Microtubule Dynamics Analysis using Kymographs and Variable-Rate Particle Filters

Ihor Smal*, *Member, IEEE,* Ilya Grigoriev, Anna Akhmanova,
Wiro J. Niessen, *Member, IEEE,* and Erik Meijering, *Senior Member, IEEE*

Abstract

Studying intracellular dynamics is of fundamental importance for understanding healthy life at the molecular level and for developing drugs to target disease processes. One of the key technologies to enable this research is the automated tracking and motion analysis of these objects in microscopy image sequences. To make better use of the spatiotemporal information than common frame-by-frame tracking methods, two alternative approaches have recently been proposed, based on either Bayesian estimation or space-time segmentation. In this paper, we propose to combine the power of both approaches, and develop a new probabilistic method to segment the traces of the moving objects in kymograph representations of the image data. It is based on variable-rate particle filtering and uses multiscale trend analysis of the extracted traces to estimate the relevant kinematic parameters. Experiments on realistic synthetically generated images as well as on real biological image data demonstrate the improved potential of the new method for the analysis of microtubule dynamics in vitro.

Index Terms

Bayesian estimation, variable-rate particle filters, multiscale trend analysis, motion analysis, biological microscopy, microtubule dynamics.

I. INTRODUCTION

Motion analysis of subcellular objects plays a major role in understanding fundamental dynamical processes occurring in biological cells. Since many diseases originate from a disturbance or failure

---

Asterisk indicates corresponding author.

*I. Smal is with the Biomedical Imaging Group Rotterdam, Departments of Medical Informatics and Radiology, Erasmus MC, Rotterdam, The Netherlands. E-mail: i.smal@erasmusmc.nl.

I. Grigoriev and A. Akhmanova are with the Department of Cell Biology, Erasmus MC, Rotterdam, The Netherlands.

W. J. Niessen is with the Biomedical Imaging Group Rotterdam, Departments of Medical Informatics and Radiology, Erasmus MC, Rotterdam, The Netherlands, and Imaging Science & Technology, Faculty of Applied Sciences, Delft University of Technology, Delft, The Netherlands.

E. Meijering is with the Biomedical Imaging Group Rotterdam, Departments of Medical Informatics and Radiology, Erasmus MC, Rotterdam, The Netherlands.

of one or more of these processes, their study is of interest not only to life scientists, but also to pharmaceutical companies in the attempt to develop adequate drugs. Even though many intracellular interaction mechanisms are well understood these days, many questions still remain unanswered. In some cases, where the analysis in living cells (in cultures or *in vivo*) is confounded by other intracellular processes, it makes sense to study the objects of interest *in vitro*, where the influence of other structures or processes is removed, reduced, or known [1], [2].

Intracellular dynamics is usually visualized using advanced fluorescence microscopy imaging techniques, including stimulated emission and depletion (STED) and total internal reflection fluorescence (TIRF) microscopy, where the objects of interest are labeled with fluorescent proteins [3]. Alternatively, non-fluorescence based techniques, such as phase contrast (PC) or differential interference contrast (DIC) microscopy can sometimes be used, which do not require labeling [4], [5]. In either case, the optical resolution of the microscope is much lower (on the order of 100 nm) than the size of the objects of interest (on the order of nanometers), causing the latter to be imaged as blurred spots (without sharp boundaries) due to diffraction. The quality of the images is further reduced by high measurement noise levels [3], [5]. Both types of distortions contribute to the ambiguity of the data, making automated quantitative image analysis an extremely difficult task.

In time-lapse microscopy, where hundreds to thousands of 2D or 3D images are acquired sequentially in time, the main task is to track the objects of interest (proteins, vesicles, microtubules, etc.) and compute relevant motion parameters from the extracted trajectories. In practice, manual tracking is labor intensive and poorly reproducible, and only a small fraction of the data can be analyzed this way. The vast majority of automatic tracking methods [6]–[12] developed in this field consist of two stages: 1) *detection* of objects of interest (independently in each frame), and 2) *linking* of detected objects from frame to frame (solving the correspondence problem). Since the methods employed for the first stage operate on data with low signal-to-noise ratio (SNR), the linking procedure in the second stage is faced with either many false positives (noise classified as objects) or false negatives (misdetection of actually present objects).

Contrary to these two-stage tracking methods, which typically use only very few neighboring frames to address the correspondence problem, methods that make better use of the available temporal information usually show better results. Such trackers are either built within a Bayesian estimation framework [13], [14], which in any frame uses all available temporal information up to that frame, or they consider the 2D+t or 3D+t image data as one spatiotemporal 3D or 4D image, respectively, and translate the estimation of trajectories into a segmentation of spatiotemporal structures [15], [16].

In this paper, we propose to combine the power of the latter two approaches, and develop a variable-rate particle filtering method that implements the Bayesian estimation framework for tracing spatiotemporal structures formed by transforming the original time-lapse microscopy image data into a special type of spatiotemporal representation: kymographs [17]–[21]. This combined approach, which to the best of our knowledge has not been explored before, results in more accurate extraction of the spatiotemporal structures (edge-like image structures in our case) compared to particle filtering applied directly to the image sequences on a per-frame basis [13].

The paper is organized as follows. In Section II, we describe the biological application considered in this paper and the proposed methods to model, acquire, transform, preprocess, and analyze the image data. In Section III, we present experimental results of applying our method to synthetic image sequences, for which ground truth was available, and to real DIC microscopy image data of microtubule dynamics. A concluding discussion of the main findings is given in Section IV.

II. METHODS

A. *In Vitro Microtubule Dynamics Model*

Microtubules (MTs) are polymers of tubulin, which assemble into hollow tubes (diameter ∼25 nm) in the presence of guanosine triphosphate (GTP), both *in vivo* and *in vitro* [22], [23]. *In vivo*, MTs are responsible for the support and shape of the cell and play a major role in several intracellular processes such as cell division, internal cell organization, and intracellular transport. MT dynamics (also referred to as dynamic instability) is highly regulated, both spatially and temporally, by a wide family of microtubule-associated proteins (MAPs) [24]. To understand the specific interactions between regulatory factors and microtubules is of great interest to biologists. Misregulation of MT dynamics, for example, can lead to erroneous mitosis, which is a characteristic feature in neurodegenerative diseases.

Microtubule dynamic instability is a stochastic process of switching between growth and shrinkage stages, regulated by MAPs [25]. The growth velocity, ν^+, depends on soluble tubulin concentration available for polymerization and GTP-tubulin association and dissociation rates. The shrinkage velocity, ν^-, which is usually much higher than the growth velocity, is independent of tubulin concentration and is characterized only by the dissociation rate of guanosine diphosphate (GDP) tubulin from the depolymerizing end. The growth velocity *in vivo* can be up to 10 times faster than *in vitro*. Two other important events that characterize dynamic instability are *rescue* (switching from shrinkage to growth) and *catastrophe* (switching from growth to shrinkage) [25]. In practice, the analysis of MT dynamics includes estimation of ν^+, ν^-, and the rescue and catastrophe frequencies, f_{res} and f_{cat}. The rescue rate *in vitro* is very low unless specific rescue factors are added to the assay and might be difficult to estimate reliably [2].

Recent studies reveal a special class of MAPs, plus-end-tracking proteins (+TIPs), that are able to accumulate at MT growing ends [24], [26]–[28]. The mechanisms by which +TIPs recognize MT ends have attracted much attention and several explanations have been proposed [24], [27], [29]. One way to understand the mechanism employed by individual +TIPs and the molecular mechanisms underlying

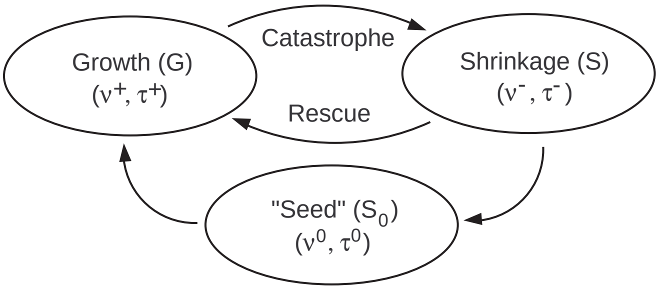

Fig. 1. Dynamics model describing microtubule behavior *in vitro*.

their functions is by measuring the distribution and displacement of +TIPs in time. However, due to lack of robust and accurate automatic methods, the manual analysis usually is a labor intensive procedure which very likely leads to user bias and loss of important information. In the case of experiments in living cells it is extremely hard to decouple the effect of other regulators while studying +TIPs influence on MT dynamics. The advantage of *in vitro* investigation is the minimal environment in which the influence of various +TIPs can be dissected individually. Recent *in vitro* studies start to reveal the mechanisms of +TIPs end-tracking and the regulation of MT dynamics by individual +TIPs [30]. This can potentially lead to combining multiple +TIPs in order to reconstitute the *in vivo* MT dynamics and observe the collective effect of +TIPs.

The stochastic behavior of the MT tip can be modeled using a dynamical system with three states (Fig. 1): G (growth), S (shrinkage), and S_0 (no dynamic activity). Each state is characterized by a velocity parameter $\tilde{\nu} \in \{\nu^+, \nu^-, \nu^0\}$ and a duration time interval $\tilde{\tau} \in \{\tau^+, \tau^-, \tau^0\}$, describing the duration of the corresponding stage. The following state transitions are allowed: $S_0 \to G$ (the MT starts to grow), $G \to S$ (catastrophe), $S \to G$ (rescue), and $S \to S_0$ (the MT is completely disassembled). At each time point the MT can "stay" only in one of the states and for a period of time no longer than the corresponding $\tilde{\tau}$ for that state. In our simulations, the time and velocity parameters are generated randomly (Section III-A), and because of that it is allowed to "leave" the state S sooner than τ^- if the MT is completely disassembled in shorter time. If after time τ^- the MT was not disassembled completely (did not reach state S_0), a rescue occurs ($S \to G$) and the MT switches to growing. A similar three-stage model of MT dynamics can be designed for the *in vivo* situation. In this case, state S_o should be replaced with a state that corresponds to a "pause" event [22], and all the transitions (arrows in Fig. 1) should be bidirectional.

B. Imaging Technique and Kymographs

In our studies, the dynamic behavior of MTs is imaged using DIC microscopy [31], which is effectively used for biological specimens that cannot be visualized with sufficient contrast using bright-

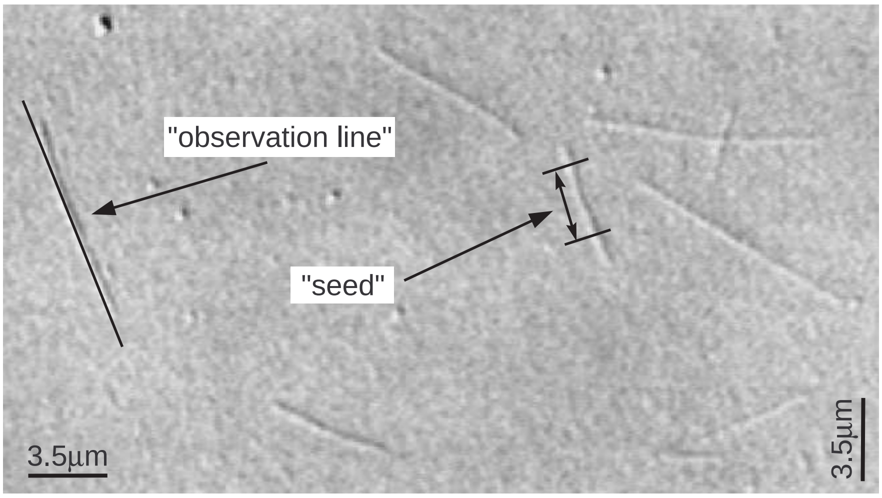

Fig. 2. Example of a DIC microscopy image. Microtubule nucleation initiates from stable tubulin “seeds”. In the experiments, “observation lines” are defined along MT bodies to construct kymographs.

field microscopy. The resulting images (see Fig. 2 for an example) are similar to those obtained with phase-contrast microscopy and depict objects as black/white shadows on a gray background with good resolution and clarity. The advantages of DIC over fluorescence microscopy is that the samples do not have to be stained. This also eliminates the possible influence of fluorescent proteins on dynamical processes being studied. The main limitation of this imaging technique is its requirement for a thin and transparent sample of fairly similar refractive index to its surroundings.

