BAYESIAN IMAGE ANALYSIS WITH ON-LINE CONFIDENCE ESTIMATES AND ITS APPLICATION TO MICROTUBULE TRACKING

J. Cardinale[*], *A. Rauch*[†], *Y. Barral*[†], *G. Székely*[‡] *& I. F. Sbalzarini*[*]

[*]Institute of Theoretical Computer Science & Swiss Institute of Bioinformatics, ETH Zurich, Switzerland
[†]Institute of Biochemistry, ETH Zurich, Switzerland
[‡]Computer Vision Laboratory, ETH Zurich, Switzerland

ABSTRACT

Automated analysis of fluorescence microscopy data relies on robust segmentation and tracking algorithms for sub-cellular structures in order to generate quantitative results. The accuracy of the image processing results is, however, frequently unknown or determined a priori on synthetic benchmark data. We present a particle filter framework based on Markov Chain Monte Carlo methods and adaptive annealing. Our algorithm provides on-line per-frame estimates of the detection and tracking confidence at run time. We validate the accuracy of the estimates and apply the algorithm to tracking microtubules in mitotic yeast cells. This is based on a likelihood function that accounts for the dominant noise sources in the imaging equipment. The confidence estimates provided by the present algorithm allow on-line control of the detection and tracking quality.

Index Terms— particle filter, confidence estimate, adaptive annealing, microtubule, tracking

1. INTRODUCTION

Statistical analysis of dynamic intracellular structures using fluorescence microscopy requires acquisition and processing of large sets of images. Computer-based segmentation and tracking procedures are essential to sustain reproducibility, speed, and accuracy of the analysis [1, 2]. Statistical interpretation of the data relies on confidence estimates of the segmentation and tracking results. In general, however, ground truth data are not available, which limits validation and error analysis of image processing algorithms. Typically, errors are estimated on artificially generated data at different Signal to Noise Ratios (SNR). This only considers the mean (expected) error at a certain SNR.

Per-frame confidence estimates can be obtained on-line (at run time) directly on the specific image. Carefully designed likelihood functions model the imaging process, accounting for the Point Spread Function (PSF) of the microscope and the dominant noise sources. Using a particle filter as a recursive Bayesian estimator, we calculate a posterior (filtering) probability density function (pdf), accounting for prior knowledge about the imaged process and the likelihood function. This posterior is represented using a particle-based function approximation. Particles live in state space and have an assigned weight, reflecting a normalized pdf value. Moments such as mean and covariance can be computed from the approximated posterior at each time point. Using the estimated posterior moments allows determining confidence intervals for the estimated states. This on-line error quantification is based on the particular realization of the noise distribution in the current image. Analysis of the posterior also enables detecting breakdowns of the tracking algorithm.

In tracking applications using particle filters, it is often the case that only few particles carry significant weight. This degeneracy problem is usually addressed using a resampling procedure. After resampling, all particles carry non-degenerate weights, but cover less of the state space. S. Godsill and T. Clapp proposed Markov Chain Monte Carlo (MCMC) methods in order to relax this sample impoverishment problem [3]. We propose here a variant of this framework, featuring an adaptive annealing scheme capable of handling multiscale likelihood functions with peak widths varying in a broad spectrum. This presents an alternative to iteratively using a subset of the particles to search for high likelihood function values [4].

We validate the confidence estimates of our method on synthetic data at different SNR and present an application to the analysis of the intrinsically stochastic motion of astral microtubules during metaphase in yeast cells. Particle filtering without on-line confidence estimates has been previously used to successfully track microtubules in live cells [5]. We track spindle pole bodies and microtubule tips in 3D fluorescence confocal microscopy movies (see Fig. 1a). Using knowledge about the imaging system, we perform 3D tracking and on-line error analysis with sub-voxel resolution.

2. ALGORITHM

In order to estimate the hidden states and confidence intervals, we analyze the likelihood function $\mathcal{L}(I_k|x_k^i)$, indicat-

JC was supported through CTI grant 9325.2 PFLS-LS.

ing how likely it is to observe the image I_k at a certain time point (frame) k given a certain state x_k^i. In our framework, the prior $p(x_k^i|x_{k-1}^i)$ is used to constrain the search space by filtering the likelihood, preserving its shape in order to ensure unbiased statistics. We first present the generic part of the algorithm and then the application-specific likelihood function and prior distribution for microtubule tracking.

