

A NEW HYBRID BAYESIAN-VARIATIONAL PARTICLE FILTER WITH APPLICATION TO MITOTIC CELL TRACKING

Ricard Delgado-Gonzalo, Nicolas Chenouard, and Michael Unser

Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Switzerland

ABSTRACT

Tracking algorithms are traditionally based on either a variational approach or a Bayesian one. In the variational case, a cost function is established between two consecutive frames and minimized by standard optimization algorithms. In the Bayesian case, a stochastic motion model is used to maintain temporal consistency. Among the Bayesian methods we focus on the particle filter, which is especially suited for handling multimodal distributions. In this paper, we present a novel approach to fuse both methodologies in a single tracker where the importance sampling of the particle filter is given implicitly by the optimization algorithm of the variational method. Our technique is capable of outlying nuclei and tracking the lineage of biological cells using different motion models for mitotic and non-mitotic stages of the life of a cell. We validate its ability to track the lineage of HeLa cells in fluorescence microscopy.

Index Terms— Active contour, snake, ellipse, ovuscul, motion, mitosis, HeLa.

1. INTRODUCTION

Because biological systems are dynamic, it is highly desirable to quantify their evolution through time in order to improve our understanding of their behavior. Current efforts in cell tracking rely on a large variety of methodologies [1]. Among them, two main paradigms can be identified: Bayesian framework and variational methods. The former involves a probabilistic reasoning grounded in a motion model [2]. The latter localizes the target accurately at each frame by optimizing a cost function that depends exclusively on the current image, often employing a standard minimization algorithm [3].

In bioimaging, the variational approach is usually preferred [4, 5, 6]. Nonetheless, several attempts have been made in the Computer Vision community to take advantage of Bayesian and variational methods simultaneously. Most of these methods make use of parametric active contours and rely on kernel density estimators, which can degrade the computational performance of the algorithm [7, 8]. The use of such estimators was avoided in [9], but active contours of the geometric variety were introduced, which are significantly slower than their parametric counterparts.

In this paper, we present a tracking algorithm that merges the capabilities of Bayesian tracking with parametric active contours (a.k.a. snakes). Among the Bayesian methods, our interest lies in the particle filter, which performs a multimodal random search guided by a motion model [2]. The fact that the search is multimodal is important when modeling uncertainties of association in dividing targets. In the role of the parametric contours, we chose an elliptical

This work was funded by the Swiss SystemsX.ch initiative under Grant 2008/005.




Fig. 1. Ovuscul. The outer ellipse Γ is shown in darker gray, while the inner ellipse Γ' is shown in lighter gray. These ellipses are entirely determined by the triplet of points $\{p, q, r\}$.

snake named the ovuscul that is robust against noise and has a convenient parameterization [10]. As shown in Figure 1, this snake is defined by only three points. Its energy function, is defined in such a way that the snake will get attracted by bright blobs in the image. In doing so, fast gradient-based optimization algorithms can be used. We embed the ovuscul in the particle filter in a way that the importance sampling of the particle filter is defined implicitly by the optimization algorithm of the variational method, and the particle weights correspond to the optimal values of the energy function of each individual particle. This construction drastically reduces the number of particles needed to have an accurate description of the target. We make use of the information provided by the ovuscul in order to detect the start and end of the mitotic stage within a simplified cell cycle, and use different motion models accordingly.

The paper is structured as follows: We first recall elements of the particle-filter framework and the ovuscul in Section 2. Then, we describe our algorithm in Section 3, and illustrate its capabilities by tracking mitotic HeLa cells and outlying their nuclei in Section 4. We finally conclude in Section 5.

2. METHODS

2.1. Particle-Filter Framework

Before describing our tracking method, we review the basic principles of particle filtering. The Bayesian-tracking framework provides a methodology to infer hidden states of a dynamic system $\mathbf{x}_{1:t} = \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$, using a sequence of noisy measurements $\mathbf{z}_{1:t} = \{\mathbf{z}_1, \dots, \mathbf{z}_t\}$. Bayesian estimation is used to recursively compute a time-evolving posterior distribution $p(\mathbf{x}_t | \mathbf{z}_{1:t})$. This distribution can be estimated by assuming a Markovian model of the state evolution, $D(\mathbf{x}_t | \mathbf{x}_{t-1})$, and a likelihood that relates the noisy measurements to the hidden state $L(\mathbf{z}_t | \mathbf{x}_t)$. Then, the probability density function (pdf) $p(\mathbf{x}_t | \mathbf{z}_{1:t})$ is estimated in two steps: prediction of the state and update after the new measurement \mathbf{z}_t is

available.