Automatic analysis of MT behavior *in vitro* using time-lapse DIC microscopy images is a complicated task. The goal is to follow (track) the fast-growing (so called “plus”) end of each MT so as to obtain 2D paths in the image plane, from which all the parameters of interest (velocity and frequency estimates) can be computed. One of the main problems is that due to the nonlinear image formation process in DIC microscopy, the object appearance (and especially the MT tip) depends on the imaging conditions (for example, the relative angle between the sample and the microscope polarization prism) and cannot be easily modeled by appearance models, as in the case of fluorescence microscopy imaging. Additionally, the real object location is further obscured by diffraction, modeled by the point-spread function (PSF) of the microscope.

Another issue that requires careful consideration is the temporal sampling rate. In our experiments, images are acquired at a rate of one per second, which is quite high for imaging the microtubule growth velocities *in vitro* (30–40 nm per second). The shrinkage, on the other hand, happens with much higher velocities, and with the mentioned sampling rate this process is imaged in only 3–10 image frames (the time between the catastrophe event and arriving at state S_0). On average, microtubule growth is usually observed during 30–100 frames, and in this case the process is highly oversampled. It is not possible simply to lower the sampling rate, as this causes undersampling of the shrinkage processes.

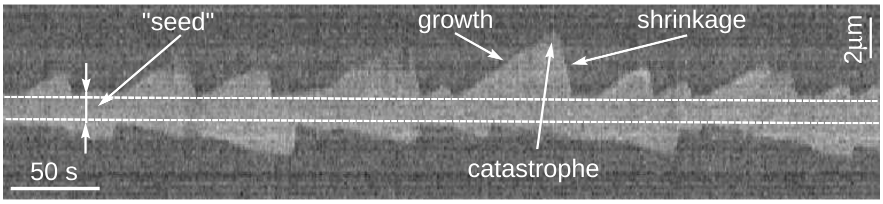

Fig. 3. Example of a kymograph obtained from the DIC microscopy images, showing the dynamics of both microtubule ends.

This problem arises due to the fact that we are imaging two processes (growth and shrinkage), which may occur simultaneously (for different microtubules) on timescales that differ approximately by a factor of ten, with one constant sampling rate. The relatively high sampling rate is both a blessing and a curse. It is a blessing because it allows one to observe the motion in more detail and possibly detect rare and extraordinary movements that would otherwise go unnoticed. It is also a curse, however, as the growth and shrinkage velocities are usually such that the change in MT length from one frame to the next is (much) less than one pixel, even if the *spatial* sampling is done at the Nyquist rate. This is on the same order as the positional estimation errors made by manual or automatic approaches [13]. As a result, instant velocity estimates (ν^+ or ν^-) computed as the ratio of positional change over elapsed time between two consecutive frames, are doomed to be highly inaccurate.

In order to exploit all image data and at the same time obtain more accurate results, we abandon the idea of frame-by-frame tracking of objects directly in the original data, and we propose to base the estimation of motion parameters on a transformation of the data that is more amenable to multiscale analysis. Specifically, we propose to use a kymograph representation [17], [21] (also called a kymoimage in this paper) for each MT. It is constructed by defining (see Section III-A5) an “observation line” L (Fig. 2) in the original image along the MT body. The length of L should approximately equal the maximum expected length of MTs in the sample. Image intensity values are then sampled equidistantly along L, yielding a vector of “measurements” at time t, $J_t = \{J_t(j) : j = 1, \ldots, Y\}$, where Y is the number of samples for the selected MT in every image frame. In practice, to increase the SNR, the measurements $J_t(j)$ are obtained by averaging pixel values in the vicinity of j, along a line perpendicular to L. The resulting kymoimage (see Fig. 3 for an example), $I(t, y) = \{J_t : t = 1, \ldots, T\}$, is the collection of measurement vectors, where every column t contains the measurements J_t as pixel values, and T is the number of frames in the image sequence. In our experiments, MT nucleation from stable tubulin oligomers was studied [32]. These “seeds” always remain present and cannot be completely disassembled. In the kymoimages (Fig. 3) they are clearly visible as a bright horizontal strip.

To estimate the kinematic parameters of interest from the kymoimages, the edge location $y(t)$ (cor-

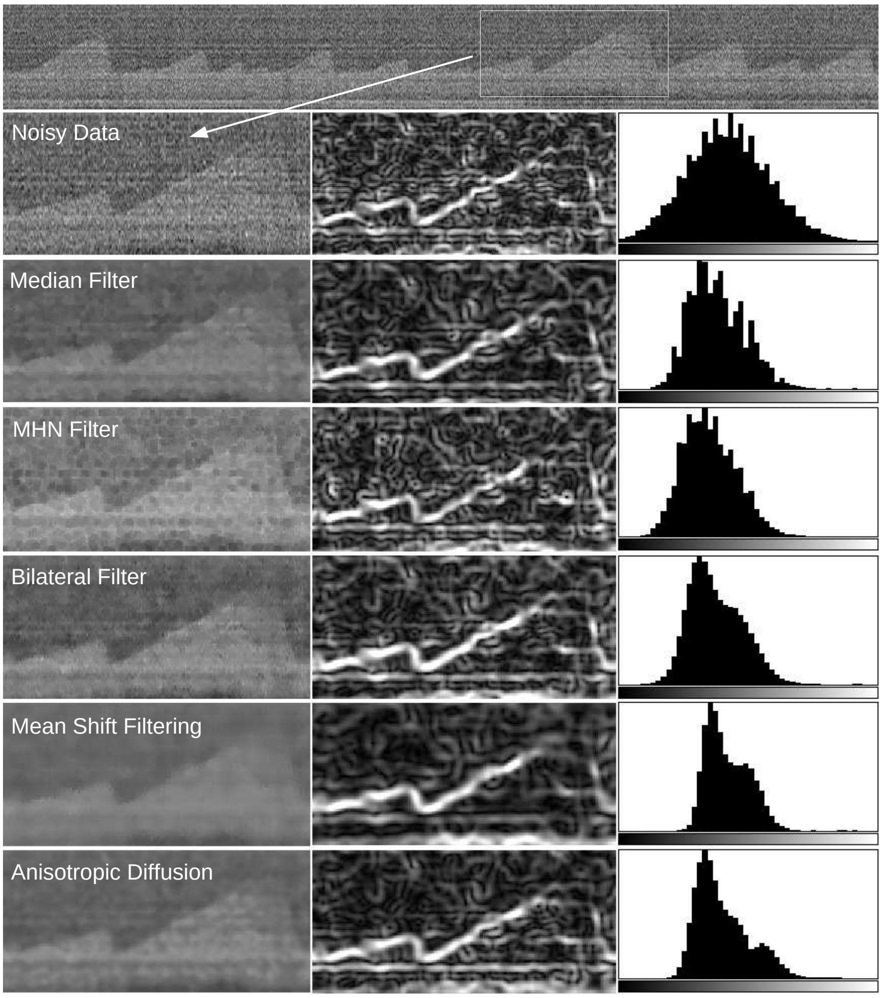

Fig. 4. Application of various edge preserving smoothing methods to our image data (top). The left column shows the results of smoothing, the middle column depicts the edge information extracted using Gaussian derivatives, and the right column shows the distribution of intensity values in the smoothed images.

responding to the MT tip) should be accurately extracted (slopes should be preserved). In kymoimages, the instant velocity ν at any time t' is estimated as $\nu = (dy/dt)_{t=t'} = \tan(\varphi)$, where φ is the angle between the time axis and the tangent to $y(t)$ at $t = t'$. As a result, small errors in the angle estimates may lead to large errors in the velocity estimation, due to the nonlinearity introduced by the tangent (the closer φ is to 90 degrees, the larger the errors).

In our method, the analysis is conducted in three subsequent steps: 1) preprocessing, 2) edge extraction, and 3) multiscale trend analysis. Step 1 enhances the quality of the image using edge preserving filtering. Step 2 traces the edges by a particle filter capable of using multiscale measurements. Finally, step 3 analyzes the extracted edges by splitting them into relevant parts and performing linear approximation in order to compute all the necessary parameters. The three steps are described in more detail in the following subsections.

C. Edge Preserving Smoothing

The main challenge in estimating the growth velocity ν^{+}, shrinkage velocity ν^{-}, and the two transition frequencies f_{res} and f_{cat}, is to accurately segment the edges in the kymoimages (Fig. 3). Two main approaches to edge detection are differentiation and model fitting. In practice, differentiation, being a noise enhancing operation, requires some form of smoothing, which in turn entails the risk of blurring edge information. Better results may be obtained by the use of nonlinear, edge preserving filters. Fig. 4 shows the results of applying the most frequently used nonlinear filtering techniques to our image data: the median filter [33], the maximum homogeneity neighbor (MHN) filter [34], the bilateral filter [35], the mean-shift filter [36], and anisotropic diffusion [37]. The examples clearly demonstrate that noise can be reduced to some extent while preserving edge information. However, they also show that edges may still not be clearly defined in (parts of) the image. Subsequent edge extraction by means of Gaussian differentiation [38] may result either in detection of noisy background structures (at small scales), or in too much positional uncertainty (at larger scales), neither of which is acceptable for accurate slope estimation of the linear parts of the edge $y(t)$.

To overcome the problems caused by differentiation, we propose to use model fitting for edge detection, using particle filtering (PF) methods. The PF can be exploited to reduce the overload of fitting the model in every pixel position, by incorporating information about the edge model, the image noise distribution, and the probability of finding the edge in the neighborhood of a pixel, by taking into account the probability of edge existence at neighboring pixels. In this case, the use of edge preserving prefiltering is still advantageous. The PF mainly replaces the edge extraction part, which in differentiation based approaches such as Canny's algorithm [39] is usually based on hard thresholding.

D. Variable-Rate Particle Filtering

The prefiltered kymoimage is an input for the next step, where particle filtering (PF) is performed to estimate the edge location $y(t)$. Particle filters [40], [41] implement the concept of Bayesian estimation, where at each time point t a system state $\mathbf{x}_t$ is estimated on a basis of previous states, noisy measurements $\mathbf{z}_t$ obtained from sensors, and prior knowledge about the underlying process [40]. For our application, the simplest working implementation of PF can be constructed with the state vector $\mathbf{x}_t$, which describes the position of the edge in every column t of the image $I(t, y)$, and the measurements $\mathbf{z}_t$, which are the intensity values in the corresponding column t of $I(t, y)$. Prior knowledge about the system is specified by the dynamics model, which describes the state transition process, and the observation model:

$$\mathbf{x}_t = f_t(\mathbf{x}_{t-1}, \mathbf{v}_t), \qquad \mathbf{z}_t = g_t(\mathbf{x}_t, \mathbf{u}_t), \tag{1}$$

where f_t and g_t are possibly nonlinear functions and $\mathbf{v}_t$ and $\mathbf{u}_t$ are white noise sources. The choice of these functions is application specific and is given below. Alternatively, the same state estimation problem can be formulated by specifying two distributions, $p(\mathbf{x}_t|\mathbf{x}_{t-1})$ and $p(\mathbf{z}_t|\mathbf{x}_t)$, instead of (1) [40], [41].

The solution of the state-space problem given by (1) is the posterior probability distribution function (pdf) $p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t})$, where $\mathbf{x}_{0:t} = \{\mathbf{x}_0, \ldots, \mathbf{x}_t\}$ and $\mathbf{z}_{0:t} = \{\mathbf{z}_0, \ldots, \mathbf{z}_t\}$, which can be found either exactly (when f_t and g_t are linear and $\mathbf{v}_t$ and $\mathbf{u}_t$ are Gaussian) using the Kalman filter [41] or, in the most general case, using approximations such as sequential Monte Carlo (MC) methods [40], [42]. In the latter case, the posterior pdf is approximated with a set of N_s MC samples (referred to as "particles"), $\{\mathbf{x}_{0:t}^{(i)}, w_t^{(i)}\}_{i=1}^{N_s}$, as

$$p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t}) = \sum_{i=1}^{N_s} w_t^{(i)} \delta(\mathbf{x}_{0:t} - \mathbf{x}_{0:t}^{(i)}), \tag{2}$$

where $\mathbf{x}_{0:t}^{(i)}$ describes one of the possible state sequences (path) and $w_t^{(i)}$ is the weight indicating the probability of realization of that path. The solution using PF is given by a recursive procedure that predicts the state from time $t-1$ to t and updates the weights based on newly arrived measurements $\mathbf{z}_t$ as

$$\mathbf{x}_t^{(i)} \sim p(\mathbf{x}_t|\mathbf{x}_{t-1}^{(i)}) \text{ and } w_t^{(i)} \propto w_{t-1}^{(i)} p(\mathbf{z}_t|\mathbf{x}_t^{(i)}), \tag{3}$$

$i = 1, \ldots, N_s$. The minimum mean square error (MMSE) or maximum a posteriori (MAP) estimators of the state can be easily obtained from $p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t})$ [40].