2.1. Generic Part

We use the adaptive procedure outlined in Algorithm 1 to relax the sample impoverishment problem and allow for multiscale likelihood functions. The algorithm starts by defining the initial set $\mathcal{P}$ of n particles (line 1 of Algorithm 1), which are then sampled from the proposal distribution q (line 4). This is followed by computation of the normalized particle weights (lines 6–7) and MCMC iterations (lines 8–18) comprising two phases: annealing (lines 10–13) and classical MCMC moves (line 17). The annealing phase provides a good starting point and proposal distribution for the classical MCMC moves.

Within the MCMC sub-routine (line 17), the particle positions are updated by a Metropolis-Hastings algorithm using the posterior pdf as its stationary distribution. This yields the new particle set $\mathcal{P}_u$ with unchanged weights. We use a Gaussian with a diagonal covariance matrix as the proposal distribution for the MCMC moves. Good proposal distributions should be similar to the desired stationary distribution [6]. During annealing, the covariance matrix of $\mathcal{P}_u$ (line 9) approximates the covariance matrix of the stationary distribution. We estimate the covariance matrix as:

$$\tilde{\Sigma} = \sum_{i=1}^{n} w^i (x^i - \tilde{x})(x^i - \tilde{x})^T , \qquad \tilde{x} = \sum_{i=1}^{n} w^i x^i . \quad (1)$$

Since the initial proposal distribution is very broad, most particles are degenerate. In order to cluster the particles at positions of high likelihood, we perform a resampling operation (line 11). After resampling, Σ is adapted (annealed) in line 12. The parameter $0 < c < 1$ is the exponential learning factor of the adaptation of Σ. Annealing ends as soon as the adaptation becomes insignificant. We quantify significance by the Kullback-Leibler divergence D_{KL} between two normal distributions with means 0 and covariance matrices Σ and Σ' (line 14). If its derivative D'_{KL} is smaller than a user-defined threshold ϵ, annealing stops (lines 14–16).

As soon as a user-defined termination criterion (e.g. maximum number of iterations) is met (line 18), the desired estimates $\hat{x}$, $\hat{\Sigma}$ for the moments of the posterior are computed based on the union of all particle sets since the end of the annealing phase (line 19).

2.2. Specific Part

We design the likelihood function $\mathcal{L}(I_k, x_k^i)$ and the prior distribution $p(x_k^i|x_{k-1}^i)$ for the specific application of astral microtubule tracking during metaphase in yeast cells. 3D digital videos were acquired using a confocal microscope and show the microtubule tip and Spindle Pole Body (SPB) proteins Spc72p and Bik1p, labeled with green fluorescent protein (Fig. 1a–c). The dynamics of the system are driven by microtubules randomly switching between phases of assembly and disassembly. As fluorescent proteins are constantly transported along the microtubules, they may cause object intensities to increase over time, even though the total intensity in the whole image decreases due to photobleaching. We model the system as 3 connected, diffraction-limited objects (microtubule tip plus 2 SPBs) that appear as scaled 3D PSFs. In order to serve as biologically relevant read-out, the peak intensities of the objects are explicitly represented in the state vector, resulting in a 12-dimensional state space.

Algorithm 1 Particle Filter with Adaptive Annealing

1: Initialize sample set $\mathcal{P}_1 = \{x_0^i\}$
2: **for** $k = 1 \dots N_{\text{frames}}$ **do**
3: $u = 1, annealing = true$
4: Draw samples $x_k^i \sim q, w^i = 1/n$
5: Initialize Σ using Eq. 1 on $\mathcal{P}_1$
6: Compute weights $w^i = \mathcal{L}(I_k, x_k^i) p(x_k^i|x_{k-1}^i)$
7: Normalize weights $w^i = w^i / \sum_j w^j$
8: **repeat**
9: Estimate Σ' using Eq. 1 on $\mathcal{P}_u$
10: **if** $annealing$ **then**
11: $\mathcal{P}_u$ = RESAMPLE($\mathcal{P}_u$)
12: Adapt $\Sigma = (1-c)\Sigma + c\Sigma'$
13: **end if**
14: **if** $D'_{\mathrm{KL}}(\mathcal{N}(0,\Sigma)||\mathcal{N}(0,\Sigma')) < \epsilon$ **then**
15: $annealing = false, u' = u$
16: **end if**
17: MCMC($\mathcal{P}_u, \Sigma_{j,j}$), $u = u + 1$
18: **until** not(termination criterion)
19: Compute $\hat{x}$ and $\hat{\Sigma}$ using Eq. 1 on $\cup_{j=u'}^{u} \mathcal{P}_j$
20: **end for**

