradient-descent procedure iteratively evolves the coefficients for the two B-spline level-set functions $\phi(\cdot)$ and $\psi(\cdot)$, going from iteration i to $i+1$ as

$$c_{\phi}^{(i+1)}(\mathbf{k}) = c_{\phi}^{(i)}(\mathbf{k}) - \alpha \mathcal{G}_{\phi}^{(i)}(\mathbf{k})$$

$$c_{\psi}^{(i+1)}(\mathbf{k}) = c_{\psi}^{(i)}(\mathbf{k}) - \alpha \mathcal{G}_{\psi}^{(i)}(\mathbf{k}), \quad (10)$$

where α is the step-size (automatically adapted as described below), and $\mathcal{G}_{\phi}^{(i)}(\mathbf{k})$ and $\mathcal{G}_{\psi}^{(i)}(\mathbf{k})$ are the energy gradients for the two level-set functions, respectively. Instead of numerically approximating the gradients, we directly compute and use the regularized analytical gradients.

3.2.1. Analytical gradients—The gradient of the energy cannot directly be computed because of the Dirac and Heaviside distributions in H_{ε} . In order to get a continuously differentiable energy, we hence regularize them as previously described (Bernard et al., 2009):

$$H_{\varepsilon}(\mathbf{x}) = \frac{1}{1 + e^{-\varepsilon_H \mathbf{x}}}$$

$$\delta_{\varepsilon}(\mathbf{x}) = \frac{1}{\pi} (\varepsilon_{\delta} / (\mathbf{x}^2 + \varepsilon_{\delta}^2)). \quad (11)$$

Since the level function is normalized to $[-1, 1]$, we fix the regularization parameters to $\varepsilon_H = 20$ and $\varepsilon_{\delta} = 0.1$ (Bernard et al., 2009).

The parametric representation of the level-set functions as B-splines allows us to directly compute the analytical gradients of the energy with respect to the spline coefficients using the chain rule for differentiation. For the data-fitting term $\mathcal{E}_d(\cdot)$, we have

$$\mathcal{G}_{\phi}^d(\mathbf{k}) = \frac{\partial \mathcal{E}_d(\mathbf{x}, \boldsymbol{\beta})}{\partial c_{\phi}[\mathbf{k}]} = (\beta_1 - \beta_2) \int_{\Omega_{\mathcal{I}}} \mathbf{B}'_{p^*}(\cdot \| \cdot) K * \delta'_{\varepsilon}(\phi) H_{\varepsilon}(\psi) b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right) d\mathbf{x},$$

$$\mathcal{G}_\psi^d(\mathbf{k}) = \frac{\partial \mathcal{E}_d(\mathbf{x}, \boldsymbol{\beta})}{\partial c_\psi[\mathbf{k}]} = (\beta_1 - \beta_2) \int_{\Omega, \mathcal{J}} \mathbf{B}'_{p^*}(\cdot \| \cdot) K * \delta_\varepsilon(\phi) \delta_\varepsilon(\psi) b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right) d\mathbf{x}. \quad (12)$$

For compactness of notation, we omit the arguments of the Bregman divergence, hence read $\mathbf{B}'_{p^*}(\cdot \| \cdot)$ as $\mathbf{B}'_{p^*}(\mathbf{u}(\mathbf{x}) \| \mathbf{g}^{-1}(K * \mu(\mathbf{x}, \boldsymbol{\beta})))$ and \mathbf{B}'_{p^*} is the derivative with respect to the B-spline coefficients in μ . For the regularizer term $\mathcal{E}_r(\cdot)$, we have

$$\mathcal{G}_\phi^r(\mathbf{k}) = \frac{\partial \mathcal{E}_r(\mathbf{x}, \boldsymbol{\beta})}{\partial c_\phi[\mathbf{k}]} = \int_{\Omega, \mathcal{J}} K * \|\nabla \delta'_\varepsilon(\phi) H_\varepsilon(\psi)\|_2 b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right) d\mathbf{x},$$

$$\mathcal{G}_\psi^r(\mathbf{k}) = \frac{\partial \mathcal{E}_r(\mathbf{x}, \boldsymbol{\beta})}{\partial c_\psi[\mathbf{k}]} = \int_{\Omega, \mathcal{J}} K * \|\nabla \delta'_\varepsilon(\phi) \delta_\varepsilon(\psi)\|_2 b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right) d\mathbf{x}. \quad (13)$$

The prime ($'$) means derivative with respect to the only variable of that function. The total gradient is the weighted sum of these two parts:

$$\mathcal{G}_\phi(\mathbf{k}) = \mathcal{G}_\phi^d(\mathbf{k}) + \lambda \mathcal{G}_\phi^r(\mathbf{k}),$$

$$\mathcal{G}_\psi(\mathbf{k}) = \mathcal{G}_\psi^d(\mathbf{k}) + \lambda \mathcal{G}_\psi^r(\mathbf{k}). \quad (14)$$

These analytical gradients can be efficiently evaluated numerically. Specifically, Eq. (14) can be rearranged as follows (taking $\mathcal{G}_\phi(\mathbf{k})$ as an example):

$$\mathcal{G}_\phi(\mathbf{k}) = \int_{\Omega, \mathcal{J}} K * \underbrace{\left((\beta_1 - \beta_2) \mathbf{B}'_{p^*}(\cdot \| \cdot) \delta'_\varepsilon(\phi) H_\varepsilon(\psi) + \lambda \|\nabla \delta'_\varepsilon(\phi) H_\varepsilon(\psi)\|_2 \right)}_{\text{a function } \omega_\phi(\mathbf{x})} \cdot b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right) d\mathbf{x}.$$

(15)

It can hence be interpreted as a convolution of the function $\omega_\phi(\mathbf{x})$ with the spline basis $b_n(\cdot)$. This simplifies the gradient calculation for each coefficient in $c_\phi[\mathbf{k}]$ to

$$\mathcal{G}_\phi(\mathbf{k}) = \omega_\phi(\mathbf{x}) * b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right),$$

$$\mathcal{G}_\psi(\mathbf{k}) = \omega_\psi(\mathbf{x}) * b_n \left(\frac{\mathbf{x}}{h+1} - \mathbf{k} \right). \quad (16)$$

3.2.2. Implementation details—A first technical issue is that any gradient-descent strategy is sensitive to the value of the step size α . In our implementation, we dynamically adapt the step size so as to guarantee that the energy decreases in each step. For this, we do a simple line search to find the optimal step size in each iteration. The search uses recursive bisection (binary search) starting from an initial value $\alpha_0 = 1$ (recall that the coefficients are normalized to $[-1, 1]$). By construction, the resulting α is guaranteed to decrease the energy.

A second technical issue is the presence of strong local distortions that may be introduced into the level functions by the energy minimization process. The level-set functions may in principle develop arbitrarily steep or flat gradients, challenging the numerical stability of the algorithm. Traditional implementations of level-set methods address this problem by periodically re-initializing the level function to a signed-distance function from the zero level-set. This, however, comes with two important drawbacks: extra computational costs and insufficient topological flexibility (Tsai and Osher, 2003). In our case, we do not require a signed-distance property, but simply wish to bound the gradient magnitudes for numerical stability. Thanks to the parametric B-spline representation of the level functions, this can be done by simple coefficient renormalization. This idea was originally introduced in Gelas et al. (2007). We hence normalize the coefficients $c_\phi(\cdot)$ and $c_\psi(\cdot)$ in each iteration by dividing with the ℓ_∞ -norm over all coefficients after each gradient-descent iteration:

$$c_\phi^{(i)}(\mathbf{k}) \leftarrow \frac{c_\phi^{(i)}(\mathbf{k})}{\|c_\phi^{(i)}(\mathbf{k})\|_\infty},$$

$$c_\psi^{(i)}(\mathbf{k}) \leftarrow \frac{c_\psi^{(i)}(\mathbf{k})}{\|c_\psi^{(i)}(\mathbf{k})\|_\infty}. \quad (17)$$

After normalization, the norm of the gradient of the level-set function is bounded, as proven in Bernard et al. (2009). This normalization has two main advantages: the first is its modest computational cost, the second is that it does not counteract the creation of new zero-level components, thereby conserving topological flexibility.

The complete optimization algorithm is summarized in Algorithm 1. The algorithm terminates as soon as the energy improvement achieved in the last iteration is less than a fixed tolerance of $\xi = 10^{-3}$, indicating that the algorithm has converged at the global optimum. The output of Algorithm 1 are the two optimal B-spline level functions $\phi^*(\cdot)$ and $\psi^*(\cdot)$. The following post-processing stage then extracts all zero-level curves from these level functions and stores them as polynomials.

Algorithm 1

Optimization Procedure

Input: Image $u(x)$, Noise type $p(\cdot)$, imaging kernel K

Output: Optimal $c_\phi^*[\mathbf{k}]$ and $c_\psi^*[\mathbf{k}]$

- 1: $i \leftarrow 0$, initialize $c_\phi^{(i)}(\mathbf{k})$ and $c_\psi^{(i)}(\mathbf{k})$ and $\mathcal{E}_T^{(i)}$;
- 2: **while** (tol > ξ) **do**
- 3: $\alpha \leftarrow \alpha_0$;
- 4: Line search for α ;
- 5: $c_\phi^{(i+1)}[\mathbf{k}] \leftarrow c_\phi^{(i)}[\mathbf{k}] - \alpha \mathcal{G}_\phi^{(i)}$;
- 6: $c_\psi^{(i+1)}[\mathbf{k}] \leftarrow c_\psi^{(i)}[\mathbf{k}] - \alpha \mathcal{G}_\psi^{(i)}$;
- 7: calculating $\mathcal{E}_T^{(i+1)}$ (call *Region Statistics Solver*);
- 8: $c_\phi^{(i+1)}[\mathbf{k}] \leftarrow \frac{c_\phi^{(i+1)}[\mathbf{k}]}{\|c_\phi^{(i+1)}[\mathbf{k}]\|_\infty}$, $c_\psi^{(i+1)}[\mathbf{k}] \leftarrow \frac{c_\psi^{(i+1)}[\mathbf{k}]}{\|c_\psi^{(i+1)}[\mathbf{k}]\|_\infty}$;
- 9: update tol = $\mathcal{E}_T^{(i+1)} - \mathcal{E}_T^{(i)}$, $i \leftarrow i + 1$;
- 10: **end while**

3.3. Post-processing

The optimal level functions $\phi^*(\cdot)$ and $\psi^*(\cdot)$ jointly encode all filaments in a single optimal indicator function $H_f^*(x)$. Due to the regularization in Eq. (11), this function is smooth and takes values in the interval $[0, 1]$, rather than being binary. The post-processing stage serves a double purpose: it determines the total number of filaments detected and then represents each filament as a parametric polynomial.