Commonly, the state sampling rate is determined by the rate at which the measurements arrive. In the application under consideration, where the MT dynamics is characterized by prolonged periods of smoothness (growth and shrinkage stages) with infrequent sharp changes (rescue and catastrophe), it is possible to obtain a much more parsimonious representation of the MT tip trajectory if the state sampling rate is adapted to the nature of the data – more state points are allocated in the regions of rapid variation and relatively fewer state points to smoother sections. Unfortunately, this idea cannot be implemented using the standard PFs because the number of state points, which would typically be much smaller than the number of observations, is random and unknown beforehand. In order to deal with this randomness, variable-rate particle filtering (VRPF) methods have been proposed recently [43], [44]. The VRPF can be compared to the more conventional interactive multiple models (IMM) approach, which uses switching between a discrete set of candidate dynamical models [8], [45], but was shown to outperform IMM in most cases [43]. The VRPF, which was initially proposed for tracking of highly maneuvering targets [43], is nowadays successfully applied in other fields, for example DNA sequencing [46], but has not been investigated before in microscopy.

Contrary to the standard state-space approach, where the state variable $\mathbf{x}_t$ evolves with time index t, within the VRPF framework the state $\mathbf{x}_k$ is defined as $\mathbf{x}_k = (\theta_k, \tau_k)$, where $k \in \mathbb{N}$ is a discrete state index, $\tau_k \in \mathbb{R}^+ > \tau_{k-1}$ denotes the arrival time for the state k, and θ_k denotes the vector of variables necessary to parametrize the object state. In tracking applications, the vector θ_k includes variables such as position, velocity, heading, etc. For our application, we define $\theta_k = (y_k, v_k)$, where y_k is the edge position at time τ_k along the observation line L, and $v_k = (dy/dt)_{t=\tau_k}$ describes the direction of the edge at $t = \tau_k$ in the image $I(t, y)$. Similar to the standard PF, it is assumed that the state sequence is a Markov process, so the successive states are independently generated with increasing k according to

$$\begin{aligned} \mathbf{x}_k &\sim p(\mathbf{x}_k|\mathbf{x}_{k-1}) \\ &= p_\theta(\theta_k|\theta_{k-1}, \tau_k, \tau_{k-1}) p_\tau(\tau_k|\theta_{k-1}, \tau_{k-1}). \end{aligned} \tag{4}$$

These assumptions and models, apart from the constraint $\tau_k > \tau_{k-1}$, are very general, and the specific choices are dictated by the application under investigation.

The measurements $\mathbf{z}_t$, $t \in \mathbb{N}$, occur on a regular time grid and in the case of the standard PF can be uniquely associated with the corresponding state $\mathbf{x}_t$. In the VRPF framework, the underlying state process is asynchronous with the measurement process and the rate of arrival of the measurements is typically (but not necessarily) higher than that of the state process. In order to define the appropriate observation model (also called the likelihood) in this case, where there may be no corresponding state variable for the measurement at time t, the data points $\mathbf{z}_t$ are assumed to be independent of all other data points, conditionally upon the neighborhood $\mathcal{N}_t$ of states $\mathbf{x}_{\mathcal{N}_t} = \{\mathbf{x}_k; k \in \mathcal{N}_t\}$, that is

$$\mathbf{z}_t \sim p(\mathbf{z}_t|\mathbf{x}_0, \ldots, \mathbf{x}_\infty) = p(\mathbf{z}_t|\mathbf{x}_{\mathcal{N}_t}). \tag{5}$$

The neighborhood $\mathcal{N}_t$ is constructed as a deterministic function of the time index t and the state sequence $\mathbf{x}_{0:\infty}$ and thus it is a random variable itself (this feature is not present in the standard state-space models). For practical (computational) reasons, the neighborhood $\mathcal{N}_t$ will contain only states whose times τ_k are "close" to the observation time t. Furthermore, the interpolated state $\hat{\theta}_t = h_t(\mathbf{x}_{\mathcal{N}_t})$ is used, where $h_t(.)$ is a deterministic function of the state in the neighborhood $\mathcal{N}_t$. The observation density (5) is then expressed as

$$p(\mathbf{z}_t|\mathbf{x}_{\mathcal{N}_t}) = p(\mathbf{z}_t|\hat{\theta}_t). \tag{6}$$

In general, the construction of the state process and the neighborhood structure is not unique and for any given model and different choices will lead to different algorithmic trade-offs.

Having all the definitions, we aim to recursively estimate the sequence of variable-rate state points as new measurements become available. Similar to the standard PF, the VRPF distribution $p(\mathbf{x}_{0:\mathcal{N}_t^+}|\mathbf{z}_{0:t})$ can be obtained using the two-step predict-update procedure, where $\mathcal{N}_t^+$ denotes the index of the state variable in $\mathcal{N}_t$ that has the largest time index τ_k. Using the factorization (4), we model the MT dynamics with the transition priors

$$\begin{aligned} & p_\theta(\theta_k|\theta_{k-1},\tau_k,\tau_{k-1}) = \\ & \quad p(v_k|v_{k-1},y_k,y_{k-1},\tau_k,\tau_{k-1}) \times \\ & \quad p(y_k|v_{k-1},y_{k-1},\tau_k,\tau_{k-1}) = \\ & \qquad p(v_k|v_{k-1})\delta(y_k - y_{k-1} - v_{k-1}(\tau_k - \tau_{k-1}), \end{aligned} \tag{7}$$

$$p_\tau(\tau_k|\theta_{k-1},\tau_{k-1}) = \mathcal{U}_{[\tau_{k-1}+\tau^0,\tau_{k-1}+\tau^1]}, \tag{8}$$

where $\mathcal{U}_{[a,b]}$ denotes the uniform distribution in the range $[a,b]$. Thus, the states $\mathbf{x}_k$ for the prediction-update procedure are sampled as

$$\begin{aligned} \tau_k - \tau_{k-1} &\sim \mathcal{U}_{[\tau^0,\tau^1]}, \\ y_k &= y_{k-1} + v_{k-1}(\tau_k - \tau_{k-1}), \\ v_k &\sim p(v_k|v_{k-1}). \end{aligned} \tag{9}$$

The sampling of the new states $\mathbf{x}_k$ at time t is performed only for those particles $\mathbf{x}_{k-1}^{(i)}$ for which $\tau_{k-1}^{(i)} \le t$, which also reduces the computational load compared to the standard PF implementation.

The crucial point here is to efficiently model the prior $p(v_k|v_{k-1})$ in order to catch the rapid changes in edge orientation (corresponding to the state transitions described in Section II-A). The underlying assumption about the MT dynamics in this study is that the MT end can either grow with nearly constant velocity ν^+, shrink with nearly constant velocity ν^-, or show almost no activity ($\nu^0 \approx 0$). This idealization of reality can be justified by specifying additionally the variances for the velocity estimates, $\sigma^2_{\nu^+}$, $\sigma^2_{\nu^-}$, $\sigma^2_{\nu^0}$, which account for small deviations in the measured velocities from the average values ν^+, ν^-, and ν^0. Taking into account three possible types of motion, we define the following prior $p(v_k|v_{k-1})$ for the velocity component v_k

$$p(v_k|v_{k-1}) = \begin{cases} (1-a)\mathcal{N}(v_{k-1},\sigma^2_{\nu^+}) + \\ \quad a\mathcal{N}(\nu^-,\sigma^2_{\nu^-}), \\ \qquad \text{for } v_{k-1} > V_{\text{th}} \\ (1-a)\mathcal{N}(v_{k-1},\sigma^2_{\nu^-}) + \\ \quad a\left(\mathcal{N}(\nu^+,\sigma^2_{\nu^+}) + \mathcal{N}(\nu^0,\sigma^2_{\nu^0})\right)/2 \\ \qquad \text{for } v_{k-1} < -V_{\text{th}} \\ (1-a)\mathcal{N}(v_{k-1},\sigma^2_{\nu^+}) + \\ \quad a\mathcal{N}(\nu^+,\sigma^2_{\nu^+}) \\ \qquad \text{for } |v_{k-1}| < V_{\text{th}} \end{cases} \tag{10}$$

where $0 < a < 1$ is a weighting parameter that balances the mixture components corresponding to different types of motion in the transition pdf (in tracking applications, a would correspond to the probability of object/target birth). The threshold V_{th} defines which prior should be used: it defines the smallest velocity below which all the small changes in the MT length are considered to belong to state S_0. Since all three types of MT motion are quite different, the performance of the algorithm is not influenced by possible inaccuracies in setting up the threshold V_{th}, which can be estimated in advance from the experimental data. Additionally, the thresholding at V_{th} does not imply that at every time point we assume that the system evolves according to only one model. Due to the probabilistic nature of the VRPF, at every time step the posterior pdf describes the probability to find the MT in each of the three states.

In order to define the likelihood $p(\mathbf{z}_t|\mathbf{x}_{\mathcal{N}_t})$, we model the edge appearance using an observation model that we have previously used successfully for tracking of tubular structures in noisy medical images [47], [48]. The proposed model describes a small perfectly sharp edge and consists of two rectangular regions, S_B and S_F (black and white rectangles in Fig. 5, respectively). For each intermediate state $\hat{\theta}_t = h_t(\mathbf{x}_{\mathcal{N}_t})$, which is required for the likelihood computation, the neighborhood is defined as $\mathcal{N}_t = \{k, k-1; \tau_{k-1} \leq t < \tau_k\}$. For the MT length changes, linear interpolation between two neighboring states θ_k and θ_{k-1} is used, $y_t = y_{k-1} + v_{k-1}(t - \tau_{k-1})$, and the orientation of the rectangles for each time point $t \in [\tau_{k-1}, \tau_k)$ is defined by the velocity component $v_{k-1}^{(i)}$. The regions S_B and S_F are defined as follows

$$S_B(\hat{\theta}_t) = S_B(\tau_{k-1}, \tau_k, v_{k-1}) = \left\{ \left(\tfrac{l - v_{k-1} b}{\sqrt{1+v_{k-1}^2}}, \tfrac{l v_{k-1} + b}{\sqrt{1+v_{k-1}^2}} \right) : l \in [0, l_v], b \in [0, d] \right\}, \tag{11}$$

$$S_F(\hat{\theta}_t) = S_F(\tau_{k-1}, \tau_k, v_{k-1}) = \left\{ \left(\tfrac{l + v_{k-1} b}{\sqrt{1+v_{k-1}^2}}, \tfrac{l v_{k-1} - b}{\sqrt{1+v_{k-1}^2}} \right) : l \in [0, l_v], b \in [0, d] \right\}, \tag{12}$$

where $l_v = (\tau_k - \tau_{k-1})\sqrt{1 + v_{k-1}^2}$.

