

Microtubule Dynamics Analysis using Kymographs and Variable-Rate Particle Filters

It is a truth very certain that when it is not in our power to determine what is true we ought to follow what is most probable.

— RENÉ DESCARTES,
Discours de la Méthode (1637)

Abstract — Studying the dynamics of intracellular objects is of fundamental importance in understanding healthy life at the molecular level and to develop 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 chapter, 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*.

5.1 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 of one or more of these processes, their study is of interest not only to life scientists, but also to pharmaceutical companies in their attempts 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 [14, 101].

Intracellular dynamics is usually visualized using advanced microscopy imaging techniques, such as fluorescence confocal microscopy, where the objects of interest are labeled with fluorescent proteins. Alternatively, differential interference contrast (DIC) microscopy can sometimes be used, which does not require labeling [112, 156]. 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 levels of measurement noise [112, 185]. 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 [52, 71, 95, 96, 132, 160, 161] 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 [141, 142], 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 [17, 128].

In this chapter, 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 [12, 22, 64, 130]. 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.

The chapter is organized as follows. In Section 5.2, we describe the biological application considered in this chapter and the proposed methods to model, acquire, transform, preprocess, and analyze the image data. In Section 5.3, 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 5.4.

5.2 Methods

5.2.1 *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* [37, 107]. *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 MT associated proteins (MAPs) [67]. 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 [99]. 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 an order of magnitude 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) [99]. 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 [101].

Recent studies reveal a special class of MAPs, plus-end-tracking proteins (+TIPs), that are able to accumulate at MT growing ends [6, 27, 67, 137]. The mechanisms by which +TIPs recognize MT ends have attracted much attention and several explanations have been proposed [4, 27, 67]. One way to understand the mechanism employed by individual +TIPs and the molecular mechanisms underlying 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




Figure 5.1. Dynamics model describing microtubule behavior *in vitro*.

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 [182]. 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. 5.1): 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: $S_0 \rightarrow G$ (the MT starts to grow), $G \rightarrow S$ (catastrophe), $S \rightarrow G$ (rescue), and $S \rightarrow S_0$ (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 5.3.1), 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 ($S \rightarrow G$) 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_0 should be replaced with a state that corresponds to a “pause” event [37], and all the transitions (arrows in Fig. 5.1) should be bidirectional.

5.2.2 Imaging Technique and Kymographs

In our studies, the dynamic behavior of MTs is imaged using DIC microscopy [102], which is effectively used for biological specimens that cannot be visualized with sufficient contrast using bright-field microscopy. The resulting images (see Fig. 5.2 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. DIC microscopy works by separating a polarized light source into two beams that take slightly different paths through the sample and then converting changes in optical path length to a visible change in brightness [102]. The advantages of DIC over fluorescence microscopy is that the samples do not have to be stained. The main limitation of this imaging technique is its requirement for a thin and transparent sample of fairly similar refractive index to its surroundings.




Figure 5.2. Example of a DIC microscopy image. Microtubule nucleation initiates from stable tubulin “seeds”. In the experiments, “observation lines” are drawn along MT bodies to construct kymographs.

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 in DIC microscopy images, the object appearance (and especially the MT tip) depends on the imaging conditions (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 every second, which is in fact a quite high sampling rate taking into account how slowly microtubules grow *in vitro* (30-40 nm per second). This 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 (in our experiments, the pixel size is $80 \times 80 \text{ nm}^2$), 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 [141]. 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 [130] (also called a kymoimage in this




Figure 5.3. Example of a kymograph obtained from the DIC microscopy images, showing the dynamics of both microtubule ends.

chapter) for each MT. It is constructed by defining (manually or automatically) an “observation line” L (Fig. 5.2) 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 , $J_t = \{J_t(j) : j = 1, \dots, Y\}$, 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. 5.3 for an example), $I(t, y) = \{J_t : t = 1, \dots, T\}$, 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 [28]. These “seeds” always remain present and cannot be completely disassembled. In the kymoimages (Fig. 5.3) they are clearly visible as a bright horizontal strip.

To estimate the kinematic parameters of interest from the kymoimages, the edge location $y(t)$ (corresponding 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 this chapter, 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.