In the prediction step, the system model and the estimated posterior density from the previous frame $p(\mathbf{x}_{t-1}|\mathbf{z}_{1:t-1})$ are combined in the Chapman-Kolmogorov equation to obtain the prior density

$$p(\mathbf{x}_t|\mathbf{z}_{1:t-1}) = \int D(\mathbf{x}_t|\mathbf{x}_{t-1}) p(\mathbf{x}_{t-1}|\mathbf{z}_{1:t-1}) d\mathbf{x}_{t-1}. \quad (1)$$

Next, in the update step, Bayes rule is used to modify the prior density and obtain the desired posterior

$$p(\mathbf{x}_t|\mathbf{z}_{1:t}) \propto L(\mathbf{z}_t|\mathbf{x}_t) p(\mathbf{x}_t|\mathbf{z}_{1:t-1}). \quad (2)$$

The solution of the problem defined by (1) and (2) is analytically tractable in a limited number of cases (e.g., linear Gaussian models). For most practical models, sequential Monte Carlo methods are used as an efficient approximation. In these methods, the posterior $p(\mathbf{x}_t|\mathbf{z}_{1:t})$ is represented with a set of N_s random weighted samples $\{\mathbf{x}_t^{(i)}, w_t^{(i)}\}_{i=1}^{N_s}$ as

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

where $\delta(\cdot)$ is the Dirac delta and the sum of the weights is normalized to one.

The particles are chosen using the principle of importance sampling. This principle relies on the availability of an importance function $q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{z}_t)$ that describes the state space. The idea is to sample the areas of the state space where the importance function is large and to avoid generating samples with low weights, since they provide a negligible contribution to the posterior. Thus, we generate a set of new particles using the importance function, that is

$$\mathbf{x}_t^{(i)} \sim q(\mathbf{x}_t|\mathbf{x}_{t-1}^{(i)}, \mathbf{z}_t). \quad (3)$$

Generally, the importance function can be chosen arbitrarily. The only requirements are the possibility to easily draw samples from it, and to have the same support as $p(\mathbf{x}_t|\mathbf{z}_{1:t})$. When using the importance density function $q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{z}_t)$, the expectation of any function $f(\mathbf{x}_t)$ with respect to the probability $p(\mathbf{x}_t|\mathbf{z}_{1:t})$ can be rewritten as

$$\begin{aligned} & \int f(\mathbf{x}_t) p(\mathbf{x}_t|\mathbf{z}_{1:t}) d\mathbf{x}_t \\ &= \int f(\mathbf{x}_t) \frac{p(\mathbf{x}_t|\mathbf{z}_{1:t})}{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{z}_t)} q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{z}_t) d\mathbf{x}_t, \end{aligned}$$

where the integration is performed over the support of $p(\mathbf{x}_t|\mathbf{z}_{1:t})$ and $q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{z}_t)$. By drawing N_s samples as in (3), the expectation can be approximated as

$$\int f(\mathbf{x}_t) p(\mathbf{x}_t|\mathbf{z}_{1:t}) d\mathbf{x}_t \approx \sum_{i=1}^{N_s} f(\mathbf{x}_t^{(i)}) w_t^{(i)}, \quad (4)$$

where

$$w_t^{(i)} \propto \frac{p(\mathbf{x}_t^{(i)}|\mathbf{z}_{1:t})}{q(\mathbf{x}_t^{(i)}|\mathbf{x}_{t-1}^{(i)}, \mathbf{z}_t)},$$

and $\sum_{i=1}^{N_s} w_t^{(i)} = 1$. Thus, the Chapman-Kolmogorov equation can be approximated using the right hand side of (4). Taking advantage of the fact that we have a good observation model given by the variational method, we propose to replace the classical importance sampling function by the optimizer of the variational scheme. This

novel approach drives the particles towards regions in the state space with high probability.

2.2. The Ovuscule

In a previous paper, we introduced a minimalistic active contour named the ovuscule [10]. This snake takes the shape of an ellipse; it is the simplest active contour capable of capturing orientation and anisotropy. The elliptic shape of the ovuscule makes it a very robust cell-segmentation algorithm in poor imaging conditions. Moreover, it is fast to compute.