To measure the likelihood of edge existence at some image position with an orientation defined by the velocity component of the state vector, the average image intensity values, μ_B and μ_F, are computed over the regions S_B and S_F. The likelihood is defined as

$$p(\mathbf{z}_t|\mathbf{x}_{\mathcal{N}_t}) \propto \begin{cases} \exp\left(\frac{\mu_F - \mu_B}{\gamma}\right) - 1, & \mu_F - \mu_B > 0, \\ 0, & \mu_F - \mu_B \leq 0, \end{cases} \tag{13}$$

which defines the pdf of the edge location and favors sharp edges over smoother noisy intensity transitions. Two model parameters that control the sensitivity to the edge location, the width d and the scaling factor γ, should be specified. The length l_v is automatically defined by the time sampling functions (8). The variety in the length of the observation model adds a multiscale property to the analysis. In general, for small values of l_v the estimation of μ_B and μ_F is less accurate than for larger values of l_v. Additionally, for large l_v the sensitivity of the observation model to the edge orientation

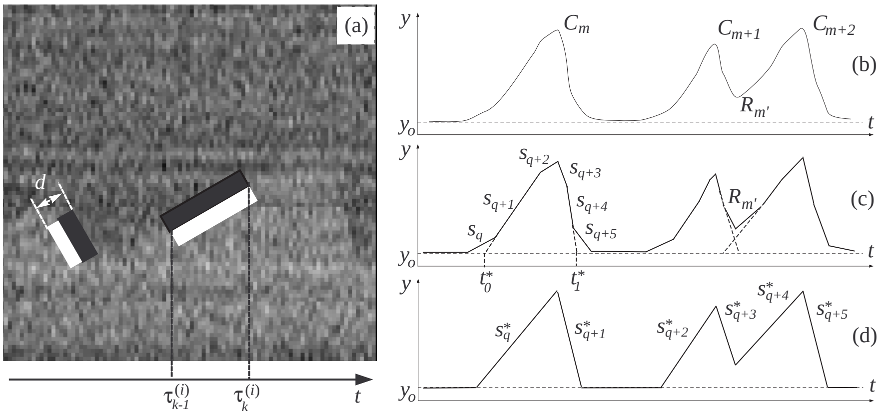

Fig. 5. The observation model used in the experiments, which compares the intensity distribution in two rectangular regions (black and white strips) and defines the likelihood of edge existence (a). Examples of applying the MTA to the extracted edge using the VRPF in order to compute the kinematic parameters (b-d).

increases – the likelihood decreases rapidly for small misalignments of the observation model with the edge. Usually this is a desirable property, because the edge can be located more precisely. The disadvantage of using only large l_v is the disability of the observation model to capture the fast motion transition stages.

Alternatively, the gradient image can be used as measurements for the VRPF, which represents the edges computed using the Gaussian derivatives. In this case, the pixel value at some position in the gradient image is the likelihood for finding the edge. Depending on the scale at which the derivatives are computed, the slopes of the tangent lines, which are related to the velocity values, can be accurately estimated, but only in regions having the same motion type. It can be seen from Fig. 4 that in the regions of the gradient image where catastrophes are present, the edge appearance is distorted – the transition between the growth and shrinkage is smoothed. This leads to a lowering of the angles of the tangent lines and, as a result, to underestimation of the velocity values. Due to the mentioned nonlinearity, this underestimation is especially severe for the shrinkage velocity.

In order to derive the MMSE estimator, the principle of fixed-lag smoothing is used, which greatly improves the final results. Here, the MMSE estimate of the state at time $t - \Delta t$ is computed using the posterior as distribution $p(\mathbf{x}_{0:\mathcal{N}_t^+}|\mathbf{z}_{0:t})$, that is

$$\hat{y}_{t-\Delta t} = \sum_{i=1}^{N_s} w_t^{(i)} h_t(\mathbf{x}_{N_{t-\Delta t}}). \tag{14}$$

In other words, the estimation of the edge position at time t is delayed until the measurements at time $t + \Delta t$ will be processed and the posterior updated.

E. Multiscale Trend Analysis

Having the estimated edge $\hat{y}_t$ after applying the VRPF, we employ multiscale trend analysis (MTA) [49] in order to automatically compute all the parameters of interest. At this stage of our analysis, it is necessary to detect all the catastrophe and rescue events and split the live history $\hat{y}_t$ into parts of growth and shrinkage, possibly separated by stages of no activity (state S_0),

The MTA was originally proposed for analysis of trends in time series and was recently successfully applied for analysis of MT transport in melanophores [50]. Compared to methods that try to construct an optimal piecewise linear approximation $L_\epsilon(t)$ with a minimal number of segments for a given error ϵ, the MTA builds a multi-level hierarchy of consecutively more detailed piecewise linear approximations of the analyzed time series at different scales. In general, it is not known beforehand which scale should be used for the analysis, but some prior knowledge about the application can significantly narrow down the range of levels that should be analyzed after applying MTA.

The following robust procedure was experimentally found to produce accurate estimates of the kinematic parameters using MTA. First, MTA decomposition is performed for a number of levels, $l = \{1, \ldots, N_L\}$, where N_L is a fixed (large) number. Each level in the decomposition can be represented with a set of nodes $\{s_q\}_{q=1}^{l}$ that partition $\hat{y}(t)$ on the interval $[0, T]$, where each node is given by four parameters, $(t_0^q, t_1^q, \alpha^q, \tilde{y}^q)$, and describes the linear approximation of $\hat{y}(t)$ on the interval $[t_0^q, t_1^q]$ with slope α^q and intercept $\tilde{y}^q = \hat{y}(t_0^q)$. In our implementation of MTA, the number of nodes (piecewise linear approximations) at level l is equal to l, and the first level ($l = 1$) is given by the base line $y = y_0$, where $y_0 = \min_t \hat{y}(t)$. At each level l, the number of catastrophes (local maxima in the approximation of $\hat{y}(t)$ at that level) $N_{\text{cat}}(l)$, is computed. Due to the nature of the signal $\hat{y}(t)$ and the way MTA works, for some range of hierarchy levels the number of catastrophes will stay constant ($dN_{\text{cat}}/dl = 0$). In general, the function $N_{\text{cat}}(l)$ is non-decreasing. By finding the maximum in the histogram of $\{N_{\text{cat}}(l) : l = \{1, \ldots, N_L\}\}$, which shows how many levels contain the same number of catastrophes, we can obtain the number of actual catastrophe events N_{cat}^*. From the set of levels $\{l_j\}$ that correspond to N_{cat}^* (satisfying $N_{\text{cat}}(l_j) = N_{\text{cat}}^*$), the median is selected, l^*, as the level for further parameter computations.

For the selected decomposition level and each catastrophe event C_m, $m = \{1, \ldots, N_{\text{cat}}^*\}$, which occurs at time t_m^c, the two sets of neighboring nodes, $\{s_q : t_{m-1}^c < t_0^q < t_m^c \cap \alpha^q > 0,\ q = 1, \ldots, l^*\}$ and $\{s_q : t_m^c < t_1^q < t_{m+1}^c \cap \alpha^q < 0,\ q = 1, \ldots, l^*\}$ are analyzed (see Fig. 5(c)), where $t_0^c = 0$ and $t_{N_{\text{cat}}^*+1}^c = T$. On both sides of the local maximum C_m, the nodes with the steepest slope α^q are selected and the linear approximations corresponding to those nodes are extrapolated until the intersection with $y = y_0$, giving the values $\tilde{t}_m^0$ and $\tilde{t}_m^1$. The rescue event $R_{m'}$ ($m' \in \mathbb{N}$) is detected between two catastrophes C_m and C_{m+1} if $\tilde{t}_m^1 > \tilde{t}_{m+1}^0$. In this case, the local minimum in the approximation of $\hat{y}(t)$ on the interval $[\tilde{t}_{m+1}^0, \tilde{t}_m^1]$ gives the position of the rescue, $t_{m'}^R$. Then, the approximation is recomputed for $\hat{y}(t)$ on the intervals $[\tilde{t}^0_m, t^c_m]$ and $[t^c_m, \tilde{t}^1_m]$. If the rescue event is positioned between two catastrophes C_m and C_{m+1}, the approximation is recomputed on the interval $[t^c_m, t^R_{m'}]$. The new approximation is given by a new set of nodes $S^* = \{s^*_q\}_{q=1}^{2N^*_{\text{cat}}}$ (see Fig. 5(d)), which is used to compute the kinematic parameters: the total growth and shrinkage times (T^+, T^-) and the corresponding velocity ($\hat{\nu}^+$, $\hat{\nu}^-$) and frequency f_{cat} and f_{res} estimates:

$$T^+ = \sum_{\substack{\forall s^*_q \in S^* \\ \alpha^q > 0}} (t^q_1 - t^q_0), \quad \hat{\nu}^+ = \frac{1}{T^+} \sum_{\substack{\forall s^*_q \in S^* \\ \alpha^q > 0}} (t^q_1 - t^q_0)\alpha^q, \tag{15}$$

$$T^- = \sum_{\substack{\forall s^*_q \in S^* \\ \alpha^q < 0}} (t^q_1 - t^q_0), \quad \hat{\nu}^- = \frac{1}{T^-} \sum_{\substack{\forall s^*_q \in S^* \\ \alpha^q < 0}} (t^q_1 - t^q_0)\alpha^q, \tag{16}$$

$$f_{\text{cat}} = N^*_{\text{cat}}/T^+, \quad f_{\text{res}} = N^*_{\text{res}}/T^-, \tag{17}$$

where N^*_{res} is the number of rescue events. In practice, the VRPF outputs a good piecewise linear approximation of the edges, so that the described procedure based on MTA runs robustly and accurately.

III. Experimental Results

The performance of the proposed VRPF-based method was evaluated using synthetic images (Section III-A) and real data from studies of MT dynamics *in vitro* (Section III-B) and in comparison with two other methods that were adapted for edge extraction in kymoimages. The synthetic images, for which the ground truth was available, were used to explore the accuracy and robustness of the method depending on the image quality (different SNR levels) and the parameter values that model the MT dynamics. The experiments on real data enabled us to compare the estimated kinematic parameters with manual analysis by expert biologists.

A. Evaluation on Synthetic Data

1) Simulation Step: The proposed technique was evaluated using computer generated kymoimages for different SNRs. The dynamics of the MT tip was simulated according to the model described in Section II-A (Fig. 1). The values of the model parameters were randomly generated each time the MT changes its state, by drawing a sample from the Gamma distribution, $\tau \sim \mathcal{G}(4, 1)$, and, depending on which state the MT was entering, the duration times were defined as $\tau^+ = 20\tau$, $\tau^- = 10\tau$, $\tau^0 = 10\tau$. The corresponding velocity values were drawn from the Gaussian distribution, $\nu^+ \sim \mathcal{N}(0.5, 0.005)$, $\nu^- \sim \mathcal{N}(-3, 0.005)$, $\nu^0 \sim \mathcal{N}(0, 0.05)$. These model values are representative of practical values.

Having the simulated dynamics $y(t)$, $0 < t < T$ (see Fig. 6(a) for an example), we created corresponding images of size $T \times Y$, where $T = 1000$ and $Y = \max_t y(t) + 2y_0$ for several SNR levels. Padding with a strip of size $T \times y_0$, $y_0 = 20$, was applied to the top and bottom of the image to avoid border problems when using the described rectangular observation model (Section II-D). The height

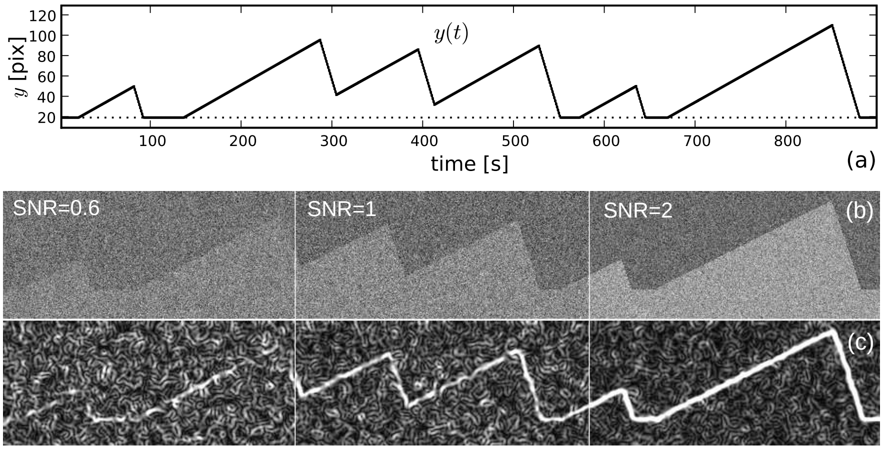

Fig. 6. Examples of the synthetic images used in the experiments. The simulated MT tip dynamics (a) is used to create the synthetic images for different SNR levels (b), for which the gradient images (c) are computed using the Gaussian derivatives at scale $\sigma_G = 3$.

of the generated images was in the range of 100–150 pixels, which corresponds to 8–12μm ($\Delta_t = 1s$ and $\Delta_y = 80$ nm). For all t, the background image intensity was set to $I_B = 100$ if $j > y(t) + y_0$ and the remainder of the image was set to the foreground intensity $I_F = I_B + \sigma\text{SNR}$, where $\sigma = 10$. To create the final noisy image, each pixel value was replaced with a random sample from the distribution $\mathcal{N}(I(t, j), \sigma^2)$. For the chosen values of I_B and $\sigma = 10$, this corresponds to the Poisson noise model, which is dominant in light microscopy images [3]. Examples of synthetic images for various SNRs are shown in Fig. 6(b). Again, for visual comparison, the edge information (the gradient magnitude) obtained using the Gaussian derivatives at scale $\sigma_G = 3$ is shown in Fig. 6(c).