5.2.3 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. 5.3). 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 in-

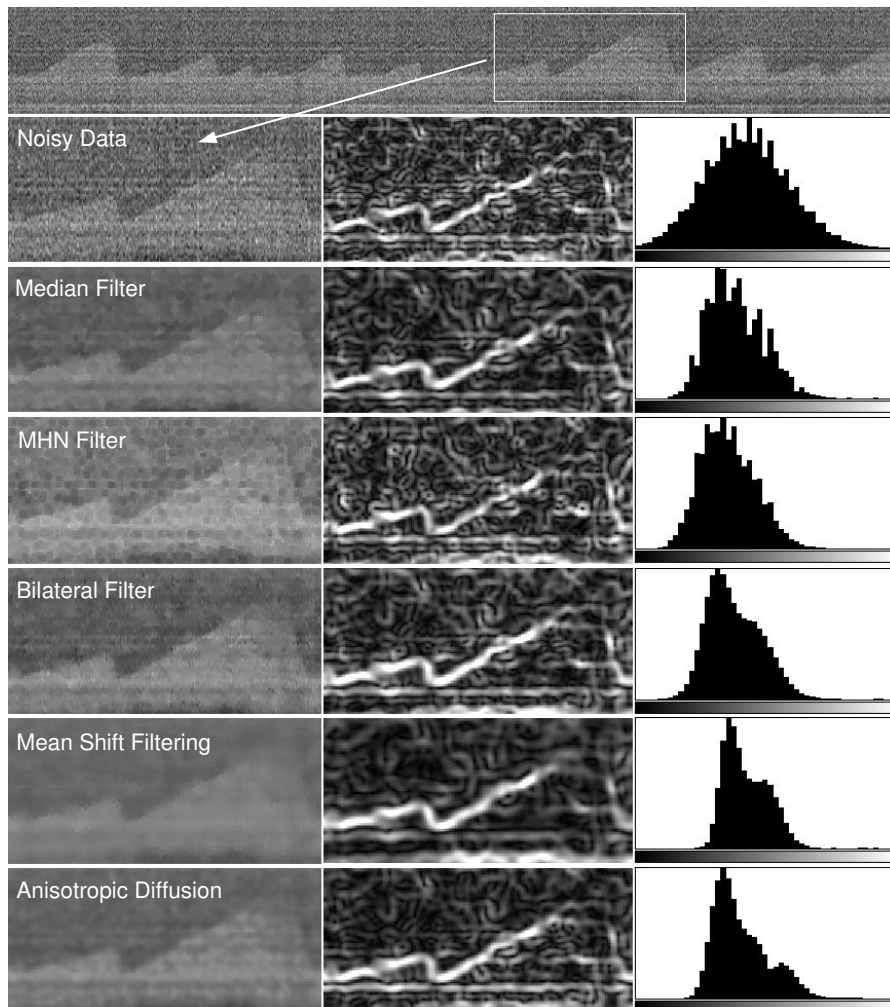


Figure 5.4. Application of various edge preserving smoothing methods to our image data (top). The left column shows the results of smoothing, the middle column depicts the edge information extracted using the Gaussian derivatives, and the right column shows the distribution of intensity values in the smoothed images.

formation. Better results may be obtained by the use of nonlinear, edge preserving filters. Fig. 5.4 shows the results of applying the most frequently used nonlinear filtering techniques to our image data: the median filter [149], the maximum homogeneity neighbor (MHN) filter [49], the bilateral filter [163], the mean-shift filter [34], and anisotropic diffusion [116]. 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 [159] 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 [25] is usually based on hard thresholding.

5.2.4 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 [9, 126] 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 [9]. 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:

$$\mathbf{x}_t = f_t(\mathbf{x}_{t-1}, \mathbf{v}_t), \quad \mathbf{z}_t = g_t(\mathbf{x}_t, \mathbf{u}_t), \quad (5.1)$$

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 (5.1) [9, 126].