The total number of filaments is determined by thresholding the smooth $H_f^*(x)$ followed by extracting the foreground regions according to their topological connectivity. Representing each filament by a polynomial starts from extracting all local maxima from the original $H_f^*(x)$.

3.3.1. Filament splitting—In order to split the global vector level function into individual filaments, we threshold the smooth $H_f^*(x)$ according to the optimal energy value \mathcal{E}_T^* achieved by Algorithm 1. For this, $H_f^*(x)$ is thresholded with an arbitrary initial threshold ξ . After thresholding, we calculate the resulting energy \mathcal{E}_T^ξ . Performing a line search, we determine the threshold ξ^* that leads to an energy closest to \mathcal{E}_T^* . Since the set of binary

functions is a subset of the set of functions over the interval $[0, 1]$, this threshold is optimal. Thus, the optimally thresholded binary indicator function is $\bar{H}_f(\mathbf{x})$.

All points $s_i(\mathbf{x})$ with $\bar{H}_f(\mathbf{x}) = 1$ that are connected by reachable paths lie on the same filament. We hence identify regions of disjoint filaments (not the filament themselves!) as connected components in a 10-fold oversampled image of $\bar{H}_f(\mathbf{x})$ (i.e., 100 grid points per pixel) using the MATLAB routine `bwconncomp`.

3.3.2. Filament extraction—After identifying the region of each $s_i(\mathbf{x})$, the body of the filament is identified by all local maxima of the original $H_f^*(\mathbf{x})$ in that region. These are, to sub-pixel accuracy, the points of highest localization probability of the filament. Let $\tilde{s}_i(\mathbf{x})$ denote the so-determined finite set of tip and body points of filament i .

In order to provide a parametric description of each filament, we fit $\tilde{s}_i(\mathbf{x})$ with a B-spline, which is then converted to a single polynomial. While this step is not strictly necessary, it facilitates later access to geometric features of the filaments, such as normals and curvatures. Fitting $\tilde{s}_i(\mathbf{x})$ does not degrade the accuracy of the solution because the fitted positions already minimize the energy \mathcal{E}_T . This problem has been studied before, and we solve it using cubic smoothing splines as, e.g., implemented in the MATLAB routine `csaps` to solve the fitting problem:

$$\min_{\mathbf{c}_i(\mathbf{x})} \int_{\mathbf{x} \in \tilde{s}_i(\mathbf{x})} p \|\mathbf{c}_i(\mathbf{x}) - \tilde{s}_i(\mathbf{x})\|_2 d\mathbf{x} + (1-p) \int_0^1 \|\partial \mathbf{c}_i(t)\|_2 dt, \quad (18)$$

which yields for each filament i a polynomial $\mathbf{c}_i(\cdot)$ that minimizes the distance to the point set $\tilde{s}_i(\mathbf{x})$, penalizing the total variation of the curve. The coefficient p balances data-fitting and curve length, and $t \in [0, 1]$ is the curve parameter.

The pseudo code of the entire post-processing step is summarized in Algorithm 2.

Algorithm 2

Post-Processing

Input: optimal indicator function $H_f^*(\mathbf{x})$ with energy \mathcal{E}_T^*
Output: parametric curve set $\{\mathbf{c}_i(\mathbf{x})\}$
% filament splitting

- 1: $\bar{H}_f^*(\mathbf{x}) \leftarrow \arg \min_{\xi \in [0,1]} \|\mathcal{E}_T^\xi - \mathcal{E}_T^*\|_2$;
- 2: split $\{\mathbf{x} : \mathbf{x} = 1\}$ into subsets $\mathcal{S} = \{s_i(\mathbf{x})\}$;
% filament extraction
- 3: $\tilde{s}_i(\mathbf{x}) \leftarrow$ local maxima in $H_f^*(\mathbf{x})$ along $s_i(\mathbf{x})$
- 4: **for** $\forall \tilde{s}_i(\mathbf{x}) \in \mathcal{S}$ **do**
- 5: find $\mathbf{c}_i(\mathbf{x})$ by solving Eq. (18);

6: end for

3.4. Summary

We use two B-spline level sets to represent filament objects without restricting their topology. This is neither an open parametric curve nor a pixel mask. This representation enables our method to detect and segment previously unknown numbers of filaments, while preserving sub-pixel curve resolution for each filament. Additionally, the B-spline formulation enables analytically computing the energy gradients for global optimization, and extracting parametric polynomials for all detected filaments.

Our proposed algorithm optimizes the coefficients of the two level functions using gradient information. Each iteration consists of three steps: the gradient-descent procedure, the photometric estimation, and the coefficient normalization. A post-processing step at the end then extracts the parametric polynomial descriptions of all detected filaments. Since all problems are convex, the final result is globally optimal with respect to the imaging model postulated.

4. Theoretical error bound

A natural question to ask is how closely this globally optimal reconstruction represents the true filaments that were imaged. The imaging process and the pixel noise irreversibly destroy image information. In addition, the geometric representation of the filaments as B-spline level sets, as well as the model of the transfer function of the microscope, may not be appropriate. Thus we ask the question: “What is the theoretical bound of the localization error of the filament, considering the information contained in an image, and what is the expected performance of an optimal and unbiased algorithm?”

This is the question about a theoretical lower bound on the segmentation error. For point sources, this bound is known: given M photons that are all emitted from the same point source and imaged with a Gaussian PSF of standard deviation σ , the location of the point source can only be estimated with an error of at least σ / \sqrt{M} (Ober et al., 2004). This is a direct consequence of the Central Limit Theorem. For non-Gaussian PSFs and more complex noise models, similar bounds have also been derived (Chao et al., 2013; Ober et al., 2015). No unbiased algorithm could possibly break them.

Such a theoretical lower bound is not known for filaments, where the situation is considerably more complex than for point sources. The photon sources along a filament are correlated through the filament geometry, which is unknown. A theoretical bound can hence only be derived by assuming something about the smoothness of the filaments, that is the function space in which they live.

Here, we study this problem and provide a theoretical lower bound on the filament segmentation error. First, we consider the general case, where we do not assume anything about how the filaments are represented in the computer. This bound is valid for any unbiased filament segmentation algorithm. Then, we derive the lower bound for the special

case where the filaments are represented by vector B-spline level sets, as in our framework. This characterizes the capability of any vector level-set algorithm to segment filaments and leads to a bound that is actually computable.

Since we segment filaments, the lower bound is on the error of the filament geometry. A similar lower bound for the estimated filament intensity, or even a joint lower bound for geometry plus intensity, could possibly be derived in a similar fashion.

4.1. The general case

Without loss of generality, we only consider the case where a single filament is represented in the image. Since our algorithm assumes that filaments never cross, the multi-filament case can always be reduced to the single-filament case by cropping or decomposing the image accordingly.

Let a curve $\gamma(s)$ denote the true imaged filament, where $s \in [0, 1]$ is the curve parameter. We assume $\gamma(s) \in C^2$, i.e., continuous and at least twice differentiable. A segmentation algorithm reconstructs the filament as $\hat{\gamma}(s)$ from a single image $u(\mathbf{x})$ that represents the intensity values in a 2D image domain $\Omega_{\mathcal{I}}$.

The posterior p.d.f. of $\gamma(s)$ in a Bayesian sense is

$$p(u(\mathbf{x}), \gamma(s)) \propto p(u(\mathbf{x})|\gamma(s))p(\gamma(s)), \quad (19)$$

where the first term on the right-hand side is the likelihood of $u(\mathbf{x})$ given the filament $\gamma(s)$, and the second term is the prior probability of $\gamma(s)$.

The perfect (i.e., noise-free and perfect imaging identity function) image one would expect to see from $\gamma(s)$ is

$$u(\mathbf{x})=I(s) \cdot \delta(d(\mathbf{x}, \gamma(s))), \quad \forall \mathbf{x} \in \Omega_{\mathcal{I}}, \quad (20)$$

where $I(s)$ is the intensity of the filament, $\mathcal{d}(\cdot)$ is the closest-point transform (CPT) yielding the shortest distance from \mathbf{x} to the filament, and $\delta(\cdot)$ is the Dirac delta distribution. If \mathbf{x} is on $\gamma(s)$, then $\mathcal{d}(\cdot) = 0$ and $\delta(\cdot) = 1$, otherwise $\delta(\cdot) = 0$.

As an error metric, we consider the mean squared error (MSE) between the estimated filament $\hat{\gamma}(s)$ and the true filament $\gamma(s)$. Let the operator \mathbb{E} denote expectation with respect to its subscript. The MSE then is

$$\Sigma=\mathbb{E}_{u,\gamma} \left\{ [\hat{\gamma}(s) - \gamma(s)]^2 \right\}. \quad (21)$$

The Bayesian Cramér Rao Bound (B-CRB) \mathbf{C} (van Trees, 2004) provides a lower bound on the MSE Σ . It is the inverse of the Bayesian information \mathbf{J} of the posterior in Eq. (19). Thus, we have

$$\Sigma \geq C \equiv J^{-1}, \quad (22)$$

indicating that $\Sigma - C$ (or equivalently $\Sigma - J^{-1}$) is a positive semi-definite matrix.

We further elaborate the Bayesian information matrix J . Let Δ_γ denote the 2nd-order differential operator (i.e., the Laplacian) with respect to its subscript. The Bayesian information matrix then is

$$J = \mathbb{E}_{u, \gamma} \{-\Delta_\gamma \ln p(u(\mathbf{x}), \gamma(s))\}. \quad (23)$$

This can be interpreted as a sum of two terms after the “ln” operation:

$$J = J_\ell + J_p, \quad (24)$$

where J_p is the *a-priori* information matrix

$$J_p = \mathbb{E}_p \{-\Delta_\gamma \ln p(\gamma(s))\}, \quad (25)$$

and J_ℓ is the contribution from the give image data to the information. It is the expected value of the standard *Fisher* information matrix $\mathcal{Q}_\ell(\gamma(s))$ with respect to the prior p.d.f. $p(\gamma(s))$:

$$J_\ell = \mathbb{E}_{u, \gamma} \{-\Delta_\gamma \ln p(u(\mathbf{x})|\gamma(s))\} = \mathbb{E}_\gamma \{\mathcal{I}_\ell(\gamma(s))\}, \quad (26)$$

where

$$\mathcal{I}_\ell(\gamma(s)) = \mathbb{E}_{u|\gamma} \{-\Delta_\gamma \ln p(u(\mathbf{x})|\gamma(s))\}. \quad (27)$$

Given the image data, the lower bound in Eq. (22) generally characterizes the minimum estimation error any *unbiased* estimator must make. The decomposition into *a-priori* and data information also nicely illustrates that one can obtain more accurate estimated by either assuming stronger prior knowledge about the true filament, or by having a more informative (e.g., less noisy) data image. This general lower bound, however, is difficult to concretely evaluate due to the CPT term in $u(\mathbf{x})$. The only way of evaluating the lower bound in the general case is by applying numerical simulations.