The parameters of the described VRPF algorithm were fixed to the following values: $\nu^+ = 0.5$, $\nu^- = -3$, $\sigma^2_{\nu^0} = 0.5$, $\sigma^2_{\nu^+} = 0.05$, $\sigma^2_{\nu^-} = 0.5$, $V_{\text{th}} = 0.15$, $d = 6$, $\tau^0 = 3$, $\tau^1 = 10$, $\Delta t = 20$, $N_s = 500$, $N_L = 80$, $a = 0.01$, $\gamma = 10$. Since the ground truth was available in these experiments, the accuracy of extracting the edges was evaluated using a traditional quantitative performance measure: the root mean square error (RMSE) [51]:

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{T}|}\sum_{t\in\mathcal{T}}(y_t - \hat{y}_t)^2}, \tag{18}$$

where y_t defines the true position of the edge at time t, $\hat{y}_t$ is a MMSE estimate of y_t given by the VRPF, $\mathcal{T}$ is the set of time points for which the edge exists, and $|.|$ denotes the set size operator.

2) Results: The proposed VRPF method was evaluated using 20 synthetically generated images. Examples of edge extraction for SNR $= 0.6$ are shown in Fig. 7. In addition to the proposed VRPF, we also implemented two standard particle filters, denoted PF$_1$ and PF$_3$, in which the state transition process is synchronous with the measurement process (see Section II-D). PF$_1$ uses only one state transition model, $p(\mathbf{x}_t|\mathbf{x}_{t-1})$, which describes nearly-constant velocity motion [13]. To capture abrupt changes in the edges, the variance of the process noise in this transition model had to be made rather large.

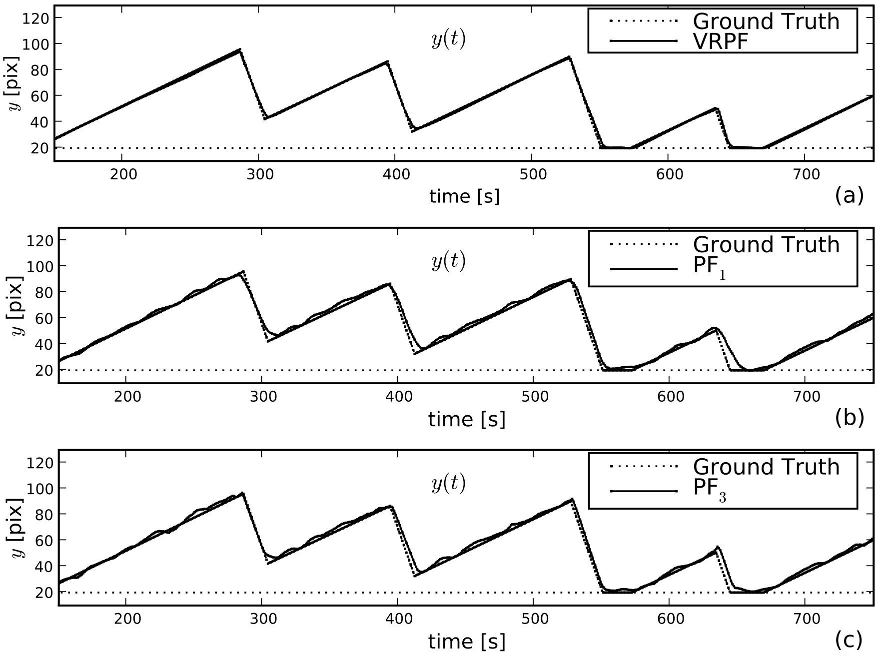

Fig. 7. Sample results of extracting the edge information from the noisy synthetic images using the proposed VRPF and two types of standard PFs in comparison with the ground truth.

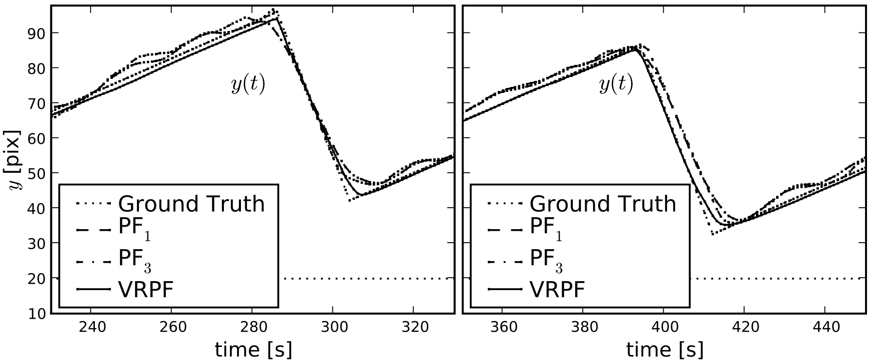

Fig. 8. More detailed results of extracting the edge information from the noisy synthetic images using the proposed VRPF and two types of standard PFs in comparison with the ground truth. The plots are zooms of the first two peaks in Fig. 7 and show the results combined.

Due to this high variance, the typical overshoots just after the catastrophe events (see Fig. 7(b)) highly corrupted the slope estimates, in particular the estimation of the shrinkage velocity. Additionally, for the low SNR image data, the filter frequently lost the edge and traced spurious background structures. PF$_3$ uses the same set of transition models as the VRPF. Contrary to the observation model used in the VRPF, however, a rectangular observation model of the same width d but fixed length $l_v = 5$ was used. The zoomed results in Fig. 8 clearly show that the edge $\hat{y}(t)$ estimated using the standard PFs is typically less smooth and piecewise linear.

The results of applying MTA for kinematic parameter estimation based on the edges extracted using PF$_3$ and VRPF are shown in Table I (results for PF$_1$ are not given here, since this filter frequently failed to find the edges at all, as indicated above). The RMSEs for both PF$_3$ and VRPF in finding the edge are approximately the same, but the velocity estimates computed using the linear approximation are different. This difference depends on the absolute value of the velocity, and for higher velocity values (especially the shrinkage velocity), VRPF is about 3–7% more accurate than PF$_3$. The results also show that prefiltering of the images does not improve the estimates significantly. This indicates that the observation model robustly estimates the mean intensities in the regions S_B and S_F even at very low SNRs. Prefiltering in this case worsens the estimation by blurring the already hardly visible edges before applying the VRPF.

3) Sensitivity Analysis: We also assessed the sensitivity of the proposed VRPF method to changes in the expected velocities. To this end, the parameter values ν^+ and ν^- were varied. It was observed that deviation of these parameters from the ground truth values decreased the accuracy of the method. In order to reduce the inaccuracy in the parameter estimation, which is caused by suboptimal initialization of the VRPF, the following procedure was used in practice. The proposed algorithm was executed iteratively, in a "bootstrapping" fashion. First, the initial velocity values ν^+ and ν^- are approximately specified, with large standard deviations σ_{ν^+} and σ_{ν^-}. After the first run, these parameters, which are still inaccurate but now closer to the optimal values, are reestimated using the MTA. Then, the algorithm is initialized with the new estimates and rerun. The number of iterations can be either user defined or specified by the desired RMSE between two edges estimated in consecutive iterations. In the experiments, we found that this approach always resulted in estimates in the range $(\nu \pm \sigma_\nu)$ defined by the ground truth.

In order to assess the robustness of the proposed iterative initialization procedure, we ran a number of iterations from different initial values of ν^+ and ν^-. In the most non-informative case, the estimates were fixed to $\nu^+ = 0$ and $\nu^- = 0$ and the standard deviations σ_{ν^+} and σ_{ν^-} were taken to be larger than 1. Typical convergence results of the velocity estimates to the ground truth values are shown in Fig. 9. After a short "burn-in" period of 1–3 first iterations, the velocity estimation is stabilized and every next run of the algorithm from the initial values given by the previous iteration does not improve the accuracy of the velocity estimates. Increasing the standard deviations σ_{ν^+} and σ_{ν^-} has influence only on the estimates during the "burn-in" period, and finally leads to the same ground truth velocity estimates. If the algorithm is initialized with the estimates that are close to the target values, the "burn-in" period is usually shorter – one or two iterations. The proposed VRPF method was also initialized with extreme initial values that would not be observable in practice: $\nu^+ = 5$, $\sigma_{\nu^+} = 1$, $\nu^- = -30$, $\sigma_{\nu^-} = 3$ and already after the first iteration, the obtained estimates were close to the target values, in the range $\nu^+ \in (0.3, 0.6)$ and $\nu^- \in (2, 3)$. In practice, for our biological application, we normally run the VRPF estimation for five iterations and the initial velocity estimates are fixed to $\nu^+ = 0$ and $\nu^- = 0$.

TABLE I

Results of parameter estimation in synthetically generated images of microtubule dynamics using MTA based on the edges extracted with different combinations of prefiltering and particle filtering methods. The velocity estimates are given in [pix/frame].

SNR	RMSE	$\nu^{+}\pm$ sd	$\nu^{-}\pm$ sd	f_{cat}	f_{res}
	Ground truth values				
-	-	0.50±0.005	-3.00±0.005	0.009	0.018
	VRPF without prefiltering				
0.4	2.54	0.47±0.07	-2.41±0.79	0.011	0.019
0.6	1.43	0.50±0.03	-3.03±0.61	0.009	0.018
0.8	1.23	0.49±0.02	-2.91±0.62	0.009	0.017
1.0	1.15	0.50±0.01	-2.96±0.37	0.009	0.017
1.2	0.96	0.49±0.01	-2.95±0.34	0.009	0.018
	VRPF with bilateral prefiltering				
0.4	2.01	0.48±0.07	-2.44±0.83	0.010	0.017
0.6	1.86	0.50±0.02	-2.86±0.40	0.009	0.015
0.8	1.64	0.49±0.02	-2.93±0.34	0.009	0.017
1.0	1.33	0.49±0.03	-3.05±0.36	0.009	0.017
1.2	1.25	0.49±0.02	-2.98±0.32	0.009	0.018
	VRPF with anisotropic diffusion prefiltering				
0.4	2.41	0.47±0.08	-2.14±0.56	0.010	0.019
0.6	2.55	0.49±0.08	-2.91±0.64	0.010	0.021
0.8	1.44	0.49±0.03	-2.98±0.39	0.009	0.018
1.0	1.13	0.49±0.02	-2.91±0.44	0.009	0.018
1.2	1.05	0.49±0.02	-2.91±0.34	0.009	0.018
	PF_3 without prefiltering				
0.4	2.72	0.47±0.08	-2.44±1.02	0.006	0.026
0.6	1.46	0.50±0.05	-2.71±0.92	0.011	0.014
0.8	1.12	0.50±0.05	-2.73±0.21	0.009	0.017
1.0	0.98	0.49±0.02	-2.81±0.27	0.009	0.015
1.2	1.02	0.49±0.02	-2.79±0.31	0.009	0.018

As can be seen from Table I and Fig. 9, for low SNR kymoimages (SNR < 0.6), the shrinkage velocity is always severely underestimated, which is also true for other methods (see Table II). The estimate of ν^{-} does not converge to the correct value even when the VRPF is initialized using the ground truth values for ν^{-} and ν^{+}. The error is larger for larger speed values due to the nonlinear error propagation during the slope estimation (as it was pointed out in Section II-B). At the same time, using the [pixel/frame] units for the velocity in the proposed VRPF gives the possibility to accurately estimate any velocity values from the real experiments, given in [μm/s]. This can be achieved by adjusting the

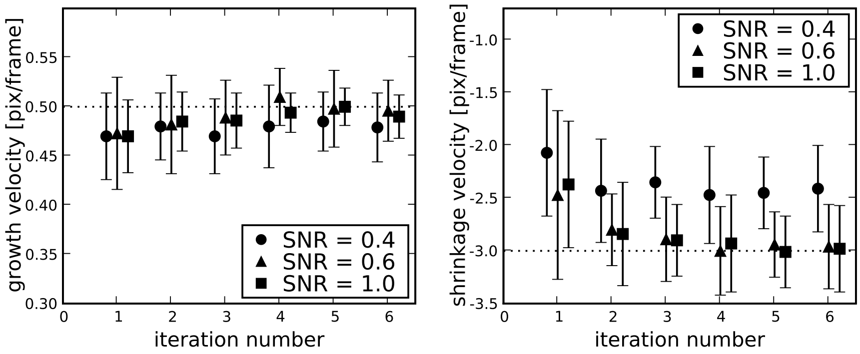

Fig. 9. Results of the velocity estimation using the proposed iterative initialization procedure. The initial values of $\nu^+ = \nu^- = 0$ and $\sigma_{\nu^+} = \sigma_{\nu^-} = 1$ are used for the initialization of the proposed algorithm (the first run). In each iteration, the obtained velocity estimates are used as initial values for the next iteration leading to the convergence of the estimates to the ground truth values after 3–4 iterations.

the parameters Δ_t and Δ_y during the imaging in order to obtain the velocity values ν (in [pixel/frame]) in the desired range (for example, for the shrinkage velocity $|\nu| \leq 3$) according to $\nu = \nu_m \Delta_t / \Delta_y$, where ν_m is the real velocity in [μm/s].