The likelihood function is designed to model the process of image acquisition. This involves both a good model for the dominant noise sources of the imaging equipment and the PSF of the microscope.

We used an electron multiplying charge-coupled device (EMCCD) camera, where electrons exiting the CCD sensor are multiplied in a pipeline before the voltage is read out. The dominant noise sources are Poisson-distributed shot noise and excess noise (multiplication noise). Excess noise is introduced by the stochastic process of impact ionization in the electron-multiplying pipeline, while shot noise is imposed by the discrete photon counts. The pdf of the number n of electrons exiting the multiplying pipeline for a number N of electrons entering is modeled as [7]:

$$p(n, N) = \frac{1}{\sqrt{2\pi F^2 M^2 N}} \exp \frac{-(n - MN)^2}{2F^2 M^2 N} , \quad (2)$$

where $M > 1$ is the linear digital gain and F the excess noise factor (ENF). The number N of electrons entering the pipeline can be determined from the grayscale intensity values O in the image as:

$$N = (O - D)\frac{C}{M \cdot Q_\lambda}\,. \tag{3}$$

D is the mean of the "dark image", recorded using the camera system without any light source, C is a camera-specific constant provided by the manufacturer, and Q_λ the quantum efficiency of the CCD sensor at the recording wavelength λ.

Using this noise model, we formulate for each pixel (or voxel) $v \in I$ in the observed image I the marginal $p(I_v|J_v)$, the pdf of the pixel having intensity I_v given an expected intensity J_v. For a specific realization of J, i.e. an expected image, $p(J_v) = 1$, such that the marginal is given by the joint pdf in Eq. 2 with $N = J_v$. The expected image J (in the absence of noise) is computed from the current state vector x_k of the particle filter and the PSF of the microscope as:

$$J(x_k) = B + \sum_{i=1}^{2} J_{\mathrm{SPB}_i}\mathrm{PSF}(x_{\mathrm{SPB}_i}) + J_{\mathrm{Tip}}\mathrm{PSF}(x_{\mathrm{Tip}})\,. \tag{4}$$

$J_{\mathrm{SPB}_{1,2}}$ and J_{Tip} are the expected intensities of the two SPBs and the microtubule tip as stored in the state vector, $\mathrm{PSF}(y)$ is the microscope PSF centered at position y. All pixels of the background image B are set to the most frequently occurring pixel intensity in the observed I. Construction of the marginals is completed by modeling the PSF of the microscope. As a first approximation, we use a Gaussian whose width is fitted to recorded images of point sources. In addition, we also measured the true PSF by imaging fluorescent beads, and generated a high-resolution lookup table for the PSF as described in Sec. 3.1. The two models are compared in Sec. 3.2.

Using the noise and PSF models outlined above, and assuming the noise in different pixels to be statistically independent, the likelihood function can be constructed by multiplying all marginal, thus:

$$\mathcal{L}(I_k, x_k^i) = \prod_v p(I_v|J_v(x_k))\,. \tag{5}$$

Evaluation of this likelihood function is computationally efficient since many voxels have identical intensity values and need not be considered separately.

The prior distribution has, in our framework, the role of constraining the search space to regions of high likelihood. Before the annealing step of the algorithm, it is equal to the proposal distribution q of the particle filter. We choose a uniform distribution $q = \mathcal{U}[\mathcal{I}]$ over the interval $\mathcal{I}$. For the present application, we use cell-fixed spherical coordinates. The old SPB serves as the origin, the optical axis of the microscope corresponds to zero azimuthal angle, and the x-axis in image space is used as the reference for the polar angle. The interval $\mathcal{I}$ around x_k is specified in these coordinates.