2.2.1. Parameterization

Following the traditional definition of parametric active contours, the parameterization of the ovuscule was designed so that its parameters correspond to control points on the outline of the snake. Since an ellipse is given by five parameters, we need at least three control points, named $\{\mathbf{p}, \mathbf{q}, \mathbf{r}\}$, to fully determine an ellipse. A full description of the process can be found in [10]. By using this parameterization, all parameters have equal importance; this creates a favorable landscape for the proceedings of the optimizer. This parameterization is also advantageous in the Bayesian framework if we define the state space to be $\mathbf{x}_t = (\mathbf{p}, \mathbf{q}, \mathbf{r})$. Under these circumstances, it is easier to define motion models for the three ovuscule points rather than other possible parameterizations of an ellipse such as foci and arc-length, or eccentricity, elongation, and orientation.

2.2.2. Snake Energy

The ovuscule is a region snake [10]. During the optimization process we tune the geometry of the ovuscule to increase the contrast between the intensity of the data averaged over an elliptical core, and the intensity of the data averaged over an elliptical shell, as shown in Figure 1. If Γ and Γ' represent these elliptical surfaces, with $\Gamma' \subset \Gamma$, and if I is our image data, then the criterion to minimize is $J = J_D + J_R$, where J_R is a contribution due to some regularization term in order to cope with the extra degree of freedom of the parameterization, and where the data term J_D is ideally given by

$$J_D = \frac{1}{|\Gamma|} \left(\int_{\Gamma \setminus \Gamma'} I(x, y) dx dy - \int_{\Gamma'} I(x, y) dx dy \right). \quad (5)$$

3. VARIATIONAL IMPORTANCE SAMPLING

In our setting, the ovuscule provides an accurate observation model that describes the elliptical shape of our nuclei within the image. In such circumstances, the state vector corresponds to the triplet of points $\mathbf{x}_t^{(i)} = (\mathbf{p}^{(i)}, \mathbf{q}^{(i)}, \mathbf{r}^{(i)})$ of the ovuscule, and the measurement vector corresponds to the pixel values $\mathbf{z}_t = \{I_t\}$ of the image.

At each frame, we propagate each particle of the previous frame following the state evolution model, which generates the predicted set of particles $\{\tilde{\mathbf{x}}_t^{(i)}\}_{i=1}^{N_s}$. Since each particle $\tilde{\mathbf{x}}_t^{(i)}$ is built from an ovuscule, it can be associated with an energy value $J^{(i)}$ measuring the goodness of fit of the ovuscule to the target being tracked. We optimize the energy value of the predicted set of particles following the gradient-based optimizer of the ovuscule. This defines the optimized set of particles $\{\mathbf{x}_{t,\text{opt}}^{(i)}\}_{i=1}^{N_s}$ with an optimized set of ovuscule energies $\{J_{\text{opt}}^{(i)}\}_{i=1}^{N_s}$. Following the principle of maximum entropy,

we assume that J_{opt} is a random variable with exponential distribution, which leads to assign the particle weights $w_t^{(i)}$ to

$$w_t^{(i)} \propto e^{-\lambda J_{\text{opt}}^{(i)}},$$

where λ is a parameter that controls the sharpness of $p(\mathbf{x}_t | \mathbf{z}_{1:t})$.

Using the proposed scheme, the importance sampling of the particle filter, usually performed by (3), is given implicitly by the optimization algorithm of the variational method. This interpretation arises naturally since the role of the optimizer is to attract the ovusculi, and therefore the particles, to the target under inspection. As a consequence, the weights of the particles within the region of convergence of the optimizer will gain importance compared to the ones that are not. Therefore, a much smaller set of particles is necessary to describe the high probability regions of the state space.

Finally, we perform a resampling step to eliminate particles that have small weights and to focus on particles with large weights. The resampling step involves generating a new set by sampling (with replacement) N_s times from $\{\mathbf{x}_{t,\text{opt}}^{(i)}\}_{i=1}^{N_s}$, which leads to the equiprobable set of particles $\{\mathbf{x}_t^{(i)}, \frac{1}{N_s}\}_{i=1}^{N_s}$. The estimation at each frame of the location and shape of the target being tracked at each frame can be carried out efficiently with the MAP estimator as follows:

$$\hat{\mathbf{x}}_t = \arg \max_{\mathbf{x}_t} \{p(\mathbf{x}_t | \mathbf{z}_{1:t})\} \approx \arg \max_i \{w_t^{(i)}\}.$$

Thus, the *maximum a posteriori* (MAP) estimation of the target corresponds to the optimized particle with highest weight. All these operations are summarized in Algorithm 1.