The described iterative initialization procedure breaks down for low SNR kymoimages in situations where the initial standard deviations σ_{ν^+} and σ_{ν^-} are too small (for example, in the case of our synthetic data, $\sigma_{\nu^+} < 0.05$ and $\sigma_{\nu^-} < 0.1$). In this case, the initialization is done outside the “capture range” in the parameter space and the iterative estimation will not converge to the correct (or even to the above mentioned underestimated) velocity values. This range of no convergence is typically very small and contains values of σ_{ν^+}, σ_{ν^-} that are much smaller than the values that we propose to use for the initialization. Within the Bayesian estimation framework, by specifying the initial velocity values with such high accuracy (small σ_{ν^-} and σ_{ν^+}), we ignore the observations (the information from the kymoimages) during the estimation procedure and rely mostly on the predictions from the prior transition model with $\nu^+ = 0$ and $\nu^- = 0$. In such a case, the initial velocity estimates will never be corrected by the observations and the edge will not be found. On the other hand, by specifying the initial velocity values inaccurately (large σ_{ν^-} and σ_{ν^+}), we take into account the observations more than the predictions from the motion model, and this way we can capture at least some parts of the edge correctly (which is enough for the first iteration) and approximately estimate the velocity values (which will have smaller and more realistic σ_{ν^-} and σ_{ν^+}) for the next iteration.

4) Comparison With Other Methods: The performance of the proposed VRPF method was also compared with the results of edge extraction by two other methods that were adapted for the kymograph analysis. In those experiments, the parameters of the considered methods were manually adjusted to obtain the best performance. First, we used a freely available plugin for ImageJ (National Institutes of Health, Bethesda, MD [52]), NeuronJ, which implements an interactive neurite tracing technique [53].

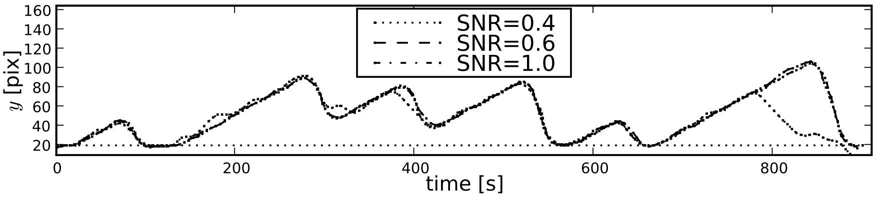

Fig. 10. Examples of edge extraction using NeuronJ applied to the gradient magnitude images obtained from the synthetic image data (see also Fig. 6) with different SNR levels.

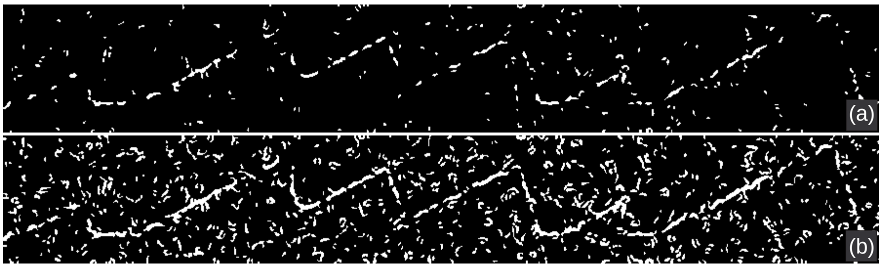

Fig. 11. Results of edge extraction from the synthetically generated kymoimage (SNR =0.6) using the Canny edge detector for two different values of the hysteresis thresholds.

NeuronJ was originally developed for accurate detection and tracing of individual neurites in 2D fluorescence microscopy images of cells in culture. It is capable of tracing bright thin elongated structures superimposed on dark, noisy background by computing for each pixel in the image a measure based on the eigenvalues of the second-order derivative matrix computed from the image intensities around that pixel [53]. Such measure reflects how likely the image intensities around that pixel resemble an elongated structure and is used as a cost function during the search for the optimal path that connects two user-specified points in the image, where “optimal” means having a globally minimal cumulative cost according to a predefined function. The algorithm can deal with very low contrast and possible gaps along the bright structure of interest. In our experiments as an input for NeuronJ, we used the gradient magnitude images obtained using the Gaussian derivatives at scale $\sigma_G = 3$, in which edges appear as neurite- or vessel-like structures (see Fig. 6(c)). The algorithm usually failed to find the correct edge if smaller scales ($\sigma_G < 3$) were used. The “shortest” paths obtained between two points, which were manually specified in the beginning and the end of the kymoimage, were analyzed using the MTA. The examples of edge extraction are shown Fig. 10. The results of the analysis using the same set of synthetic images as in the case of the VRPF method are presented in Table II. It can be seen that especially the shrinkage velocity estimates are far inferior to those of the VRPF method. For visual comparison, the edge information extracted using the Canny edge detector [39] for two different values of hysteresis thresholds is shown in Fig. 11.

The second considered approach for edge extraction is based on labeling the kymograph intensities into two classes (MT body and background), followed by extraction of the boundary that separate those

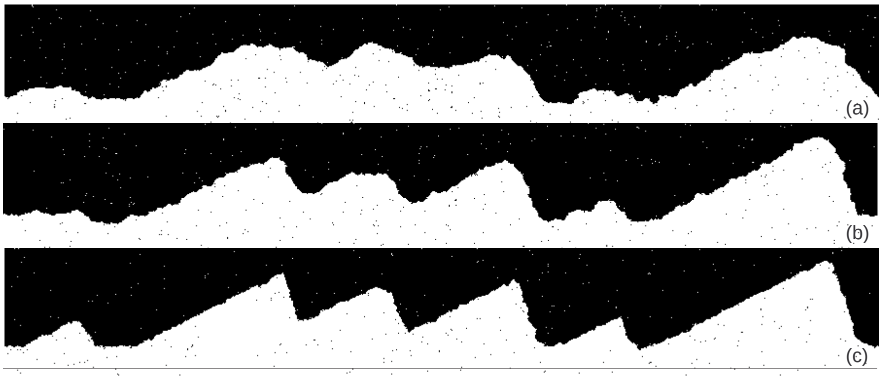

Fig. 12. Examples of segmentation results obtained using Graph Cuts applied to the synthetically generated images (see also Fig. 6) with SNR=0.4 (a), SNR=0.6 (b) and SNR=1 (c).

two regions. For this purpose we used a graph cut technique (further abbreviated as Graph Cuts), which is a robust method for globally optimal image segmentation that recently became popular and was shown to outperform similar energy minimization-based segmentation and classification methods [54]–[56]. For segmentation purposes, the method optimally minimizes the energy that consists of two terms: the data dependent term that evaluates the penalty for assigning a particular pixel to a given label and the smoothness term that evaluates the penalty for assigning two neighboring pixels to different regions. The incorporation of prior information about the relation between the neighboring pixels using the smoothness term improves the quality of the segmentation drastically compared to classifiers that do not employ Markov Random Fields (see [54], [55], [57] for more information on this subject). We employed the implementation of graph cut segmentation technique with a 8-neighbor system also used in [56] (for more details see [55] and [56]). Examples of the segmentation results are shown in Fig. 12. The velocity estimates obtained by the MTA after the edge extraction from 20 synthetically generated kymoimages using Graph Cuts are shown in Table II. The results of the segmentation show that it is very difficult to accurately separate MTs from background for low SNR kymoimages. Even for SNR levels around 1, where the proposed VRPF method has no difficulties to accurately estimate the kinematic parameters, the small inaccuracies in the edge extraction (Fig. 12(c)) led to unacceptable (about 30-50%) relative errors in shrinkage velocity estimates (see Table II).

5) Semi-Automatic Kymograph Construction Procedure: In order to apply the proposed VRPF technique, the kymographs should be constructed from the original noisy DIC image sequences (see Fig.13 for an example). In practice, biologists manually select relevant MTs for further analysis. This selectiveness is necessary for several reasons: in a movie that contains a number of growing and shrinking MTs some of them might be only partially visible (only one end, see Fig.13, MT number 3) due to the limited field of view, some might display no dynamic activity, or be positioned perpendicular to the DIC shading gradient and because of that poorly imaged compared to other MTs in the same movie. For these reasons, only a small portion of all visible MTs (typically about 10-30%) is included in

TABLE II

RESULTS OF PARAMETER ESTIMATION IN SYNTHETICALLY GENERATED IMAGES OF MICROTUBULE DYNAMICS USING MTA BASED ON THE EDGES EXTRACTED USING NEURONJ AND GRAPH CUTS. THE VELOCITY ESTIMATES ARE GIVEN IN [PIX/FRAME].

	NeuronJ		Graph Cuts	
SNR	$\nu^{+} \pm$ sd	$\nu^{-} \pm$ sd	$\nu^{+} \pm$ sd	$\nu^{-} \pm$ sd
0.4	0.47±0.05	-0.93±0.63	0.48±0.11	-0.78±0.32
0.6	0.49±0.04	-1.59±0.53	0.47±0.05	-1.52±0.47
1.0	0.49±0.03	-1.62±0.37	0.49±0.04	-2.15±0.41

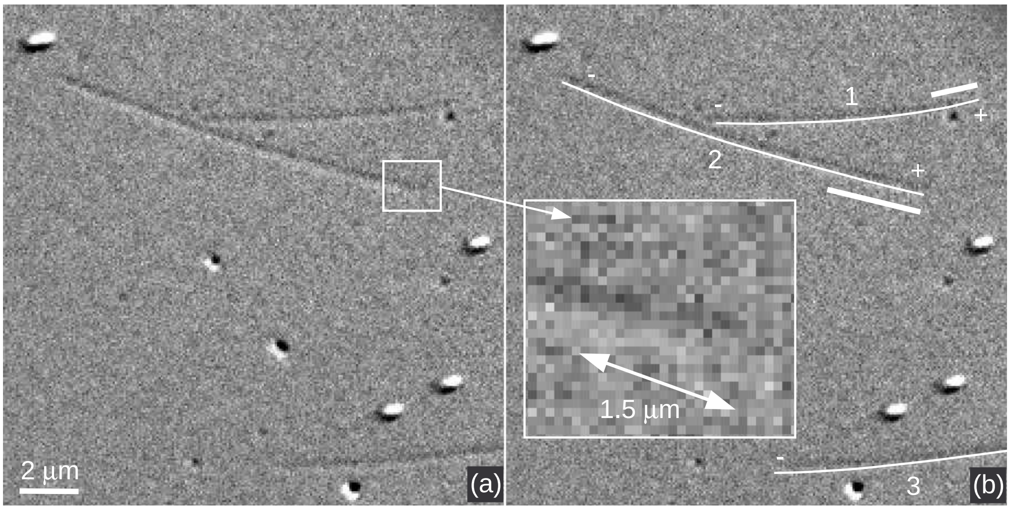

Fig. 13. Example of a single frame from typical DIC microscopy image sequence in our real experiments (a). Three MTs (number 1, 2 and 3, with indicated (+) or (-) ends), marked using the observation curves obtained by the semi-automatic kymograph construction procedure and manually truncated (for visualization purposes) to cover only the MT body (b). The thick white lines (close to the MTs) indicate the range of MT motion. The inset shows the zoomed MT tip and demonstrates the difficulty of finding the exact position of the MT end (in a single image), which is located somewhere in the region of about 1.5μm, indicated by the corresponding arrows.

the analysis. Automatic construction of kymographs for all visible MTs and the further analysis would always require a manual correction/selection procedure applied to the final estimates, which makes this approach impractical. Hence, such selection is usually done in the very beginning and kymographs are built only for a specified set of MTs.