4.2. The B-Spline level-set case

If one assumes that $\gamma(s)$ is represented by a B-Spline level set, then an analytical expression for the lower bound can be derived. By definition, this special bound is larger than or equal to the general bound above. The model $\mu(\cdot)$ in Eq. (1) is then defined by two level functions,

each of which is a B-spline as given by Eq. (4). Accordingly, the B-CRB in Eq. (23) can be used to bound the minimum error in the coefficients of the two B-spline functions, i.e., in $c_\phi[\mathbf{k}]$ and $c_\psi[\mathbf{k}]$:

$$\mathbf{J}_{c_\phi, c_\psi} = \mathbb{E}_{u, c_\phi, c_\psi} \left\{ -\Delta_{c_\phi c_\psi} \ln p(u(\mathbf{x}) | K * \mu(\mathbf{x}, \boldsymbol{\beta})) \right\} = \mathbb{E}_{u, c_\phi, c_\psi} \left\{ -\Delta_{c_\phi c_\psi} \int_{\Omega_{\mathcal{I}}} p^*(u(\mathbf{x}) | K * \mu(\mathbf{x}, \boldsymbol{\beta})) d\mathbf{x} \right\}, \quad (28)$$

where $p^*(\cdot)$ is the p.d.f. of the noise model. In our case, it has to be a member of the EF.

For any p.d.f. from the EF, this B-CRB is asymptotically attainable (Wijsman, 1973). It is hence tight. We demonstrate this below, where we experimentally evaluate the segmentation error and compare it with the theoretical bound.

5. Experimental results

We evaluate the applicability and accuracy of the present method on a number of different image types. The tests consider images from phase-contrast, fluorescence, and dark-field microscopy with varying numbers of filaments in them. For quantitatively evaluating the segmentation accuracy, we use synthetic images, in which the underlying shape is known. In the following, we will be referring to this true shape as the *ground truth*. For all tests, we use the MATLAB implementation of the algorithm.

To create synthetic benchmark data, we generate random polynomial curves of order 3 in a 150×150 continuous domain. Based on these ground truth filament objects, we simulate an expected image using the two imaging models, fluorescence and phase contrast. The simulated continuous image is then sampled onto a 150×50 discrete pixel grid. The fluorescence imaging model is continuous convolution with a Gaussian PSF of $\sigma = 3.32$ pixel, corresponding to the experimentally measured PSF FWHM = 500 nm (full width at half maximum) of the microscope used when imaging 25 nm fluorescent microtubules. The phase-contrast imaging model from Yin et al. (2010) is used with parameters measured from an image of the back focal plane of the microscope used here (phase ring inner diameter 2 mm, outer diameter 2.5 mm, back aperture 5.8 mm). After that, we add noise to the image using the corresponding noise model (Gaussian or Poisson). We choose filament and background intensities β_1 and β_2 such that a pre-defined signal-to-noise ratio (SNR) is achieved. In order to make the SNR values comparable across different noise models, we use the distribution-independent definition based on the Bhattacharyya distance \mathcal{B} between the noise-free image histogram and the noisy image histogram: $\text{SNR} = \sqrt{8\mathcal{B}}$ (Goudail et al., 2004). For the fluorescence model, we keep the background fixed at $\beta_2 = 10$ and vary the filament intensity β_1 to change the SNR. For the phase-contrast model, we keep the filament fixed at $\beta_1 = 10$ and vary the background intensity β_2 . In order to test how the algorithm behaves when the filament intensity is not uniform, we also generate complete sets of test images with linearly shaded intensity of 50% and 70% magnitude along the filament axis.

For the 50% shading, the intensity linearly varies along the filament from $1.5 \beta_1$ to $0.5 \beta_1$, such that the average intensity (and hence the average SNR) remains the same. The 70% shading case uses a linear intensity gradient from $1.7 \beta_1$ to $0.3 \beta_1$. Even though it is impossible to exhaustively test all possible spatial intensity distributions, this linear test assesses how the algorithm behaves when the actual images strongly depart from the theoretical assumption of uniform intensity. Some examples of synthetic images for different SNR are shown in Fig. 3. All images are stored as 8bit TIFF files. The complete set of 4000 benchmark images (1000 per case) can be downloaded from the MOSAIC Group's web site `mosaic.mpi-cbg.de`.

All segmentations start from the same initialization with the two level functions $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$ being Euclidean distance functions from two uniformly randomly placed points α_ϕ and α_ψ in the image domain, respectively. This describes a random filament with one bend, illustrating how the topology and geometry of the segmentation can change during algorithm evolution. However, since the problem is convex, the result is independent of the initialization.

We first provide an illustrative example for better understanding. After that, we demonstrate the differences between our present method and other state-of-the-art filament segmentation methods. We then quantitatively evaluate the segmentation errors of our proposed method on synthetic images with known ground truth. Finally, we show applicability of our proposed method to real images.

5.1. An illustrative example

We start with a simple illustrative example to visualize the working of our proposed algorithm. We artificially generate an image of two straight filaments with an SNR = 10, Gaussian noise, and blur with a Gaussian PSF of $\sigma = 3$ pixel. The resulting image is shown in Fig. 4a. The evolution of the soft indicator function $H_f(\mathbf{x})$ over the iterations of the gradient-descent procedure is shown in Fig. 4b to Fig. 4h, starting from the random initialization described above that is far from ground truth.

After 60 iterations, the algorithm has converged to the optimal $H_f^*(\mathbf{x})$ shown in Fig. 4h. The algorithm converges rapidly and correctly detects the not previously specified number of two filaments. The topology of the segmentation changes several times during the process.

Thresholding $H_f^*(\mathbf{x})$ using the post-processing algorithm correctly identifies the two filaments. Since we can evaluate the continuous B-spline level functions (and hence also $H_f^*(\mathbf{x})$) to any precision we want, the segmentation results have sub-pixel resolution.

Additionally, this example illustrates that the final indicator function $H_f^*(\mathbf{x})$ is not perfectly binary, but continuous in $[0, 1]$. The value corresponds to a localization “likelihood” of the filaments.

5.2. Qualitative comparison with previous methods

We illustrate the behavior of our method by comparing it to existing methods from each of the four classes introduced previously: an optimal pixel mask method (Alternating Split-

Bregman, “ASB”) (Paul et al., 2013), an active contour method (“Act. Contour”) (Goldstein et al., 2010), an open-curve fitting method (“SOAC”) (we use the software by Smith et al. (2010) of which a later extension to 3D is also available (Xu et al., 2014)), and a semi-automatic hybrid method (“FIESTA”) (Ruhnow et al., 2011). FIESTA (Ruhnow et al., 2011) is a semi-automatic method that uses manual thresholding to circumvent the segmentation problem. After thresholding, FIESTA efficiently solves a localization/detection problem. We nevertheless show the result here as a benchmark. We use a synthetic image containing a single filament, as this is the least common denominator that can be solved by all compared methods without additional prior knowledge. The segmentation results are shown in Fig. 5.

While both the pixel-mask method (Paul et al., 2013) and the active contour (Goldstein et al., 2010) correctly identify the filament, they represent it as a region of finite thickness, rather than as a curve. Additionally, the pixel mask as in Fig. 5a can only provide pixel-level resolution. While the open-curve fitting method as in Fig. 5c segments the filament as a curve, it is sensitive to noise and ends with a non-optimal solution where the curve fails to trace the entire filament. FIESTA Fig. 5d provides the best result from all previous methods, providing a geometric curve with sub-pixel accuracy. Our method as in Fig. 5e provides a result that is qualitatively as good as that from FIESTA, but solves the full segmentation problem without requiring any manual thresholding. It handles local minima well and is insensitive to noise because the underlying optimization problem is globally convex.

After highlighting these qualitative differences, we present a quantitative comparison of the accuracy of the present method. High accuracy is often desired when monitoring the temporal growth dynamics of individual filaments, where already changes of half a pixel or less can give rise to polymerization/depolymerization switches, which are biologically or biophysically relevant.

5.3. Quantitative evaluation

We use 4000 synthetic images (described above, 100 images per case and SNR) with known ground truth to quantitatively evaluate the detection and segmentation accuracy. We first evaluate the mean-square error (MSE) in the spline coefficients and compare it with the theoretical B-CRB, showing that our method asymptotically approaches the lower error bound. Second, we quantify how the error grows as the images depart from the theoretical assumption of uniform intensity along the filaments, using both the 50% and 70% shading cases. Third, we quantify the influence of the two algorithm parameters (i.e., the scaling size h and the regularization coefficient λ) on the accuracy of filament segmentation, and we make recommendations for parameter choices.

While measuring the error in the spline coefficients can directly be done using the MSE, comparing segmented *filaments* with ground truth requires more intuitive error metrics: first, we use the standard Dice coefficient to quantify the *detection quality* of the segmentation, i.e., how well the filaments were identified. Second, we use the Fréchet distance to quantify how *accurately* the correctly detected filaments were segmented to sub-pixel accuracy. In addition, we report the percentage of cases where the filament has been correctly identified and no spurious filaments were segmented (i.e., no under- or over-segmentation). These error metrics quantify the difference between a segmentation result and ground truth, which

is not to be confused with the Bregman divergence used to quantify differences between the model image and the data.

The Dice coefficient measures the overlap between ground truth and the obtained segmentation as $2TP / ((FP + TP) + (TP + FN))$, where TP are the true positives, defined here as all pixels in the segmentation that are within ± 5 pixel from the ground truth filament. Similarly, FN are the false negatives (pixels missed from ground truth to within ± 5 pixel) and FP the false positives (pixels of spurious filaments detected farther than 5 pixel from ground truth). We use the 5 pixel detection tolerance because the Gaussian PSF has $\sigma \approx 4$ pixel blur. A Dice coefficient of 0 means that no filament pieces were correctly detected, one of 1 means that all filament pieces were correctly detected with no spurious false positives.