Algorithm 1: Ovusculi-based particle filter

input: Particle set $\{\mathbf{x}_{t-1}^{(i)}\}_{i=1}^{N_s}$ and current image I_t
output: MAP estimation $\hat{\mathbf{x}}_t$ and particle set $\{\mathbf{x}_t^{(i)}\}_{i=1}^{N_s}$

for $i \leftarrow 1$ **to** N_s **do**
 $\tilde{\mathbf{x}}_t^{(i)} \leftarrow$ Propagate $\mathbf{x}_{t-1}^{(i)}$ with the motion model;
 $\{\mathbf{x}_{t,\text{opt}}^{(i)}, J_{\text{opt}}^{(i)}\} \leftarrow$ Adjust the ovusculi to I_t ;
 $w_t^{(i)} \leftarrow e^{-\lambda J_{\text{opt}}^{(i)}}$;
end
for $i \leftarrow 1$ **to** N_s **do**
 $w_t^{(i)} \leftarrow w_t^{(i)} / \sum_j w_t^{(j)}$;
end
 $\hat{\mathbf{x}}_t \leftarrow \arg \max_i \{w_t^{(i)}\}$
 $\{\mathbf{x}_t^{(i)}\}_{i=1}^{N_s} \leftarrow$ **Resampling**($\{\mathbf{x}_{t,\text{opt}}^{(i)}, w_t^{(i)}\}_{i=1}^{N_s}$);

4. APPLICATION

In this section, we apply our ovusculi-based particle filter to construct the lineage of migrating HeLa cells, and outline their nuclei.

4.1. Biphasic Motion Model

For our particular application, two different motion models are considered depending on the state of the cell cycle. Both models are considered to be linear, with

$$\tilde{\mathbf{x}}_t = \mathbf{x}_{t-1} + \mathbf{n}_t, \quad (6)$$

where \mathbf{n}_t is a random vector that depends on the state of the cell. The two cell states are:

- Non-mitotic, where the nuclei are essentially circular, and move and deform without any preferred direction, as shown in Figure 2 (a)-(b);
- Mitotic, where nuclei are more elongated and brighter than in the non-mitotic state, and where the movement during the splitting is fast and perpendicular to the main axis of the cell, as shown in Figure 2 (c)-(d).




Fig. 2. Migrating HeLa nuclei: (a) non-mitotic state at time $(t-1)$, (b) non-mitotic state at time t , (c) mitotic state at time $(t-1)$, (d) mitotic state at time t

For the non-mitotic stage, the natural choice in (6) is to assume Gaussianity and independence for each component of \mathbf{n}_t . For the mitotic stage, we adopt a purely translational model perpendicular to the main orientation axis. A cell is considered to enter in the mitotic state if its MAP estimation is brighter and more eccentric than a certain threshold values. At that point, the motion model switches to the mitotic one, and eventually returns to the non-mitotic one once the values of the brightness and eccentricity get below the thresholds.

4.2. Experimental Results

To illustrate our method, we applied our algorithm to a time-lapse sequence of images of HeLa nuclei expressing fluorescent core histone 2B on an RNAi live cell array¹. We focused on building the cell lineage of a single cell. We only used a total of 20 particles. The thresholds, λ , and the standard deviations for \mathbf{n}_t were chosen empirically.

In Figure 3, we show the behavior of Algorithm 1 when a non-mitotic motion model is used. In particular, we observe in Figure 3 (a) the outlines of the ovusculi representing the particles from the previous frame. These particles are propagated following the non-mitotic motion model to the locations shown in Figure 3 (b). After optimizing the ovusculi, we obtained the particles shown in Figure 3 (c), and, finally, after the resampling, the particles in Figure 3 (d). Note that, after the optimization, one ovusculi converged to a local minima, but its weight was negligible compared to the others. Therefore, it was eliminated in the resampling step. In Figure 4, we show the behavior of Algorithm 1 when a mitotic motion model is used and when the cell division occurs. In particular, we observe in Figure 4 (a) the outlines of the ovusculi representing the particles from the previous frame located at the same position. These particles are propagated following the mitotic motion model to the locations shown in Figure 4 (b). After running the ovusculi optimizer we obtained the particles shown in Figure 4 (c), and after the resampling we obtained the particles shown in Figure 4 (d). Note that, after the optimization, some ovusculi converged to different targets, and this information was preserved in the resampling step.