As was described in Section II-B, the kymographs are built by defining an observation line L along the bright part of a MT in the image sequence and sampling the image intensities in time along L. In order to reduce the noise in the kymograph, the averaging of n_l pixel values along the direction perpendicular to L is used (we say that the observation line has ”width” n_l). As an example, the resulting kymographs, constructed using a straight observation line along MT number 2 in Fig.13, for several line widths n_l, are shown in Fig. 14(a-d). The averaging improves the image quality up to some extent (Fig. 14(b)), but if the line width is too large, the averaging of bright and dark parts of the MT diminishes the benefits (Fig. 14(d)). Additionally, it is possible that the straight line L will not cover the MT body properly for MTs that are slightly bent. In practice, such bending does not cause any problems, because most of the time biologists do not try to cover the whole MT body with such line, but only the relevant dynamic parts (see the thick white lines in Fig.13 that indicate the MT motion range in the whole image sequence), which indeed can be accurately approximated by a straight line.

In order to deal with potential frame-to-frame drifts of MT bodies in the image sequence, we devised the following semi-automatic procedure for kymograph construction, based on our previous work on particle filtering tracking of vessel-like structures in medical imaging [47], [48]. In order to find the observation curve L_c (see Fig.13, the white curves along the MTs) rather than a straight line, the 2D particle filter is used (further abbreviated as the 2DPF), which is capable of tracing the bright ridge of MT body if the appropriate observation model is used. The original observation model employed in [47] was designed for bright elongated structures on dark background. In order to deal with the black-and-white appearance of MTs in the DIC microscopy imaging, we used a 2D observation model, which is similar in appearance to a 2D Gabor filter, but which is a 1D Gaussian function (with the variance of 20-50) in a cross-section along the MT direction and a derivative of a 1D Gaussian function (with a variance of 5, which reflects the width of the MT appearance in the DIC microscopy imaging) in the perpendicular cross-section direction. The direction is automatically estimated by the 2DPF during the tracing, along with the location of the bright ridge. The 2DPF is manually initialized by clicking on the MT body in one of the frames, for example on the ”seed”, which is visible in every frame. The initial MT direction for that location is obtained by using the Hough transform in a small square image region around the specified point. The 2DPF runs both ways (in opposite directions) starting from the initial location until the image border is encountered. This way, two kymographs (for plus and minus MT ends) are constructed and can be further analyzed using the proposed VRPF method. The resulting kymographs built for MT number 2 (Fig.13) are shown in Fig. 14(e-h). It can be observed that due to the fact that the observation curves L_c follow the bright intensity ridges more accurately than the straight observation lines, the width n_l in this case can be safely chosen between 1 and 3 pixels, producing kymoimages of optimal image quality. In practice, we run the 2DPF for every frame, which produces a set of observation curves L_c that are used to form a kymoimage by taking the image intensities in every frame along the corresponding observation curve. For our application, the interframe displacement of those curves was much smaller than one pixel, again confirming that the straight observation lines are quite appropriate.

The 2DPF, which performs the estimation within the Bayesian framework, can accurately trace the MTs in a single frame, due to the combination of the measurements (information from the images)

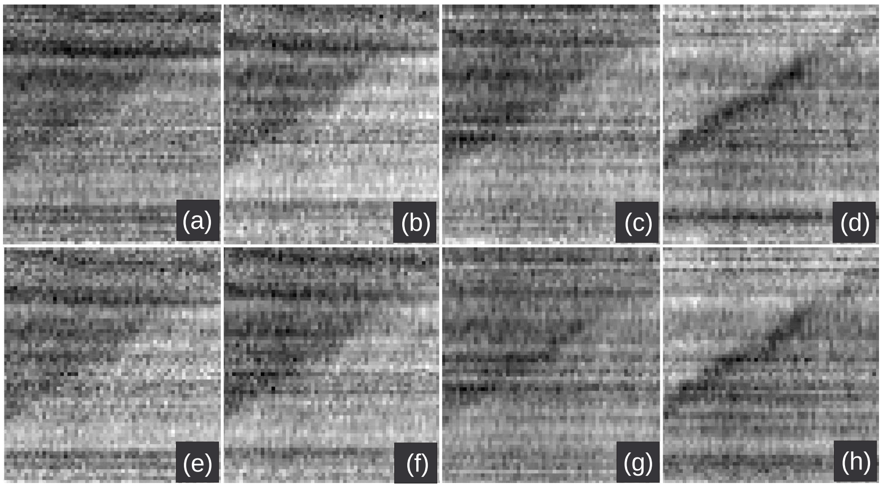

Fig. 14. Example of kymographs obtained for MT number 2 (see Fig. 13(b)) using a straight observation line L (a, b, c and d) and observation curves L_c obtained by the semi-automatic kymograph construction procedure (e, f, g and h) for different width n_l =1, 3, 5, and 7 pixels, respectively for both cases.

and the specified prior information about the MTs – on a small scale they are strictly linear structures. With this approach, even for extremely low SNR images we achieve accurate estimation (see more on validation of a similar 2DPF in [47], [48]). When the 2DPF reaches the MT tip, it continues tracing through the background until it faces the image border. During that period of the estimation there are no measurements from the MT in the image (only background), and because of that only the transition prior is used, which effectively extrapolates the observation curve from the MT tip to the image border. As a result, for that part of the observation curve we get a smooth and rather straight line, which is not critical, because the image intensities sampled over that part are coming from the background anyway. At the same time, it is very difficult to locate the exact location of the MT tip (the border where the MT ends) in a single 2D image. As can be observed from Fig.13 (see the inset) the MT end in this case is located somewhere in the region of at least 1μm, and once again, the frame-to-frame displacement of the MT during the growth event is less than 0.05μm. This demonstrates the fact that with the error in locating the MT tip of about 1 μm it is impossible to accurately estimate the velocity as interframe displacement over time and that is the main reason why we developed and applied the advanced VRPF method.

B. Evaluation on Real Data

For the validation of the VRPF method on real data we collected kymographs from three representative DIC microscopy image sequences acquired to study the influence of different concentrations of EB3 (end-binding protein 3) and GFP-EB3 (EB3 fused to the green fluorescent protein) on the MT growth and shrinkage velocities (ν^+ and ν^-) and the catastrophe rate (f_{cat}). The sequences were taken from experiments with MT nucleation from stable tubulin seeds, where 15μM of tubulin was added

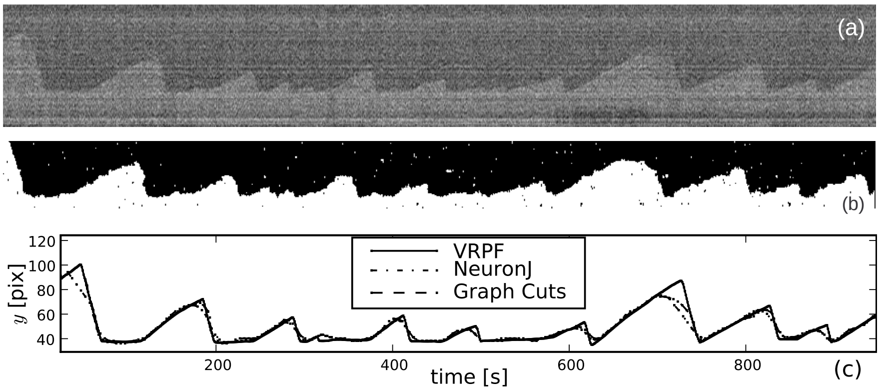

Fig. 15. Example of a kymograph generated in the experiments on real DIC microscopy image data with SNR ≈ 1 (a), the results of segmentation using Graph Cuts (b) and the results of edge estimation by Graph Cuts, NeuronJ and the proposed VRPF (c).

(Experiment I), or, in addition, 1μM of EB3 (Experiment II), or 1μM of GFP-EB3 (Experiment III), and from each sequence 10–20 kymographs were built manually [58]. The image sequences contained about 1000–1200 frames (one per second) of size 700×500 pixels (of size 86×86 nm^2). To estimate the parameters of interest, for each experiment 10 kymographs were analyzed manually and using the proposed VRPF method. The results are presented in Table III, where the speed estimates are also converted to μm/min. The usage of these units is more common in biological experiments and it also allows straightforward comparison with our recently published results [58]. An example of edge extraction using VRPF in real data together with the edges obtained using the NeuronJ and Graph Cuts are shown in Fig. 15(c). As an example, the velocity estimates for the kymograph shown in Fig. 15(a) obtained by all the considered methods are as follows (in the form $(\nu^+ \pm \sigma_{\nu^+},\ \nu^- \pm \sigma_{\nu^-})$): manually $(0.47 \pm 0.05,\ -2.79 \pm 0.42)$, VRPF $(0.48 \pm 0.06,\ -2.82 \pm 0.36)$, NeuronJ $(0.31 \pm 0.13,\ -1.23 \pm 0.35)$, Graph Cuts $(0.41 \pm 0.08,\ -2.03 \pm 0.43)$. Application of a paired Student t-test revealed no statistically significant difference between the estimates obtained by manual and VRPF-based analysis (p-values $\gg 0.05$ in all cases), suggesting that the proposed automatic method can replace the laborious manual procedures.

IV. CONCLUSIONS

In this paper we have proposed a new approach for the automatic analysis of *in vitro* microtubule dynamics imaged using time-lapse differential interference contrast microscopy. It is based on a transformation of the 2D image sequences into kymographs (space-time images) for each microtubule along a corresponding observation line. By using this representation, the task of tracking microtubule tips on a per-frame basis in the noisy images, which from our previous work is known to be a difficult and error-prone problem, is replaced by a segmentation of spatiotemporal structures (edges in our case). For the extraction of these structures from the kymographs, we have proposed a variable-rate particle

TABLE III

RESULTS OF PARAMETER ESTIMATION IN REAL DIC MICROSCOPY IMAGE DATA SETS USING MANUAL ANALYSIS VERSUS VRPF.

	$\nu^+ \pm$ sd [pix/frame]	$\nu^- \pm$ sd [pix/frame]	f_{cat}	ν^+ [μm/min]	ν^- [μm/min]
	Experiment I (pure tubulin)				
Manual	0.19±0.04	-2.06±0.43	0.0021	0.56	-10.63
VRPF	0.17±0.07	-1.89±0.52	0.0020	0.51	-9.72
	Experiment II (tubulin and EB3)				
Manual	0.52±0.05	-2.78±0.65	0.0133	2.68	-14.34
VRPF	0.49±0.07	-2.84±0.51	0.0141	2.52	-14.65
	Experiment II (tubulin and GFP-EB3)				
Manual	0.49±0.08	-2.88±0.41	0.0132	2.52	-14.86
VRPF	0.50±0.06	-2.72±0.50	0.0145	2.58	-14.03

filtering method, which is better capable of dealing with abrupt changes than standard particle filtering methods. The method is built within a Bayesian framework and optimally combines the measurements and prior knowledge about the underlying (imaging and motion) processes. For the estimation of important kinematic parameters from the extracted edges, we have adopted multiscale trend analysis.

The quantitative evaluation of the proposed method was done using realistic synthetic images as well as real microscopy image data from biological experiments. From the results of the experiments on synthetic data, where the ground truth of the microtubule tip position was available, it was concluded that the method is capable of accurate estimation of the important kinematic parameters. Moreover, it was concluded that the method is more robust and more accurate than standard particle filtering methods and several other advanced edge extraction techniques that we compared with. For the real data, the proposed method was compared to manual analysis carried out by expert biologists. The results of this comparison clearly demonstrated that the automatically estimated parameters are in good agreement with the results obtained manually. Together, these observations lead to the conclusion that the proposed method can replace laborious manual analyses.