For those filaments that were correctly detected within the 5 pixel tolerance, we quantify the *worst-case segmentation error* using the Fréchet distance:

$$\mathcal{F}(\mathbf{c}_a, \mathbf{c}_b) = \inf_{\zeta} \max_{\eta, t \in [0, 1]} \{ \|\mathbf{c}_a(\zeta(t)) - \mathbf{c}_b(\eta(t))\|_2 \},$$

where \mathbf{c}_a and \mathbf{c}_b are two parametric curves with parameter $t \in [0, 1]$. The functions ζ and η range over all monotone reparametrizations. This is a standard metric to measure the similarity between two curves. Identical curves have a Fréchet distance of zero. The Fréchet distance can be interpreted as the minimum cord length sufficient to join a point traveling forward along one curve and one traveling forward along the other, although the speed of travel for either point may not necessarily be uniform. The Fréchet distance is not computable for general continuous curves, but can be efficiently computed for curves given by polynomials, as is the case here. This discrete Fréchet distance is a metric and is also sometimes called *coupling measure*.

5.3.1. Error in the spline coefficients—We compare the algorithm performance against the B-CRB using the synthetic benchmark images of the fluorescence case with both uniform and shaded filament intensities. We run our algorithm to derive segmentation results for each image. We then compare the coefficients of the final spline level functions (i.e., the coefficients of the two spline level functions minimizing the energy) with those of the synthetic ground-truth splines. We show the resulting MSE and the theoretical B-CRB for the Gaussian noise model in Fig. 6. The result for the Poisson noise model is visually indistinguishable and hence not shown.

As shown in Fig. 6, for low SNR the algorithm performs sub-optimally and remains above the theoretical B-CRB lower bound. This is due to sub-optimal parameter choices. The optimization problem is solved optimally, so no error originates from its solution. Also, the imaging model used is the ground-truth model that also generated the synthetic data. The sub-optimality for low SNR must hence be due to parameter choices. While the optimal h could be 1.7 for example, we restricted ourselves to integer values in our tests, which may not be optimal. This shows only at low SNR, since both parameters h and λ affect the sensitivity of the method to noise, as shown below.

For increasing SNR, the algorithm is asymptotically optimal. Some of the error bars extend below the bound, because the B-CRB is for the expectation value of the error. Individual cases, however, can be below the bound. This is not a violation of the theory as long as the mean stays above the bound. In both noise cases, we can observe that the empirical MSEs are tight to the theoretical B-CRB. This shows that our method almost achieves the theoretical minimum of the estimation error. The asymptotic approach to the bound for higher SNR confirms the attainability of the B-CRB predicted by our above theoretical analysis.

If the filament intensity is not uniform along the filament axis, the SNR is different for different points on the filament and the SNR values are to be interpreted as average SNR over the entire image. The result accuracy then decreases, as shown by the 50% and 70% shading cases in Fig. 6. This is expected. For a 50% intensity shading, the MSE is about twice as far from the B-CRB as for uniform filaments. Even for a 70% departure from the theoretical assumption of uniformity, the MSE is still acceptable and asymptotically optimal. This suggests (albeit a theoretical proof seems out of reach) that while uniformity is a necessary mathematical assumption in deriving the optimal algorithm, the algorithm's sensitivity to violation of this assumption is bounded in practice.

5.3.2. Influence of scaling size h and regularizer λ —The present algorithm has two parameters: the scale size h and the regularization parameter λ . This generates a 2D parameter space $h \times \lambda$. In order to analyze the influence of these parameters, we repeat the above benchmark for many combinations of h and λ , performing a grid search. For each parameter set, we measure the Dice coefficient. The result is shown in Fig. 7. The imaging model used is that of a fluorescence microscopy (PSF $\sigma = 3.32$) with Gaussian noise of SNR = 4.0. For different SNR and different imaging models, this surface may look different. It can hence not directly be used to fix standard algorithm parameters.

The fitted smooth colored surface is only added for better visualization. The figure shows that mainly the parameter h affects the quality of the segmentation result. For this SNR, $h = 1$ is optimal for all tested λ . The performance is robust over a wide range of λ values spanning four orders of magnitude. While no standard parameter values can be given, the algorithm performance seems to be insensitive to variations of λ . We next investigate in more detail the influence of h as a function of image SNR.

5.3.3. Influence of scaling size h for different SNRs—The parameter h is easy to set and useful, as it allows tuning the pixel resolution required in the result. To give a better feeling for this, we quantitatively evaluate the influence of h as a function of SNR.

We do this using our 2000 synthetic benchmark images with SNR values ranging from 4 \rightarrow 30. For each image, we try scaling sizes $h \in [0, 3]$ for fixed $\lambda = 10^{-5}$, which provides the best Dice coefficient at SNR = 4.

For fluorescence images (Fig. 8a to d), $h = 1$ provides the best results for all tested SNR. Theoretically, one would expect that higher SNR allow lower h , hence providing higher-resolution segmentations. This, however, seems to be robust over the SNR range tested here.

For phase-contrast images (Fig. 8e to h), segmentation is generally more accurate, but less reliable (smaller percentage of perfect detection). An $h = 0$ is always the best choice for the phase-contrast images.

Comparing to FIESTA (solid black line in Fig. 8a and b), our method generally performs equally well, sometimes even with a higher accuracy (i.e., lower Fréchet distance). At low SNR, our method shows significantly better detection performance than FI-ESTA with 100% of all filaments correctly detected by our method and 73% by FIESTA (see numbers in Fig. 8b). For FIESTA, we manually tuned the threshold to give the best results, which was then used to fit the data for different SNRs. For our method, we manually tune λ and h .

Looking at the segmentation accuracy, the Fréchet distances are around 1.5 pixels for the fluorescence images and around 1 pixel for the phase-contrast images. This may seem like a low accuracy for a sub-pixel segmentation method. Keeping in mind that the Fréchet distance is the maximum distance between the two filaments, it is natural to ask *where* this maximum error occurs. Visually inspecting the results suggests that the error is localized to the filament tips. While the lateral tracing of the filaments is sub-pixel accurate, the segmentation often misses the exact filament tip location by 1 to 1.5 pixels. Since in biology we are often interested in computing the length of a filament from such segmentations, this is indeed the relevant error.

5.4. Application to microscopy images of biological filaments

We demonstrate the usefulness of our method, and further illustrate the influence of the parameters λ and h , using real biological microscopy data from fluorescence, phase-contrast, and dark-field microscopy. For each case, the same specimen has been imaged on the same microscope (Nikon Eclipse Ti) with the same objective (Nikon Plan Fluor 100 \times , NA 1.45 for fluorescence and phase-contrast, NA 0.5 for dark field) and the same camera (Andor Zyla 4.2; pixel size 6.4 μm , effective pixel size with the 100 \times lens 64 nm). This ensures comparability of the results. The specimen is a demembrated axoneme from a pf2:pf2GFP *Chlamydomonas reinhardtii* cell (Alper et al., 2013). This particular mutant is rescued with a GFP tagged dynein regulation complex (DRC), which localizes along the entire length of the axoneme. Thus, these axonemes can be imaged not only by phase contrast and dark-field, but also by fluorescence microscopy.

5.4.1. Application of the present method to fluorescence microscopy images

—Segmentation results for the 16bit fluorescence image are shown in Fig. 9. In our method, we use the standard fluorescence imaging model with a Gaussian PSF of 500 nm full width at half maximum (hence, $\sigma = 3.32$ pixel) and a Poisson noise model. Increasing the scaling size h in the present method for fixed λ changes the segmentation. For $h = 0$ the present method segments some of the noise close to the filament. The best result is obtained for $h = 1$ (shown). For $h = 2$, precision starts to be lost.

When using FIESTA (Ruhnow et al., 2011), the threshold for the initial segmentation was manually adjusted to 377 (out of 65,535 for this 16bit image).

The three previous model-based methods perform as expected. Specifically, the method from Paul et al. (2013) successfully locates the region containing the filament to pixel-size resolution. The closed active contour from Goldstein et al. (2010) provides good outline of the filament. The curve-fitting method from Smith et al. (2010) converges in a local minimum. All methods are used with their standard parameter settings as provided by the original publications.

5.4.2. Application of the present method to phase-contrast microscopy

images—Segmentation results for 16bit phase-contrast images are shown in Fig. 10. Since FIESTA (Ruhnow et al., 2011) does not include a phase-contrast imaging model, we first inverted the images before processing them with FIESTA. In our method, we use the imaging model for phase-contrast images from Yin et al. (2010) with parameters measured from an image of the back focal plane of the microscope and the Phase-3 oil condenser used (phase ring inner diameter 2 mm, outer diameter 2.5 mm, back aperture 5.8 mm). In addition, we use a Poisson noise model.

Fig. 10 shows the result using the present method with $\lambda = 10^{-4}$ and $h = 0$. Some of the background speckles are spuriously detected as short filaments and would have to be filtered based on their length in a post-processing step. The FIESTA (Ruhnow et al., 2011) result is free from false positives and correctly traces the length of the filament. The results obtained with the other methods used the respective standard parameter settings as provided by the original publications. The method from Paul et al. (2013) correctly identifies the pixel mask containing the filament, but also detects a false positive in the background. Also the method from Goldstein et al. (2010) generates artifacts in the back-ground while correctly tracing the outline of the filament. The SOAC curve fitting methods from Smith et al. (2010) fails to trace the filament, getting trapped in a local energy minimum on the inverted image.

5.4.3. Application of the present method to dark-field microscopy images

—We repeat the same evaluation and comparison using 16bit dark-field microscopy images of the same specimen. While an imaging model for dark-field microscopy exists (Mehta and Oldenbourg, 2014), it strongly depends on the geometry of the imaged object. It can hence not be formulated as a convolution kernel K , which is why we resort to using the imaging model for fluorescence microscopy here with $\sigma = 4$ and a Gaussian noise model.

Our present method successfully identifies the filament for all scaling sizes tested ($h = 1, 2, 3$). The best result is obtained for $h = 2$ (shown) and $\lambda = 10^{-4}$, as shown in Fig. 11. The result from FI-ESTA (Ruhnow et al., 2011) used a threshold of 4000 (out of 65,535 for the 16bit images). The method from Paul et al. (2013) successfully identifies the region containing the filament to pixel-size resolution. The closed active contour from Goldstein et al. (2010) also successfully traces the filament boundary with sub-pixel resolution. The curve-fitting method from Smith et al. (2010) converges in a local minimum. All methods are used with their standard parameter settings as provided by the original publications.