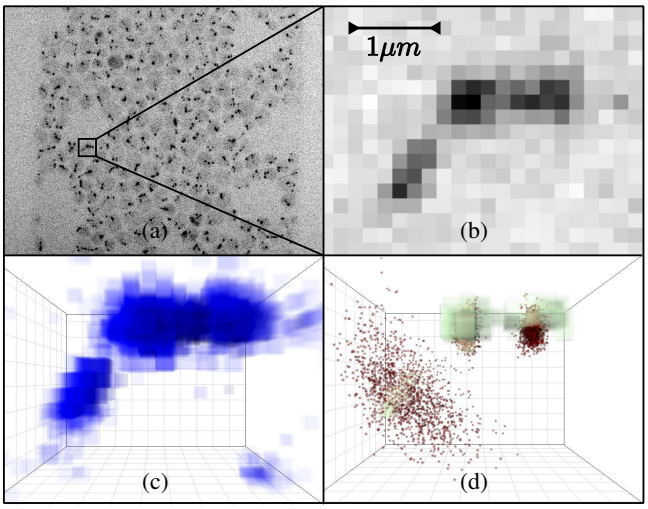

Fig. 1. (a) An example micrograph of labeled yeast cells in maximum projection. Intensities are inverted for better visualization. (b) Maximum projection of the labeled microtubules in a single cell during mitosis. (c) 3D stack of the microtubules in (b). The spots from left to right are the microtubule tip, SPB 2 (old pole), and SPB 1 (new pole). (d) Image with overlaid proposal distribution. Each particle of the particle filter gives rise to 3 dots, corresponding to the 9 estimated position dimensions.

3. VALIDATION AND RESULTS

3.1. Validation Data

We validate the error estimates on artificial feature point data. Time series of artificial images K_i are generated according to Eq. 4 based on known state vectors x and background $B = 50$. The feature point positions perform a random walk, yielding a movie of simulated Brownian motion. As a PSF, we use the measured PSF of the microscope determined from images of fluorescent beads. Assuming the PSF to be radially symmetric, the intensities are averaged along circles of different radii around the intensity centroid of the bead's image [1]. We average the PSF determined from 5 different images in order to reduce the noise. Since the width of the likelihood function decreases with increasing SNR, we test adaptation to different likelihood widths by simulating different SNRs. This is done by scaling the peak intensity $K_{\max}$ of the PSF according to: $\mathrm{SNR} = (K_{\max} - b)/(F \cdot \sqrt{K_{\max}})$. The ENF is measured to be $F = 1.5$ for our equipment. Finally we replace all pixel values K_v in K by Gaussian random numbers with mean K_v and standard deviation $FM\sqrt{K_v}$ (cf. Eq. 2) with gain $M = 40$. The resulting images correspond to a pixel size of 160×160 nm and a distance between confocal planes (voxel depth) of 200 nm.

3.2. Results

We use the present adaptive particle filter to track the Brownian motion of the feature points in the artificially generated movies and assess the quality of the confidence estimates (not of the tracking itself!). We compare the estimated standard deviation $\hat{\sigma}$ of the x and z positions (y is identical to x, not shown) – as determined by the $\hat{\Sigma}$ in Algorithm 1 – to the standard deviation with respect to the known true positions x_i of the validation data: $\sigma = \sqrt{\frac{1}{f}\sum_i^f (\hat{x}_i - x_i)^2}$. All estimates are averaged over $f = 200$ frames, leading to averaged true and estimated uncertainties $\langle\sigma\rangle$ and $\langle\hat{\sigma}\rangle$, respectively. We use $n = 40$ particles and 30 MCMC iterations after annealing. Figure 2 shows the results. The confidence estimates as determined by Algorithm 1 (solid lines) are always larger than the reference standard deviations of the benchmark data (dashed lines), providing conservative uncertainty estimates. Using a Gaussian PSF model (squares), reduces the absolute tracking quality (dashed lines) in z, but has no significant effect on the estimated tracking uncertainty (solid lines). For the true PSF, our algorithm is able to accurately estimate the tracking confidence with mean standard deviation differences $\Delta\sigma = |\langle\sigma\rangle - \langle\hat{\sigma}\rangle| < 7\,\text{nm}$ in the lateral and $\Delta\sigma < 12\,\text{nm}$ in axial direction. The standard deviation of the estimator decreases with increasing SNR and is low throughout (<5 nm laterally and <22 nm axially). In 99% of the frames, the confidence estimates are accurate within $\pm 3\hat{\sigma}$.