The solution of the state-space problem given by (5.1) is the posterior probability distribution function (pdf) $p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t})$, where $\mathbf{x}_{0:t} = \{\mathbf{x}_0, \dots, \mathbf{x}_t\}$ and $\mathbf{z}_{0:t} = \{\mathbf{z}_0, \dots, \mathbf{z}_t\}$, 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 [126] or, in the most general case, using approximations such as sequential Monte Carlo (MC) methods [9, 39]. 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

$$p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t}) = \sum_{i=1}^{N_s} w_t^{(i)} \delta(\mathbf{x}_{0:t} - \mathbf{x}_{0:t}^{(i)}), \quad (5.2)$$

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

$$\mathbf{x}_t^{(i)} \sim p(\mathbf{x}_t | \mathbf{x}_{t-1}^{(i)}) \text{ and } w_t^{(i)} \propto w_{t-1}^{(i)} p(\mathbf{z}_t | \mathbf{x}_t^{(i)}), \quad (5.3)$$

$i = 1, \dots, N_s$. The minimum mean square error (MMSE) or maximum a posteriori (MAP) estimators of the state can be easily obtained from $p(\mathbf{x}_{0:t} | \mathbf{z}_{0:t})$ [9].

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 [56,57]. 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 [11,52], but was shown to outperform IMM in most cases [56]. The VRPF, which was initially proposed for tracking of highly maneuvering targets [56], is nowadays successfully applied in other fields, for example DNA sequencing [61], 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 $I(t, y)$. Similarly 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

$$\mathbf{x}_k \sim p(\mathbf{x}_k | \mathbf{x}_{k-1}) = p_\theta(\theta_k | \theta_{k-1}, \tau_k, \tau_{k-1}) p_\tau(\tau_k | \theta_{k-1}, \tau_{k-1}). \quad (5.4)$$

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

$$\mathbf{z}_t \sim p(\mathbf{z}_t | \mathbf{x}_0, \dots, \mathbf{x}_\infty) = p(\mathbf{z}_t | \mathbf{x}_{\mathcal{N}_t}). \quad (5.5)$$

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 $h_t(\cdot)$ is a deterministic function of the state in the neighborhood \mathcal{N}_t . The observation density (5.5) is then expressed as

$$p(\mathbf{z}_t | \mathbf{x}_{\mathcal{N}_t}) = p(\mathbf{z}_t | \hat{\theta}_t). \quad (5.6)$$

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. Similarly 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, similar to (5.3) [56, 57], 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 (5.4), we model the MT dynamics with the transition priors

$$\begin{aligned} p_\theta(\theta_k | \theta_{k-1}, \tau_k, \tau_{k-1}) &= p(v_k | v_{k-1}, y_k, y_{k-1}, \tau_k, \tau_{k-1}) p(y_k | v_{k-1}, y_{k-1}, \tau_k, \tau_{k-1}) \\ &= p(v_k | v_{k-1}) \delta(y_k - y_{k-1} - v_{k-1}(\tau_k - \tau_{k-1})), \end{aligned} \quad (5.7)$$

$$p_\tau(\tau_k | \theta_{k-1}, \tau_{k-1}) = \mathcal{U}_{[\tau_{k-1} + \tau^0, \tau_{k-1} + \tau^1]}, \quad (5.8)$$

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

$$\begin{aligned} \tau_k - \tau_{k-1} &\sim \mathcal{U}_{[\tau^0, \tau^1]}, \\ y_k &= y_{k-1} + v_{k-1}(\tau_k - \tau_{k-1}), \\ v_k &\sim p(v_k | v_{k-1}). \end{aligned} \quad (5.9)$$

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 $\tau_{k-1}^{(i)} \leq t$, 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 5.2.1). 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, $\sigma_{\nu^+}^2, \sigma_{\nu^-}^2, \sigma_{\nu^0}^2$, 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

$$p(v_k|v_{k-1}) = \begin{cases} (1-a)\mathcal{N}(v_{k-1}, \sigma_{\nu^+}^2) + a\mathcal{N}(\nu^-, \sigma_{\nu^-}^2), & \text{for } v_{k-1} > V_{\text{th}}, \\ (1-a)\mathcal{N}(v_{k-1}, \sigma_{\nu^-}^2) + \\ \quad a(\mathcal{N}(\nu^+, \sigma_{\nu^+}^2) + \mathcal{N}(\nu^0, \sigma_{\nu^0}^2))/2, & \text{for } v_{k-1} < -V_{\text{th}}, \\ (1-a)\mathcal{N}(v_{k-1}, \sigma_{\nu^+}^2) + a\mathcal{N}(\nu^+, \sigma_{\nu^+}^2), & \text{for } |v_{k-1}| < V_{\text{th}}, \end{cases} \quad (5.10)$$