5.4.4. Application to images containing multiple filaments—Often images contain an unknown number of filaments. We show the application of our method to 16bit fluorescence images containing a number of fluorescently labeled filaments. We use the

same fluorescence imaging model as above and do not previously specify the number of filaments in the image. Our present method detects them automatically and without manual tuning of any threshold parameter. We start from the same initialization as before, i.e., a single curved filament. The segmentation results are shown in Fig. 12.

We test two images that contain two and six filaments, respectively. We show the results for the default parameters $h = 1$ and $\lambda = 10^{-5}$. The first image (a) contains two purified pf2:pf2GFP axonemes (Alper et al., 2013). The second image (b) shows six fluorescently labeled microtubules (Gell et al., 2010). This image has much lower SNR, yet the microtubules are correctly segmented.

This demonstrates that our method combines the benefits of multiple existing methods. It provides sub-pixel geometric curve descriptions of the filaments, automatically detects the number of filaments in an image, does not rely on manual thresholding, works for different imaging models and microscopy modalities, and provides globally optimal results under the postulated model that asymptotically achieve the Bayesian Cramér-Rao lower bound in the spline coefficients.

6. Conclusions

We have presented a new algorithm to solve the filament segmentation problem for curved one-dimensional filaments in two-dimensional digital images. The presented method has only two parameters and does not require prior knowledge about the number of filaments in the image. It provides segmentation results with sub-pixel accuracy as open polynomial curves. This enables high-resolution studies of filament dynamics and provides access to spatial and temporal derivatives of the curve and its time evolution. Open curves are represented by a vector level-set, which enables topological changes during contour evolution. The method is flexible to accommodate different noise models and image-acquisition models. The segmentation result is guaranteed to be optimal for the given model. This is the result of a convex formulation of the underlying optimization problem, using Bregman divergences to measure the similarity between images, and a GLM model for the image pixel statistics. In other words, our method exploits prior knowledge about the physical properties of the imaging system in order to improve segmentation quality. Note that the geometric and photometric problems are likely not jointly convex, but separately convex, which is enough for convergence. If needed, a jointly convex formulation could possibly be derived extending recent works (Brown et al., 2012).

We provided an efficient algorithm to compute the optimal filament segmentation under the postulated models. The algorithm combines the analytical energy gradients with a fast convolution-based gradient evaluation. We have also for the first time shown an information-theoretic lower bound on the segmentation error in filament segmentation. No algorithm can possibly do better than this bound given the information in the image. We have shown that for the specific bound of B-spline level sets, the presented algorithm is asymptotically optimal for high SNR. For low SNR, the presented algorithm is not the best one could do, as it remains slightly above the theoretical bound. This, however, is due to parameter choices,

as discussed. Quantifying the residual fitting un-certainty could, for example, be done using particle filters in an outer loop, as previously described (Cardinale et al., 2009).

Comprehensive simulations were conducted to illustrate the major differences with prior works, show main features of our method, quantitatively evaluate the performance on synthetic images, and demonstrate applications on real biological images. The benchmarks have also identified recommended choices for the algorithm parameters h and λ : for fluorescence images with SNR between 4 and 30, one should set h to 1, for phase-contrast image $h = 0$. The λ value is largely inconsequential. A default value of 10^{-5} should yield good results.

The time required by the algorithm to segment an image depends on the image size and on how far the initialization is from the final segmentation. It does not depend on the number of filaments in the image. The current MATLAB implementation processes a 512×512 pixel 8bit image in about 30 seconds using MATLAB 2013b on a single core of an Intel Core i7 2.2 GHz processor.

In its current form, the method has a number of limitations. Probably the most restrictive one is that filament crossings and overlaps are not allowed. This is an inherent limitation of the two-component vector level-set description of filaments. Handling crossing filaments would require using more than two level functions (three to allow single crossings, four for tripple points, etc.). Currently, the method is also limited to 2D images. Segmenting filaments in 3D images would require the use of three level functions. A limitation of the current theoretical derivation is that the intensities of the filaments and the background are assumed to be uniform. As we have shown, however, the segmentation is still asymptotically optimal even when these assumptions are not fulfilled, with a bounded increase in the error. Intensity variations along a filament could then be read out during post processing and used as biological readouts, but the tips of the filaments may be inaccurately estimated, especially when the filaments get dimmer towards the tips.

There are a number of directions along which future work could improve on the present framework. An obvious thing to do would be to extend to more than two level functions in order to handle crossing filaments and higher-dimensional images. Extending the method to allow for shaded fore- and background intensities in a theoretically optimal way is less straightforward and requires more research in convex relaxation theory. A more short-term goal could be to include a line search over h in a scale-space pyramid (i.e., automatically repeating the segmentation for successively lower h) to find the optimal h that leads to the lowest global energy minimum in each case. Together with a default value for λ , this would effectively render the algorithm parameter-free.

We provide an open Matlab implementation of the presented algorithm as a reference. In addition, the algorithm has also been implemented in Java, as part of the open-source MOSAIC-suite plugin for the user-friendly image-analysis environments ImageJ (Schneider et al., 2012) and Fiji (Schindelin et al., 2012). Both implementations are freely available from the MOSAIC Group's web site mosaic.mpi-cbg.de. The main contributions of the present paper were to: (1) show an information-theoretic lower bound on the segmentation

error for filaments, providing a baseline against which algorithms can be compared; (2) introduce the vector level-set idea for filament segmentation, uniting the advantages of both region-segmentation and curve-fitting methods; (3) showing that the resulting optimization problem can be solved efficiently and in a globally optimal way; (4) providing a practical framework for filament segmentation across different imaging modalities and noise models.

Acknowledgments

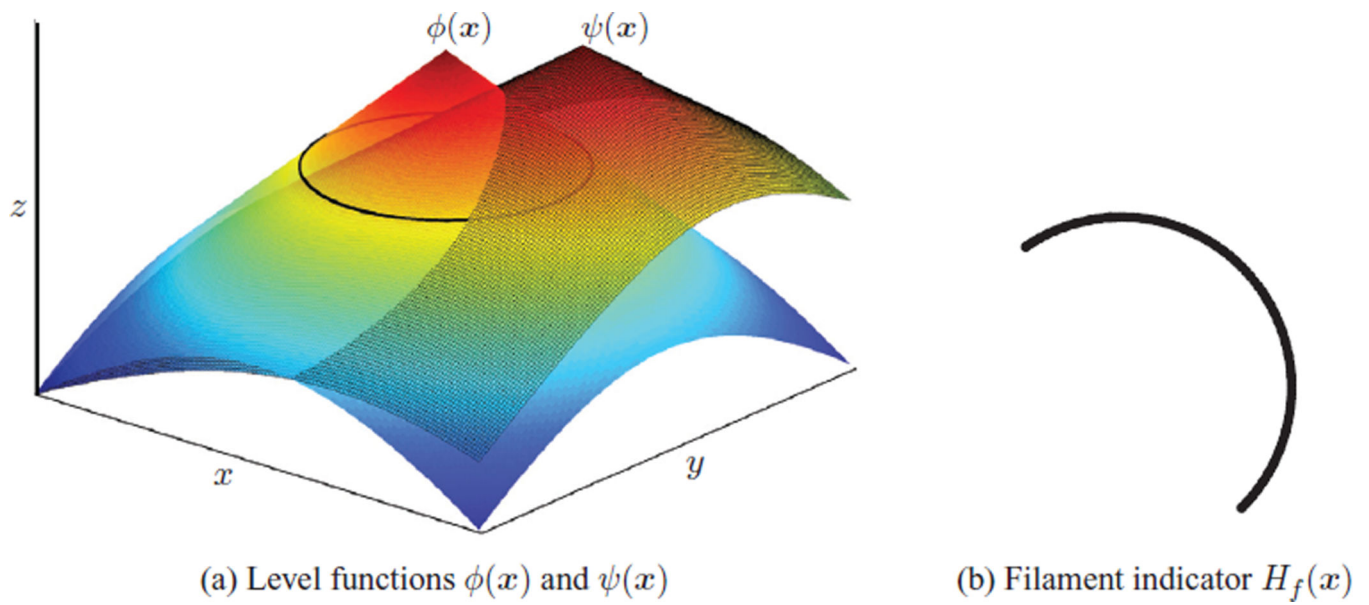
We thank all members of the MOSAIC Group for the many fruitful discussions. Particular thanks go to Ulrik Günther for helping derive the theoretical error bound. This work was supported by the German Federal Ministry of Research and Education (BMBF) under funding code 031A099.

References





- Alper J, Geyer V, Howard J. Reconstitution of flagellar sliding. *Methods Enzymol.* 2013; 524:343–369. [PubMed: 23498749]
- Banerjee A, Merugu S, Dhillon IS, Ghosh J. Clustering with Bregman divergences. *J. Mach. Learn. Resch.* 2005; 6:1705–1749.
- Bernard O, Friboulet D, Thévenaz P, Unser M. Variational B-spline level-set: a linear filtering approach for fast deformable model evolution. *Image Process. IEEE Trans. on.* 2009; 18:1179–1191.
- Boykov Y, Funka-Lea G. Graph cuts and efficient nd image segmentation. *Int. J. Comput. Vis.* 2006; 70:109–131.
- Boykov Y, Kolmogorov V. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Anal. Mach. Intell. IEEE Trans. on.* 2004; 26:1124–1137.
- Bregman LM. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.* 1967; 7:200–217.
- Brown E, Chan T, Bresson X. Completely convex formulation of the Chan-Vese image segmentation model. *Int. J. Comput. Vis.* 2012; 98:103–121.
- Butenuth M, Heipke C. Network snakes: graph-based object delineation with active contour models. *Mach. Visi. Appl.* 2012; 23:91–109.
- Cardinale J, Paul G, Sbalzarini IF. Discrete region competition for unknown numbers of connected regions. *IEEE Trans. Image Process.* 2012; 21:3531–3545. [PubMed: 22481820]
- Cardinale, J.; Rauch, A.; Barral, Y.; Székely, G.; Sbalzarini, IF. Bayesian image analysis with on-line confidence estimates and its application to microtubule tracking. *IEEE International Symposium on Biomedical Imaging (ISBI); IEEE; Boston, USA.* 2009. p. 1091-1094.
- Chao J, Ram S, Ward ES, Ober RJ. Ultrahigh accuracy imaging modality for super-localization microscopy. *Nat. Methods.* 2013; 10:335–338. [PubMed: 23455923]
- Clack NG, O'Connor DH, Huber D, Petreanu L, Hires A, Peron S, Svoboda K, Myers EW. Automated tracking of whiskers in videos of head fixed rodents. *PLoS Comp. Biol.* 2012; 8:e1002591.
- Danuser G, Tran PT, Salmon ED. Tracking differential interference contrast diffraction line images with nanometre sensitivity. *J. Microsc.* 2000; 198:34–53. [PubMed: 10781207]
- Farin, GE. *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide.* London, UK: Academic Press, Inc; 1996.
- Florin, C.; Paragios, N.; Williams, J. *MICCAI.* Springer; 2005. Particle filters, a quasi-Monte Carlo solution for segmentation of coronaries; p. 246-253.2005
- Fuller N, Abouadarham J, Bentley RD. Filament recognition and image cleaning on Meudon H α spectroheliograms. *Sol. Phys.* 2005; 227:61–73.
- Gelas A, Bernard O, Friboulet D, Prost R. Compactly supported radial basis functions based collocation method for level-set evolution in image segmentation. *Image Process. IEEE Trans. on.* 2007; 16:1873–1887.