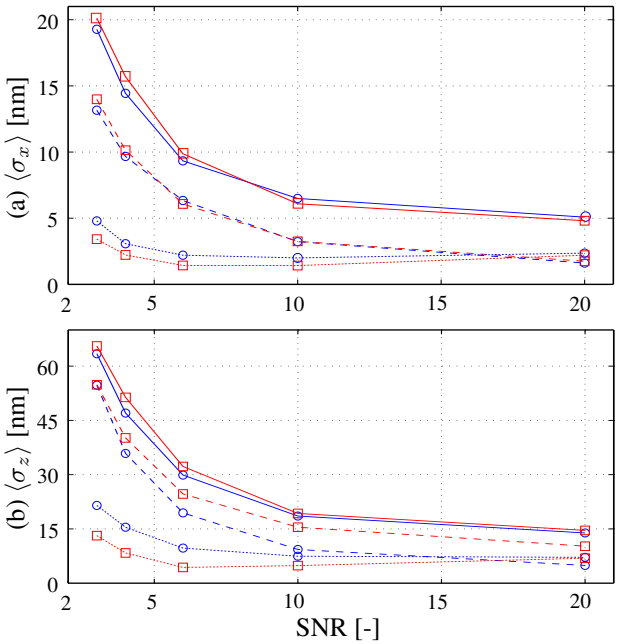

Fig. 2. Means of real (dashed lines) and estimated (solid lines) tracking uncertainties at different SNRs using a Gaussian PSF approximation (squares) and the true, measured PSF (circles). The dotted lines represent the standard deviations of the estimators. (a) In the lateral x direction, the mean estimate deviates by $3.0 \ldots 6.2\,\text{nm}$ (< 0.04 pixel). (b) In the axial z direction, the difference of the mean errors is between 8.7 and $11.2\,\text{nm}$ (< 0.06 voxel depths).

4. CONCLUSIONS

We have presented and validated an object tracking framework with on-line confidence estimates based on a particle filter with adaptive annealing. The framework provides on-line sub-pixel estimates of the tracking uncertainty at all SNRs tested. In our case, using a Gaussian PSF model had no significant effect on the estimator. The presented adaptive MCMC scheme handles well likelihood functions of different widths. On multimodal posterior distributions, however, the proposal distribution of the MCMC is bad, which leads to frequent rejection of MCMC moves and bad estimation of tracking uncertainly. Our framework, as well as the PSF estimation tool, are implemented as plug-ins to the open-source platform ImageJ.

5. REFERENCES

[1] I. F. Sbalzarini and P. Koumoutsakos, "Feature point tracking and trajectory analysis for video imaging in cell biology," *J. Struct. Biol.*, vol. 151, no. 2, pp. 182–195, 2005.

[2] K. Jaqaman, D. Loerke, M. Mettlen, H. Kuwata, S. Grinstein, S.L. Schmid, and G. Danuser, "Robust single-particle tracking in live-cell time-lapse sequences," *Nat. Methods*, vol. 5, pp. 695–702, 2008.

[3] S. Godsill and T. Clapp, "Improvement strategies for Monte Carlo particle filters," in *Sequential Monte Carlo Methods in Practice*. 2001, Springer-Verlag.

[4] S. J. McKenna and H. Nait-Charif, "Tracking human motion using auxiliary particle filters and iterated likelihood weighting," *Image Vision Comput.*, vol. 25, no. 6, pp. 852–862, 2007.

[5] 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 Trans. Medical Imaging*, vol. 27, no. 6, pp. 789–804, June 2008.

[6] Siddhartha Chib and Edward Greenberg, "Understanding the metropolis-hastings algorithm," *The American Statistician*, vol. 49, no. 4, pp. 327–335, 1995.

[7] J. Hynecek and T. Nishiwaki, "Excess noise and other important characteristics of low light level imaging using charge multiplying CCDs," *IEEE Trans. Electron Devices*, vol. 50, no. 1, pp. 239–245, 2003.