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 [133, 134]. The proposed model describes a small perfectly sharp edge and consist of two rectangular regions, S_B and S_F (black and white rectangles in Fig. 5.5, 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 is defined by the velocity component $v_{k-1}^{(i)}$. The regions S_B and S_F are defined as follows

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

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

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

$$p(\mathbf{z}_t|\mathbf{x}_{\mathcal{N}_t}) = \begin{cases} \exp\left(\frac{\mu_F - \mu_B}{\gamma}\right) - 1, & \mu_F - \mu_B > 0, \\ 0, & \mu_F - \mu_B \leq 0, \end{cases} \quad (5.11)$$

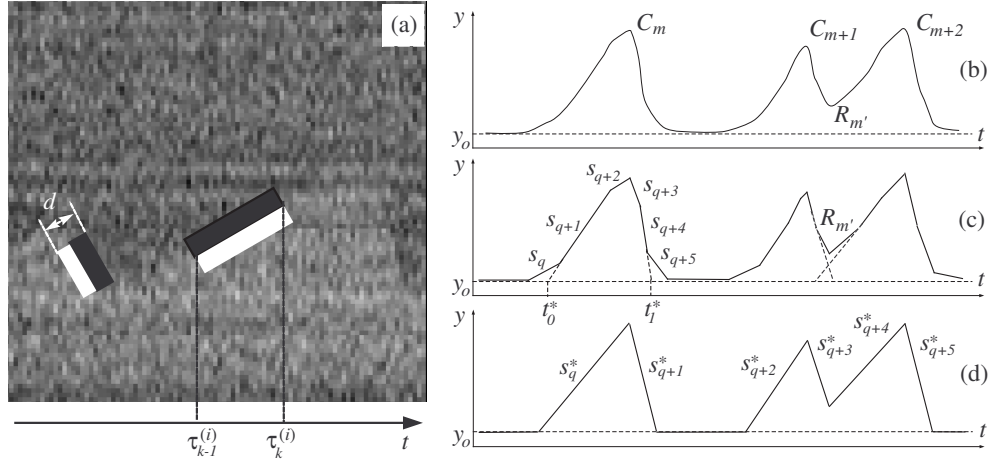


Figure 5.5. The observation model used in the experiments, which compares the intensity distribution in two rectangular regions (black and white strips) and defines the likelihood of edge existence (a). Examples of applying the MTA to the extracted edge using the VRPF in order to compute the kinematic parameters (b-d).

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 (5.8). 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 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 were 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. 5.4 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 $p(\mathbf{x}_{0:N_t^+} | \mathbf{z}_{0:t})$, that is

$$\hat{y}_{t-\Delta t} = \sum_{i=1}^{N_s} w_t^{(i)} h_t(\mathbf{x}_{N_t-\Delta t}). \quad (5.12)$$

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.

5.2.5 Multiscale Trend Analysis

Having the estimated edge \hat{y}_t after applying the VRPF, we employ multiscale trend analysis (MTA) [188] 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 [189]. 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, $l = \{1, \dots, N_L\}$, 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 $\{N_{\text{cat}}(l) : l = \{1, \dots, N_L\}\}$, 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 , $m = \{1, \dots, N_{\text{cat}}^*\}$, which occurs at time t_m^c , the two sets of neighboring nodes, $\{s_q : t_{m-1}^c < t_0^q < t_m^c \cap \alpha^q > 0, q = 1, \dots, l^*\}$ and $\{s_q : t_m^c < t_1^q < t_{m+1}^c \cap \alpha^q < 0, q = 1, \dots, l^*\}$ are analyzed (see Fig. 5.5(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 $[\tilde{t}_m^0, t_{m'}^c]$ and $[t_{m'}^c, \tilde{t}_m^1]$. If the rescue event is positioned between two catastrophes C_m and C_{m+1} , the approximation is recomputed on the interval $[t_{m'}^c, t_{m'}^R]$. The new approximation is given by a new set of nodes $S^* = \{s_q^*\}_{q=1}^{2N_{\text{cat}}^*}$ (see Fig. 5.5(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:

$$T^+ = \sum_{\substack{\forall s_q^* \in S^* \\ \alpha^q > 0}} (t_1^q - t_0^q), \quad \hat{\nu}^+ = \frac{1}{T^+} \sum_{\substack{\forall s_q^* \in S^* \\ \alpha^q > 0}} (t_1^q - t_0^q) \alpha^q, \quad (5.13)$$

$$T^- = \sum_{\substack{\forall s_q^* \in S^* \\ \alpha^q < 0}} (t_1^q - t_0^q), \quad \hat{\nu}^- = \frac{1}{T^-} \sum_{\substack{\forall s_q^* \in S^* \\ \alpha^q < 0}} (t_1^q - t_0^q) \alpha^q, \quad (5.14)$$

$$f_{\text{cat}} = N_{\text{cat}}^*/T^+, \quad f_{\text{res}} = N_{\text{res}}^*/T^-, \quad (5.15)$$

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.