- Gell C, Bormuth V, Brouhard GJ, Cohen DN, Diez S, Friel CT, Helenius J, Nitzsche B, Petzold H, Ribbe J, et al. Microtubule dynamics reconstituted in vitro and imaged by single-molecule fluorescence microscopy. *Methods Cell Biol.* 2010; 95:221–245. [PubMed: 20466138]
- Geng W, Cosman P, Berry CC, Feng Z, Schafer WR. Automatic tracking, feature extraction and classification of *C. elegans* phenotypes. *Biomed. Eng. IEEE Trans. on.* 2004; 51:1811–1820.
- Gittes F, Mickey B, Nettleton J, Howard J. Flexural rigidity of microtubules and actin filaments measured from thermal fluctuations in shape. *J. Cell Biol.* 1993; 120:923–934. [PubMed: 8432732]
- Goldstein T, Bresson X, Osher S. Geometric applications of the split Bregman method: segmentation and surface reconstruction. *J. Sci. Comput.* 2010; 45:272–293.
- Goudail F, Réfrégier P, Delyon G. Bhattacharyya distance as a contrast parameter for statistical processing of noisy optical images. *JOSA A.* 2004; 21:1231–1240. [PubMed: 15260255]
- Graham JS, McCullough BR, Kang H, Elam WA, Cao W, Enrique M. Multi-platform compatible software for analysis of polymer bending mechanics. *PloS ONE.* 2014; 9:e94766. [PubMed: 24740323]
- Gunawardana A, Byrne W. Convergence theorems for generalized alternating minimization procedures. *J. Mach. Learn. Res.* 2005; 6:2049–2073.
- Kass M, Witkin A, Terzopoulos D. Snakes: active contour models. *Int. J. Comput. Vis.* 1988; 1:321–331.
- Kolmogorov V, Zabini R. What energy functions can be minimized via graph cuts? *Pattern Analysis and Machine Intelligence, IEEE Transactions on.* 2004; 26:147–159.
- Law, MWK.; Chung, ACS. MICCAI. Springer; 2009. A deformable surface model for vascular segmentation; p. 59-67.2009
- Maeland E. On the comparison of interpolation methods. *Med. Imag. IEEE Trans. on.* 1988; 7:213–217.
- McCullagh P. Generalized linear models. *Eur. J. Oper. Res.* 1984; 16:285–292.
- Mehta SB, Oldenbourg R. Image simulation for biological microscopy: microlith. *Biomed. Opt. Expr.* 2014; 5:1822–1838.
- Mukundan V, Sartori P, Geyer VF, Jülicher F, Howard J. Motor regulation results in distal forces that bend partially disintegrated chlamydomonas axonemes into circular arcs. *Biophys. J.* 2014; 106:2434–2442. [PubMed: 24896122]
- Nelder JA, Wedderburn RWM. Generalized linear models. *J. Roy. Stat. Soc. A.* 1972:370–384.
- Ober RJ, Ram S, Ward ES. Localization accuracy in single-molecule microscopy. *Biophys. J.* 2004; 86:1185–1200. [PubMed: 14747353]
- Ober RJ, Tahmasbi A, Ram S, Lin Z, Ward ES. Quantitative aspects of single-molecule microscopy – information-theoretic analysis of single-molecule data. *IEEE Signal Proc. Mag.* 2015; 32:58–69.
- Paul G, Cardinale J, Sbalzarini IF. Coupling image restoration and segmentation: a generalized linear model/Bregman perspective. *Intl. J. Comp. Vis.* 2013:1–25.
- Ramot D, Johnson BE, Berry TL Jr, Carnell L, Goodman MB. The parallel worm tracker: a platform for measuring average speed and drug-induced paralysis in nematodes. *PloS ONE.* 2008; 3:e2208. [PubMed: 18493300]
- Rigort A, Günther D, Hegerl R, Baum D, Weber B, Prohaska S, Medalia O, Baumeister W, Hege HC. Automated segmentation of electron tomograms for a quantitative description of actin filament networks. *J. Struct. Biol.* 2012; 177:135–144. [PubMed: 21907807]
- Rikmenspoel R, Isles CA. Digitized precision measurements of the movements of sea urchin sperm flagella. *Biophys. J.* 1985; 47:395. [PubMed: 3978210]
- Ronfard R. Region-based strategies for active contour models. *Int. J. Comput. Vis.* 1994; 13:229–251.
- Ruhnnow F, Zwicker D, Diez S. Tracking single particles and elongated filaments with nanometer precision. *Biophys. J.* 2011; 100:2820–2828. [PubMed: 21641328]
- Saban, M.; Altinok, A.; Peck, A.; Kenney, C.; Feinstein, S.; Wilson, L.; Rose, K.; Manjunath, BS. Automated tracking and modeling of microtubule dynamics. *IEEE Symposium on Biomedical Imaging (ISBI); IEEE; 2006. p. 1032-1035.*

- Sarry L, Boire JY. Three-dimensional tracking of coronary arteries from biplane angiographic sequences using parametrically deformable models. *IEEE Trans. Med. Imag.* 2001; 20:1341–1351.
- Schindelin J, Arganda-Carreras I, Frise E, Kaynig V, Longair M, Pietzsch T, Preibisch S, Rueden C, Saalfeld S, Schmid B, Tinevez JY, White DJ, Hartenstein V, Eliceiri K, Tomancak P, Cardona A. Fiji: an open-source platform for biological-image analysis. *Nat. Methods.* 2012; 9:676–682. [PubMed: 22743772]
- Schneider CA, Rasband WS, Eliceiri KW. NIH image to ImageJ: 25 years of image analysis. *Nat. Methods.* 2012; 9:671–675. [PubMed: 22930834]
- Sethian, JA. *Level Set Methods and Fast Marching Methods*. Cambridge, UK: Cambridge University Press; 1999.
- Smith MB, Li H, Shen T, Huang X, Yusuf E, Vavylonis D. Segmentation and tracking of cytoskeletal filaments using open active contours. *Cytoskeleton.* 2010; 67:693–705. [PubMed: 20814909]
- van Trees, HL. *Detection, Estimation, and Modulation Theory*. New York, USA: John Wiley & Sons; 2004.
- Tsai A, Yezzi A Jr, Wells W, Tempany C, Tucker D, Fan A, Grimson WE, Willisky A. A shape-based approach to the segmentation of medical imagery using level sets. *Med. Imag. IEEE Trans. on.* 2003; 22:137–154.
- Tsai R, Osher S. Review article: Level set methods and their applications in image science. *Commun. Math. Sci.* 2003; 1:1–20.
- Valdman D, Atzberger PJ, Yu D, Valentine MT. Spectral analysis methods for the robust measurement of the flexural rigidity of biopolymers. *Biophys. J.* 2012; 102:272–284.
- Wijsman RA. On the attainment of the cramer-rao lower bound. *The Ann. Stat.* 1973:538–542.
- Wong YY, Yuen PC, Tong CS. Segmented snake for contour detection. *Pattern Recognit.* 1998; 31:1669–1679.
- Xu T, Vavylonis D, Huang X. 3D actin network centerline extraction with multiple active contours. *Med. Image Anal.* 2014; 18:272–284. [PubMed: 24316442]
- Yezzi A Jr, Kichenassamy S, Kumar A, Olver P, Tannenbaum A. A geometric snake model for segmentation of medical imagery. *Med. Imag. IEEE Trans. on.* 1997; 16:199–209.
- Yin, Z.; Li, K.; Kanade, T.; Chen, M. *Proc. MICCAI, Medical Image Computing and Computer-Assisted Intervention*. Beijing, China: Springer; 2010. Understanding the optics to aid microscopy image segmentation; p. 209-217.
- Zhang F, Zhang X, Cao K, Li R. Contour extraction of gait recognition based on improved GVF Snake model. *Comput. Elect. Engineering.* 2012a; 38:882–890.
- Zhang, Q.; Liu, L.; Liu, B. *Pattern Recognition*. Berlin Heidelberg, Germany: Springer; 2012b. Adaptive vector flow for active contour model; p. 121-128.

**Fig. 1.**

Vector level-set representation of open curves. (a) Two example level functions ϕ and ψ , each shown as a colored surface. Color is the z value (height). The bold black line shows the closed curve defined by $\{x: \phi(x) = 0\}$. The second level function ψ (with meshgrid) is then used to define the open curve $\{x: \phi(x) = 0 \text{ and } \psi(x) > 0\}$. (b) The so-defined open curve is represented by its indicator function H_f (Eq. (3)).

(a) $h = 0$ (b) $h = 1$ (c) $h = 2$ (d) $h = 3$ **Fig. 2.**

A B-spline surface with different scaling sizes h . The scaling size h allows tuning the geometric resolution of the spline, providing regularization against noise. It defines the node spacing in units of pixels with $h = 0$ (a) placing a spline node into every pixel, $h = 1$ (b) into every other pixel, and so on (c)–(d).