ACKNOWLEDGMENTS

The authors thank Laura Munteanu, Susana Montenegro Gouveia, and Marileen Dogterom for their help with the acquisition of DIC images of microtubules. Ihor Smal and Erik Meijering are supported by the Netherlands Organization for Scientific Research (NWO) through VIDI-grant 639.022.401. Ilya Grigoriev and Anna Akhmanova are supported by the NWO programs ALW-VICI and ZonMw TOP.

REFERENCES

[1] P. Bieling, L. Laan, H. Schek, L. E. Munteanu, L. Sandblad, M. Dogterom, D. Brunner, and T. Surrey, "Reconstitution of a microtubule plus-end tracking system *in vitro*," *Nature*, vol. 450, no. 7172, pp. 1100–1105, 2007.

[2] E. L. Munteanu, "Dynamics and regulation at the tip: A high resolution view on microtubule assembly," Ph.D. dissertation, Institute for Atomic and Molecular Physics (AMOLF), Amsterdam, 2008.

[3] Q. Wu, F. A. Merchant, and K. R. Castleman, *Microscope Image Processing*. Burlington, MA: Elsevier Academic Press, 2008.

[4] D. J. Stephens and V. J. Allan, "Light microscopy techniques for live cell imaging," *Science*, vol. 300, no. 5616, pp. 82–86, 2003.

[5] J. B. Pawley, *Handbook of Biological Confocal Microscopy*, 3rd ed. New York: Springer, 2006.

[6] D. Thomann, D. R. Rines, P. K. Sorger, and G. Danuser, "Automatic fluorescent tag detection in 3D with super-resolution: Application to the analysis of chromosome movement," *Journal of Microscopy*, vol. 208, no. 1, pp. 49–64, 2002.

[7] D. Thomann, J. Dorn, P. K. Sorger, and G. Danuser, "Automatic fluorescent tag localization II: Improvement in super-resolution by relative tracking," *Journal of Microscopy*, vol. 211, no. 3, pp. 230–248, 2003.

[8] A. Genovesio, T. Liedl, V. Emiliani, W. J. Parak, M. Coppey-Moisan, and J.-C. Olivo-Marin, "Multiple particle tracking in 3-D+t microscopy: Method and application to the tracking of endocytosed quantum dots," *IEEE Transactions on Image Processing*, vol. 15, no. 5, pp. 1062–1070, 2006.

[9] I. F. Sbalzarini and P. Koumoutsakos, "Feature point tracking and trajectory analysis for video imaging in cell biology," *Journal of Structural Biology*, vol. 151, no. 2, pp. 182–195, 2005.

[10] K. Jaqaman, D. Loerke, M. Mettlen, H. Kuwata, S. Grinstein, S. Schmid, and G. Danuser, "Robust single-particle tracking in live-cell time-lapse sequences," *Nature Methods*, vol. 5, no. 8, pp. 695–702, 2008.

[11] E. Meijering, I. Smal, and G. Danuser, "Tracking in molecular bioimaging," *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 46–53, 2006.

[12] E. Meijering, I. Smal, O. Dzyubachyk, and J.-C. Olivo-Marin, "Time-lapse imaging," in *Microscope Image Processing*, Q. Wu, F. A. Merchant, and K. R. Castleman, Eds. Burlington, MA: Elsevier Academic Press, 2008, pp. 401–440.

[13] I. Smal, K. Draegestein, N. Galjart, W. Niessen, and E. Meijering, "Particle filtering for multiple object tracking in dynamic fluorescence microscopy images: Application to microtubule growth analysis," *IEEE Transactions on Medical Imaging*, vol. 27, no. 6, pp. 789–804, 2008.

[14] I. Smal, E. Meijering, K. Draegestein, N. Galjart, I. Grigoriev, A. Akhmanova, M. E. van Royen, A. B. Houtsmuller, and W. Niessen, "Multiple object tracking in molecular bioimaging by Rao-Blackwellized marginal particle filtering," *Medical Image Analysis*, vol. 12, no. 6, pp. 764–777, 2008.

[15] D. Sage, F. R. Neumann, F. Hediger, S. M. Gasser, and M. Unser, "Automatic tracking of individual fluorescence particles: Application to the study of chromosome dynamics," *IEEE Transactions on Image Processing*, vol. 14, no. 9, pp. 1372–1383, 2005.

[16] S. Bonneau, M. Dahan, and L. D. Cohen, "Single quantum dot tracking based on perceptual grouping using minimal paths in a spatiotemporal volume," *IEEE Transactions on Image Processing*, vol. 14, no. 9, pp. 1384–1395, 2005.

[17] Y. Sato, J. Chen, R. Zoroofi, N. Harada, S. Tamura, and T. Shiga, "Automatic extraction and measurement of leukocyte motion in microvessels using spatiotemporal image analysis," *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 4, pp. 225–236, 1997.

[18] B. Hinz, W. Alt, C. Johnen, V. Herzog, and H.-W. Kaiser, "Quantifying lamella dynamics of cultured cells by SACED, a new computer-assisted motion analysis," *Experimental Cell Research*, vol. 251, no. 1, pp. 234–243, 1999.

[19] J. C. Bulinski, D. J. Odde, B. J. Howell, T. D. Salmon, and C. M. Waterman-Storer, "Rapid dynamics of the microtubule binding of ensconsin in vivo," *Journal of Cell Science*, vol. 114, no. 21, pp. 3885–3897, 2001.

[20] J. E. Bear, T. M. Svitkina, M. Krause, D. A. Schafer, J. J. Loureiro, G. A. Strasser, I. V. Maly, O. Y. Chaga, J. A. Cooper, G. G. Borisy, and F. B. Gertler, “Antagonism between Ena/VASP proteins and actin filament capping regulates fibroblast motility,” *Cell*, vol. 109, no. 4, pp. 509–521, 2002.

[21] V. Racine, M. Sachse, J. Salamero, V. Fraisier, A. Trubuil, and J. B. Sibarita, “Visualization and quantification of vesicle trafficking on a three-dimensional cytoskeleton network in living cells,” *Journal of Microscopy*, vol. 225, pp. 214–228, 2007.

[22] A. Desai and T. J. Mitchison, “Microtubule polymerization dynamics,” *Annual Review of Cell and Developmental Biology*, vol. 13, pp. 83–117, 1997.

[23] E. Nogales, M. Whittaker, R. Milligan, and K. Downing, “High-resolution model of the microtubule,” *Cell*, vol. 96, no. 1, pp. 79–88, 1999.

[24] J. Howard and A. A. Hyman, “Dynamics and mechanics of the microtubule plus end,” *Nature*, vol. 422, no. 6933, pp. 753–758, 2003.

[25] T. Mitchison and M. Kirschner, “Dynamic instability of microtubule growth,” *Nature*, vol. 312, no. 5991, pp. 237–242, 1984.

[26] S. C. Schuyler and D. Pellman, “Microtubule “plus-end-tracking proteins”: The end is just the beginning,” *Cell*, vol. 105, no. 4, pp. 421–424, 2001.

[27] P. Carvalho, J. S. Tirnauer, and D. Pellman, “Surfing on microtubule ends,” *Trends in Cell Biology*, vol. 13, no. 5, pp. 229–237, 2003.

[28] A. Akhmanova and M. O. Steinmetz, “Tracking the ends: A dynamic protein network controls the fate of microtubule tips,” *Nature Reviews Molecular Cell Biology*, vol. 9, no. 4, pp. 309–322, 2008.

[29] A. Akhmanova and C. C. Hoogenraad, “Microtubule plus-end-tracking proteins: Mechanisms and functions,” *Current Opinion in Cell Biology*, vol. 17, no. 1, pp. 47–54, 2005.

[30] R. Walker, E. O’Brien, N. Pryer, M. Soboeiro, W. Voter, H. Erickson, and E. Salmo, “Dynamic instability of individual microtubules analyzed by video light microscopy: Rate constants and transition frequencies,” *Journal of Cell Biology*, vol. 107, no. 4, pp. 1437–1448, 1988.

[31] D. Murphy, *Fundamentals of Light Microscopy and Digital Imaging*. New York: Wiley-Liss, 2001.

[32] N. Caudron, I. Arnal, E. Buhler, D. Job, and O. Valiron, “Microtubule nucleation from stable tubulin oligomers,” *Journal of Biological Chemistry*, vol. 227, no. 52, pp. 50 973–50 979, 2002.

[33] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, 2nd ed. Pacific Grove: PWS Publishing, 1999.

[34] Y. S. Fong, C. A. Pomalaza, and X. H. Wang, “Comparison study of nonlinear filters in image processing applications,” *Optical Engineering*, vol. 28, no. 7, pp. 749–760, 1989.

[35] C. Tomasi and R. Manduchi, “Bilateral filtering for gray and color images,” in *Proceedings of the Sixth International Conference on Computer Vision*, 1998, pp. 839–846.

[36] D. Comaniciu and P. Meer, “Mean shift: A robust approach toward feature space analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pp. 603–619, 2002.

[37] P. Perona and J. Malik, “Scale-space and edge detection using anisotropic diffusion,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–639, 1990.

[38] B. M. ter Haar Romeny, *Front-End Vision and Multi-Scale Image Analysis*. Berlin: Springer, 2003.

[39] J. F. Canny, “A computational approach to edge detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.

[40] S. M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.

[41] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston: Artech House, 2004.

[42] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Berlin: Springer-Verlag, 2001.

[43] S. Godsill and J. Vermaak, "Variable rate particle filters for tracking applications," in *Proceedings of the IEEE/SP 13th Workshop on Statistical Signal Processing*, 2005, pp. 1280–1285.

[44] S. Godsill, J. Vermaak, K.-F. Ng, and J.-F. Li, "Models and algorithms for tracking of manoeuvring objects using variable rate particle filters," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 925–952, 2007.

[45] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: Wiley, 2001.

[46] N. M. Haan and S. J. Godsill, "A time-varying model for DNA sequencing data submerged in correlated noise," in *Proceedings of the IEEE Workshop on Statistical Signal Processing*, August 2001.

[47] M. Schaap, I. Smal, C. Metz, T. van Walsum, and W. J. Niessen, "Bayesian tracking of elongated structures in 3D images," in *Proceedings of the 20th International Conference on Information Processing in Medical Imaging*, 2007, pp. 74–85.

[48] M. Schaap, R. Manniesing, I. Smal, T. van Walsum, A. van der Lugt, and W. J. Niessen, "Bayesian tracking of tubular structures and its application to carotid arteries in CTA," in *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, 2007, pp. 562–570.

[49] I. Zaliapin, A. Gabrielov, and V. Keilis-Borok, "Multiscale trend analysis," *Fractals*, vol. 12, no. 3, pp. 275–292, 2004.

[50] I. Zaliapin, I. Semenova, A. Kashina, and V. Rodionov, "Multiscale trend analysis of microtubule transport in melanophores," *Biophysical Journal*, vol. 88, no. 6, pp. 4008–4016, 2005.

[51] W. Ng, J. Li, S. Godsill, and J. Vermaak, "A hybrid approach for online joint detection and tracking for multiple targets," in *Proceedings of the 2005 IEEE Aerospace Conference*, 2005, pp. 2126–2141.

[52] W. S. Rasband, "ImageJ: Image processing and analysis in Java," http://rsb.info.nih.gov/ij/.

[53] E. Meijering, M. Jacob, J.-C. F. Sarria, P. Steiner, H. Hirling, and M. Unser, "Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images," *Cytometry*, vol. 58A, no. 2, pp. 167–176, 2004.

[54] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1222–1239, 2004.

[55] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124–1137, 2004.

[56] J. Malcolm, Y. Rathi, and A. Tannenbaum, "Multi-object tracking through clutter using graph cuts," in *Proceedings of the IEEE International Conference on Computer Vision*, 2007, pp. 1–5.

[57] G. Winkler, *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction (Stochastic Modelling and Applied Probability)*. New York: Springer-Verlag, 2006.

[58] Y. Komarova, C. O. de Groot, I. Grigoriev, S. M. Gouveia, E. L. Munteanu, J. M. Schober, S. Honnappa, R. M. Buey, C. C. Hoogenraad, M. Dogterom, G. G. Borisy, M. O. Steinmetz, and A. Akhmanova, "Mammalian end binding proteins control persistent microtubule growth," *Journal of Cell Biology*, vol. 184, no. 5, pp. 691–706, 2009.