5.3 Experimental Results

The performance of the proposed VRPF-based method was evaluated using both synthetic images (Section 5.3.1) and real data from studies of MT dynamics *in vitro* (Section 5.3.2). The synthetic images, for which the ground truth was available, gave the possibility 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.

5.3.1 Evaluation on Synthetic Data

5.3.1.1 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 5.2.1 (Fig. 5.1). 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 is 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, $\nu^+ \sim \mathcal{N}(0.5, 0.005)$,

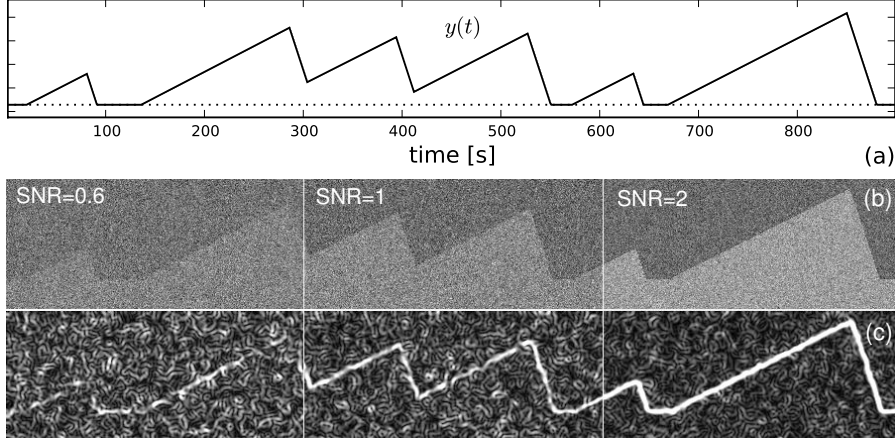


Figure 5.6. Examples of the synthetic images used in the experiments. The simulated MT tip dynamics (a) is used to create the synthetic images for different SNR levels (b), for which the gradient images (c) are computed using the Gaussian derivatives at scale $\sigma_G = 3$.

$\nu^- \sim \mathcal{N}(-3, 0.005)$, $\nu^0 \sim \mathcal{N}(0, 0.05)$. These model values are representative of practical values.

Having the simulated dynamics $y(t)$, $0 < t < T$ (see Fig. 5.6(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 5.2.4). The height of the generated images was in the range of 100 – 150 pixels, which corresponds to 8 – 12 μm ($\Delta_t = 1\text{s}$ and $\Delta_y = 80\text{nm}$). For all t , the image pixels were filled with background intensity $I_B = 100$ if $j > y(t) + y_0$ and otherwise with 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 images obtained using light microscopy [185]. Examples of synthetic images for various SNRs are shown in Fig. 5.6(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. 5.6(c).

The parameters of the described VRPF algorithm were fixed to the following values: $\nu^+ = 0.5$, $\nu^- = -3$, $\sigma_{\nu^0}^2 = 0.5$, $\sigma_{\nu^+}^2 = 0.05$, $\sigma_{\nu^-}^2 = 0.5$, $V_{\text{th}} = 0.15$, $d = 6$, $\tau^0 = 3$, $\tau^1 = 10$, $\Delta t = 20$, $N_s = 500$, $N_L = 80$, $a = 0.01$, $\gamma = 10$. 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) [104]:

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (y_t - \hat{y}_t)^2}, \quad (5.16)$$




Figure 5.7. 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.

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 $|\cdot|$ denotes the set size operator.