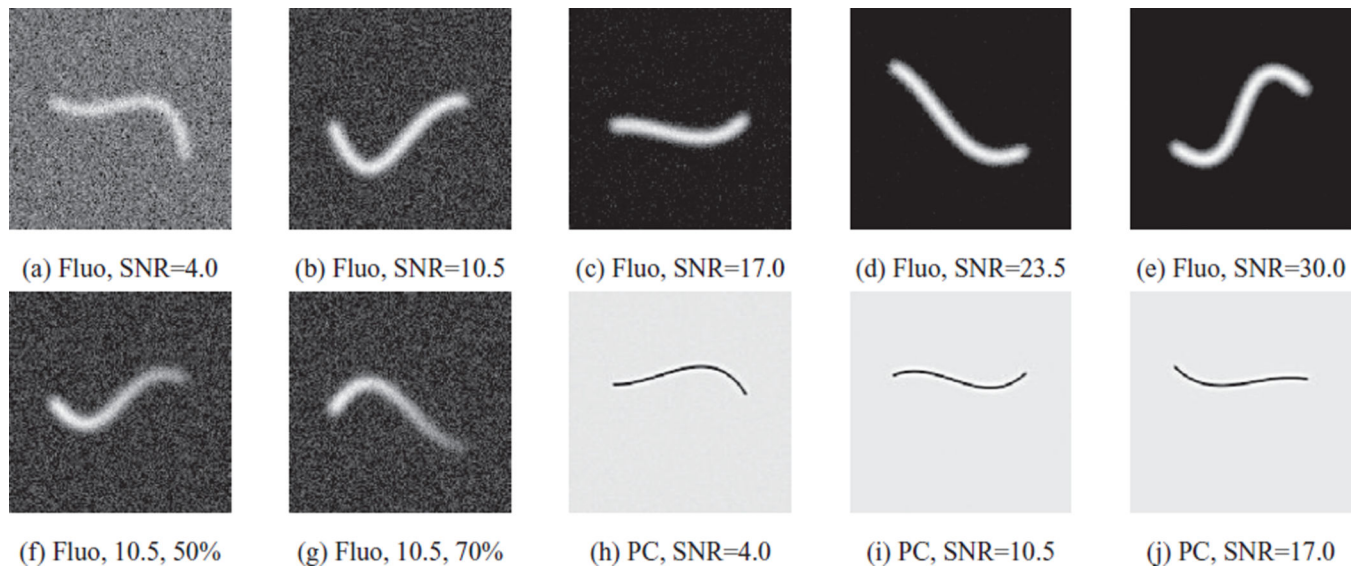


Fig. 3. Examples of synthetic benchmark images showing random filaments with known ground truth. (a)–(e) Five different SNR values using the fluorescence imaging model (Gaussian noise, $\sigma = 3.32$ pixel) with uniform filament intensity. (f)–(g) Images from the fluorescence imaging model (Gaussian noise, $\sigma = 3.32$ pixel, SNR = 10.5) with linearly shaded filaments of 50% and 70% shading, respectively. (h)–(j) Three different SNR values using the phase-contrast imaging model (Poisson noise). The complete set of 4000 benchmark images can be downloaded from the MOSAIC Group’s web site mosaic.mpi-cbg.de.

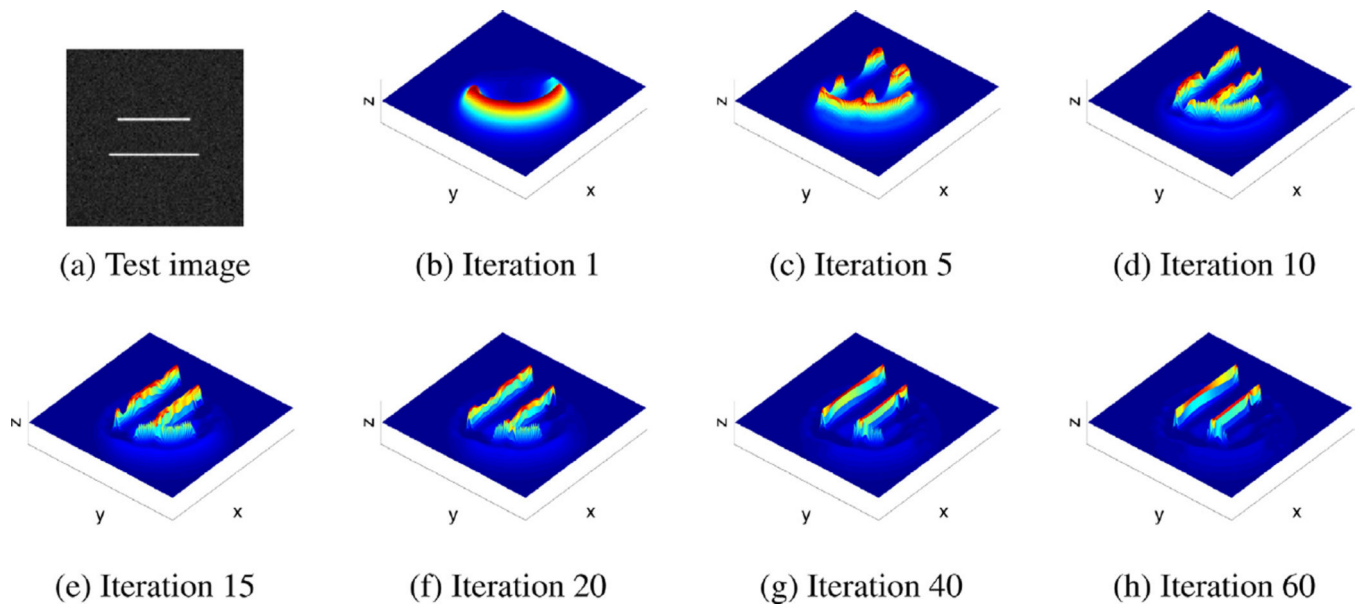


Fig. 4. Evolution of $H_f(x)$ during gradient descent. (a) A test image of size 100×100 pixel with SNR = 10, Gaussian noise, and Gaussian blur with $\sigma = 3$ pixel. (b)–(h) Snapshots of $H_f(x)$ at iterations 1, 5, 10, 15, 20, 40, 60, starting from a random initialization far from ground truth.

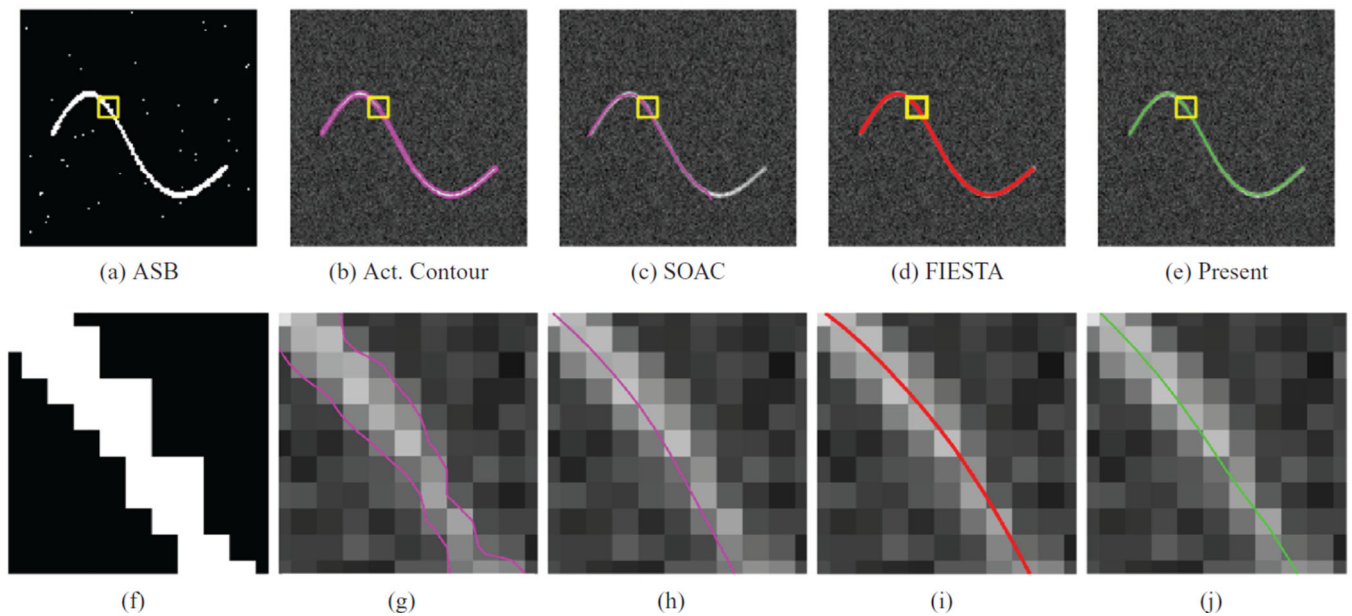


Fig. 5. Qualitative comparison of different methods on an example synthetic image (Gaussian noise with SNR = 4, Gaussian PSF with $\sigma = 3.3$ pixel). (a)–(e): the segmentation results produced by ASB (Paul et al., 2013), Active Contours (Goldstein et al., 2010), SOAC (Smith et al., 2010), FIESTA (with threshold manually adjusted to 140 out of 255 for the present 8bit images) (Ruhnow et al., 2011), and our present method. (f)–(j): Corresponding magnifications of the regions in the yellow boxes above. For FIESTA, the following settings were used: FWHM = 1100 nm, relative intensity threshold 160%, smooth after tracking, correct for focus drift, Track especially curved filaments = 30%. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