¹Courtesy of D. Gerlich, Institute of Biochemistry, ETHZ, Zürich.




Fig. 3. Different steps of Algorithm 1 with non-mitotic motion model. (a) Initial. (b) Propagated. (c) Optimized. (d) Resampled.




Fig. 4. Different steps of Algorithm 1 with mitotic motion model. (a) Initial. (b) Propagated. (c) Optimized. (d) Resampled.

We show in Figure 5 the temporal evolution of the mean value and the eccentricity of a single nucleus. We can observe a simultaneous peak in both graphs between frames 180 and 188, which corresponds to the mitotic stage of the cell.




Fig. 5. (a) Mean intensity within the elliptical MAP estimation. (b) Eccentricity of the elliptical MAP estimation. The mitotic stage occurs between Frame 180 and 188.

The use of our biphasic motion model would not have been possible if we had not used the optimized ovusculum to obtain an accurate estimation of the orientation of the cell during the mitotic stage with a reasonable number of particles. Moreover, thanks to the capability of the particle filter to describe multimodal distributions, our algorithm is capable of building the cell lineage, which the ovusculum could not have achieved on its own.

The computation time is usually directly related to the number of particles used in the particle filter. Since our variational importance sampling provides a better description of the high probability regions of $p(\mathbf{x}_t | \mathbf{z}_{1:t})$, a reduced number of particles is necessary. Moreover, the optimization of each ovusculum during the variational importance sampling stage can be carried out independently, thus, the algorithm is fully parallelizable.

The method described in this article has been programmed as a plugin for ImageJ, which is a free open-source multiplatform Java image-processing software². Our plugin does not depend in any spe-

cific imaging hardware; thanks to ImageJ, any common file format may be used.

5. CONCLUSION

We have proposed a new methodology that fuses in a single tracker the two major tracking philosophies and that retains the advantages of both. We showed that, by using a robust variational method, it is possible to replace the importance sampling function of the particle filter and obtain an alternative scheme. The resulting algorithm is capable of creating an accurate segmentation of elliptic targets with a reduced number of particles, and capable of detecting and tracking cells undergoing mitosis.

6. REFERENCES

- [1] C. Zimmer, B. Zhang, A. Dufour, A. Thébaud, S. Berlemont, V. Meas-Yedid, and J.-C. Olivo-Marin, "On the digital trail of mobile cells," *IEEE Signal Processing Magazine*, vol. 23, pp. 54–62, January 2006.
- [2] 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, pp. 174–188, February 2002.
- [3] K. Miura, "Tracking movement in cell biology," *Advances in Biochemical Engineering/Biotechnology*, vol. 95/2005, pp. 267–295, June 2004.
- [4] C. Zimmer, E. Labryère, V. Meas-Yedid, N. Guillén, and J.-C. Olivo-Marin, "Segmentation and tracking of migrating cells in videomicroscopy with parametric active contours: A tool for cell-based drug testing," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1212–1221, January 2002.
- [5] J. Cui, N. Ray, S. T. Acton, and Z. Lin, "An affine transformation invariance approach to cell tracking," *Computerized Medical Imaging and Graphics*, vol. 32, pp. 554–565, October 2008.
- [6] O. Dzyubachyk, W. van Cappellen, J. Essers, W. Niessen, and E. Meijering, "Advanced level-set-based cell tracking in time-lapse fluorescence microscopy," *IEEE Transactions on Medical Imaging*, vol. 29, pp. 852–867, March 2010.
- [7] M. Isard and A. Blake, "CONDENSATION: Unifying low-level and high-level tracking in a stochastic framework," in *Proceedings of the 5th European Conference on Computer Vision (ECCV'98)*, vol. 1, (Freiburg, Germany), pp. 893–908, June 1998.
- [8] M. Bray, E. Koller-Meier, and L. V. Gool, "Smart particle filtering for high-dimensional tracking," *Computer Vision and Image Understanding*, vol. 106, pp. 116–129, April 2007.
- [9] Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi, "Tracking deforming objects using particle filtering for geometric active contours," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, pp. 1470–1475, August 2007.
- [10] P. Thévenaz, R. Delgado-Gonzalo, and M. Unser, "The ovusculum," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 382–393, February 2011.

²<http://rsb.info.nih.gov/ij/>