5.3.1.2 Results

The proposed VRPF method was evaluated using 20 synthetically generated images. Examples of edge extraction for $\text{SNR} = 0.6$ are shown in Fig. 5.7. 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 5.2.4). PF_1 uses only one state transition model, $p(\mathbf{x}_t|\mathbf{x}_{t-1})$, which describes nearly-constant velocity motion [141]. To capture abrupt changes in the edges, the variance of the process noise in this transition model had to be made rather large. Due to this high variance, the typical overshoots just after the catastrophe events (see Fig. 5.7(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

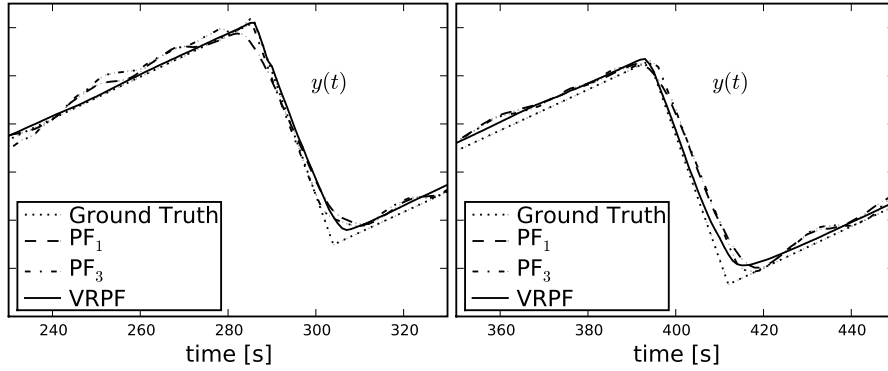


Figure 5.8. More detailed results of extracting the edge information from the noisy synthetic images using the proposed VRPF and two types of standard PFs in comparison with the ground truth. The plots are zooms of the first two peaks in Fig. 5.7 and show the results combined.




Figure 5.9. Results of edge extraction using the Canny edge detector for two different values of the hysteresis thresholds.

but fixed length $l_v = 5$ was used. The zoomed results in Fig. 5.8 clearly show that the edge $\hat{y}(t)$ estimated using the standard PFs is typically less smooth and piecewise linear. For visual comparison, the edge information extracted using the Canny edge detector [25] for two different values of hysteresis thresholds is shown in Fig. 5.9.

The results of applying MTA for kinematic parameter estimation based on the edges extracted using PF_3 and VRPF are shown in Table 5.1 (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.

Table 5.1. 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.

SNR	RMSE	$\nu^+ \pm \text{sd}$	$\nu^- \pm \text{sd}$	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

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 practice, however, this inaccuracy can be easily reduced, by running the algorithm 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. Then, the algorithm is initialized with the new estimates and rerun. In the experiments, we found that this approach always resulted in estimates in the range $(\nu \pm \sigma_\nu)$ defined by the ground truth.




Figure 5.10. Example of a kymograph generated in the experiments on real DIC microscopy image data with $\text{SNR} \approx 1$ (a) and the results of applying the proposed VRPF (b).

5.3.2 Evaluation on Real Data

For the validation on real data we collected 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 $15\mu\text{M}$ of tubulin was added (Experiment I), or, in addition, $1\mu\text{M}$ of EB3 (Experiment II), or $1\mu\text{M}$ of GFP-EB3 (Experiment III) [78]. From each sequence, 10-20 MTs were selected by biologists, and the observation lines were drawn manually. The image sequences contained about 1000-1200 frames (one per second) of size 700×500 pixels (of size $86 \times 86 \text{ nm}^2$). To estimate the parameters of interest, for each experiment 10 kymographs were constructed and analyzed manually and using the proposed VRPF method. The results are presented in Table 5.2, where the speed estimates are also converted to $\mu\text{m}/\text{min}$. The usage of these units is more common in biological experiments and it also allows straightforward comparison with the recently published results [78]. An example of edge extraction using VRPF in real data is shown in Fig. 5.10. A comparison of the estimates obtained by manual and VRPF-based analysis suggests that the proposed automatic method may replace the laborious manual procedures.

5.4 Discussion and Conclusions

In this chapter 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 filtering method, which is better capable of

Table 5.2. Results of parameter estimation in real DIC microscopy image data sets using manual analysis versus VRPF.

	$\nu^+ \pm \text{sd}$ [pix/frame]	$\nu^- \pm \text{sd}$ [pix/frame]	f_{cat}	ν^+ [$\mu\text{m}/\text{min}$]	ν^- [$\mu\text{m}/\text{min}$]
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

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 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. 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 may replace laborious manual analyses.