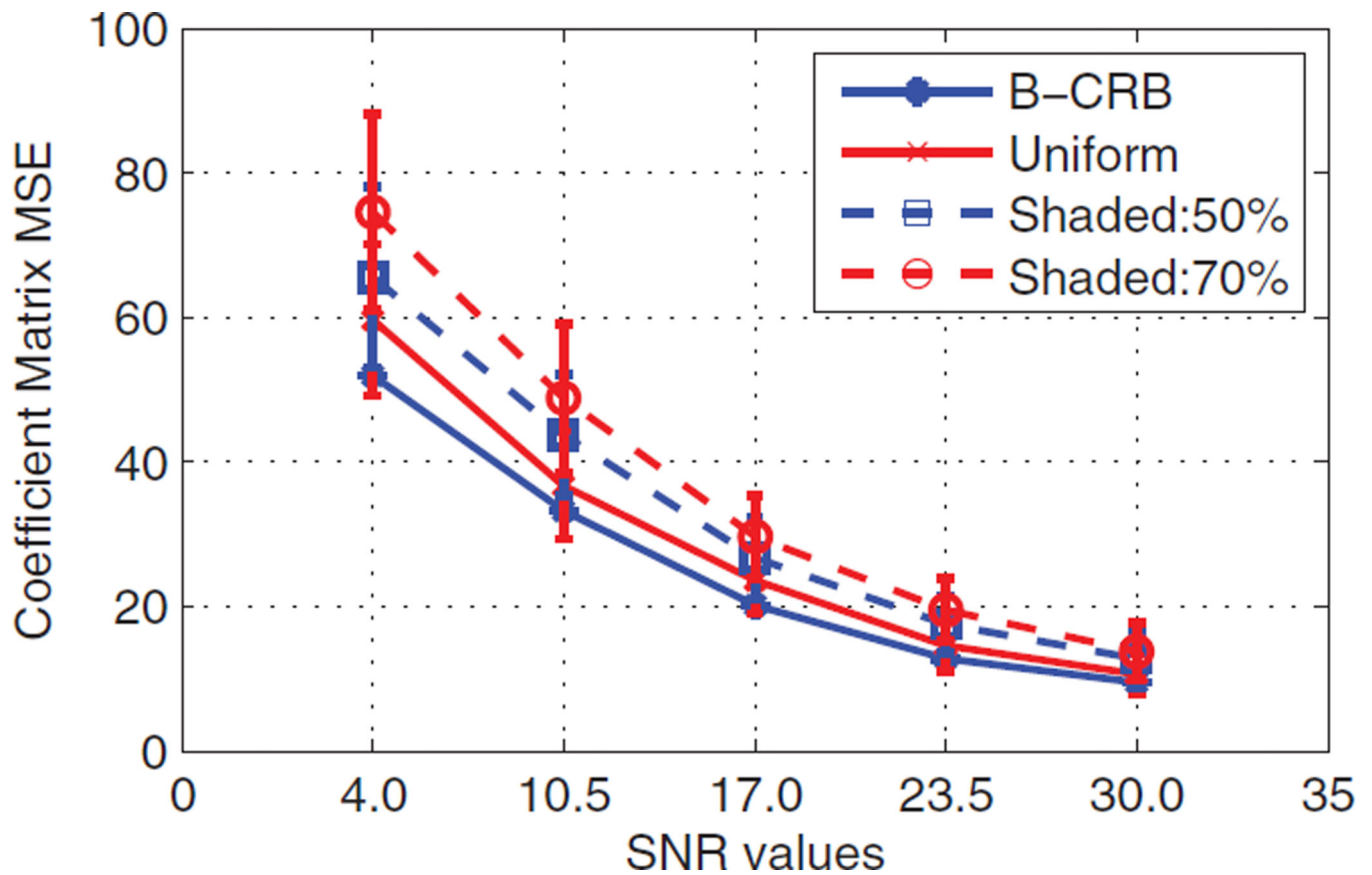


Fig. 6. Quantitative evaluation of the mean-square error (MSE, Eq. (21)) in the level-function spline coefficients (mean \pm standard deviation) for different SNR and different linear intensity gradients along the filaments, in comparison with the theoretical Bayesian Cramér-Rao Bound (B-CRB) of the error. 100 random synthetic filament images are processed for each data point with $\lambda = 10^{-5}$ and $h = 0$ in all cases. Only the result for Gaussian noise and the fluorescence imaging model is shown. The result for Poisson noise is visually indistinguishable and hence omitted.




Fig. 7. Dice coefficient from ground truth for different scaling sizes $h \in [0, 4]$ and regularizer coefficients $\lambda = [10^{-7}, 10^{-3}]$, increasing 10-fold at each step, shown for the case of fluorescence imaging with Gaussian noise of SNR = 4.

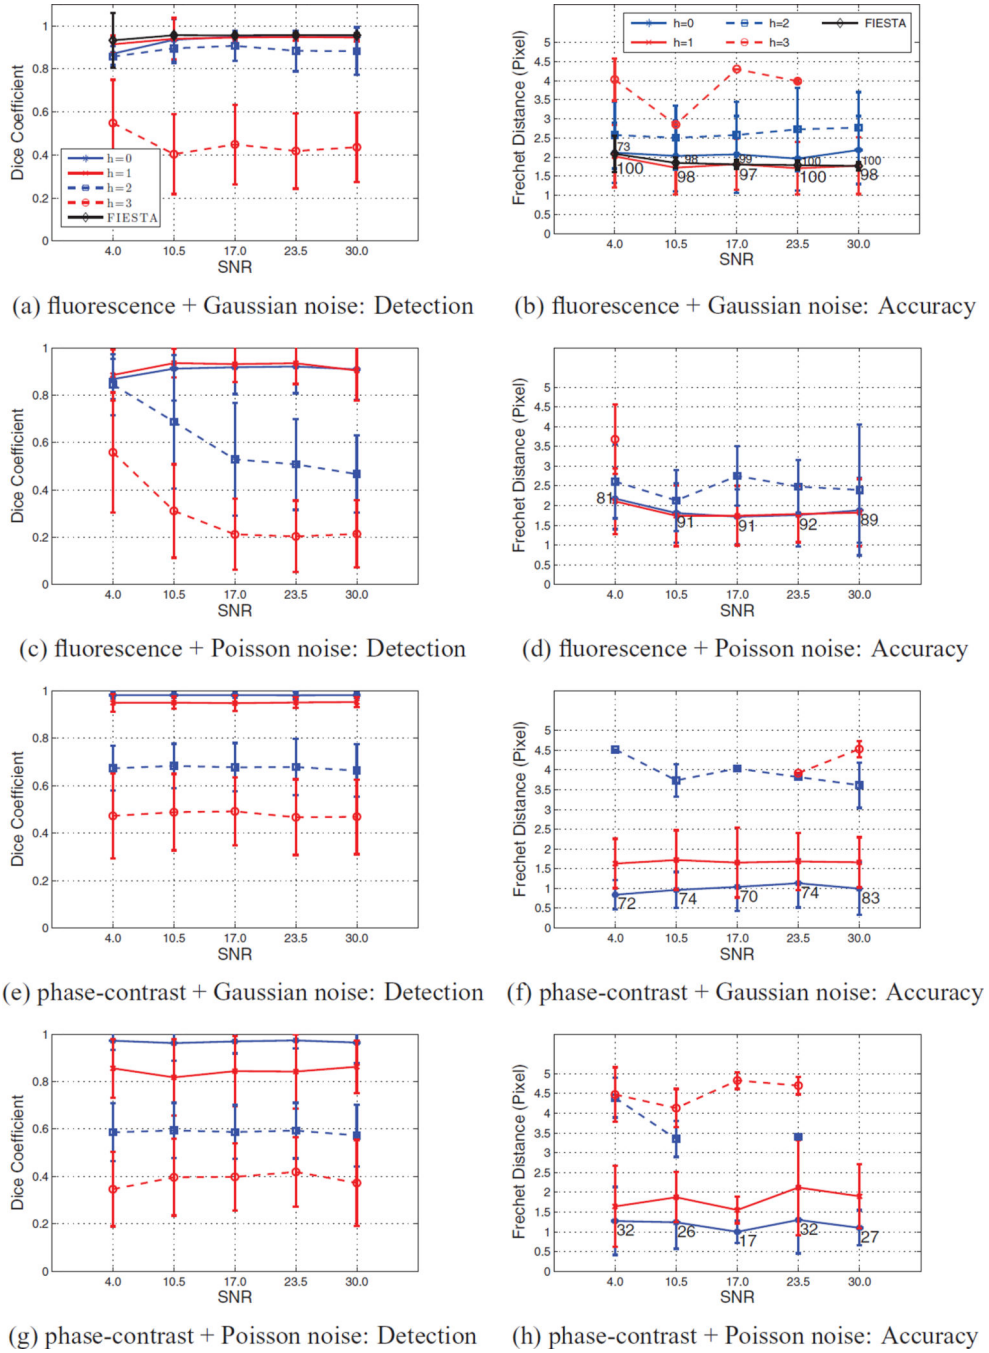


Fig. 8. Detection accuracy (Dice coefficient) and segmentation accuracy (Fréchet distance) for all tested cases. All values are given as mean \pm standard deviation vs. SNR when using the correct noise and imaging models. The numbers near the symbols are the percentage of cases in which the correct filament was correctly detected to within ± 5 pixel. Only these cases are taken into account when computing the Fréchet distance. For the present method, we use a constant $\lambda = 2.5 \cdot 10^{-3}$ for all cases and vary h as indicated in the legends. For FIESTA, we use a manually optimized relative threshold of 160% in all cases, tuned for the

lowest SNR and then applied across all SNR. Other FIESTA settings were: track filaments only, smooth after, PixelSize = 100 nm/pixel, FWHM = 1100 nm, area = 100 pixel, correct for focus drift, track especially curved filaments = 30%. Missing symbols are those cases where no acceptable segmentation within ± 5 pixel from ground truth was found.





Fig. 9. Application to a fluorescence image of a demembrated axoneme from a pf2:pf2GFP *Chlamydomonas reinhardtii* cell. (a) Segmentation result using the present method with $h = 1$ and $\lambda = 10^{-5}$. (b) Segmentation result using FIESTA (Ruhnnow et al., 2011) with threshold 377 (out of 65,535 for the present 16bit image). (c)–(e) Segmentation results using the methods from Paul et al. (2013), Goldstein et al. (2010), and Smith et al. (2010) with their standard parameters, respectively.

**Fig. 10.**

Application to a phase-contrast image of a demembrated axoneme from a pf2:pf2GFP *Chlamydomonas reinhardtii* cell. (a) Segmentation result using the present method with $\lambda = 10^{-4}$ and $h = 0$. (b) Segmentation result using FIESTA (Ruhnow et al., 2011) on the inverted image with relative threshold = 156%, FWHM = 100, track curved filaments = 30%, drift correction on, 64 nm/pixel. (c)–(e) Segmentation results using the methods from Paul et al. (2013), Goldstein et al. (2010), and Smith et al. (2010) with their standard parameters, respectively.






Fig. 11.

Application to a dark-field microscopy image of a demembrated axoneme from a pf2:pf2GFP *Chlamydomonas reinhardtii* cell. (a) Segmentation result using the present method with $h = 2$ and $\lambda = 10^{-4}$. (b) Segmentation results using FIESTA (Ruhnow et al., 2011) with threshold 4000 (out of 65,535 for this 16bit image). (c)–(e) Segmentation results using the methods from Paul et al. (2013), Goldstein et al. (2010), and Smith et al. (2010) with their standard parameters, respectively.



(a) Axonemes



(b) Microtubules

Fig. 12.

Application of the present method to two fluorescence microscopy images containing multiple filaments: (a) two purified pf2:pf2GFP axonemes (Alper et al., 2013); (b) six GMPCPP-stabilized microtubules, *in vitro* polymerized from a mixture of Rhodamine-labeled and unlabeled pig brain Tubulin (1:3) imaged by fluorescence microscopy (Gell et al., 2010). We use the default parameters $h = 1$ and $\lambda = 10^{-5}$.