

Multiple Hypothesis Tracking for Cluttered Biological Image Sequences

Nicolas Chenouard, Isabelle Bloch, *Member, IEEE*, and Jean-Christophe Olivo-Marin, *Fellow, IEEE*

Abstract—In this paper, we present a method for simultaneously tracking thousands of targets in biological image sequences, which is of major importance in modern biology. The complexity and inherent randomness of the problem lead us to propose a unified probabilistic framework for tracking biological particles in microscope images. The framework includes realistic models of particle motion and existence and of fluorescence image features. For the track extraction process per se, the very cluttered conditions motivate the adoption of a multiframe approach that enforces tracking decision robustness to poor imaging conditions and to random target movements. We tackle the large-scale nature of the problem by adapting the multiple hypothesis tracking algorithm to the proposed framework, resulting in a method with a favorable tradeoff between the model complexity and the computational cost of the tracking procedure. When compared to the state-of-the-art tracking techniques for bioimaging, the proposed algorithm is shown to be the only method providing high-quality results despite the critically poor imaging conditions and the dense target presence. We thus demonstrate the benefits of advanced Bayesian tracking techniques for the accurate computational modeling of dynamical biological processes, which is promising for further developments in this domain.

Index Terms—Particle tracking, biological imaging, multiple hypothesis tracking, target perceivability, cluttered images

1 INTRODUCTION

1.1 Particle Tracking in Biology

MAJOR biological processes such as morphogenesis, pathogens invasion strategies, and immune response to infections involve a very complex machinery that is inherently governed by highly dynamic processes both at the molecular and the cellular levels. These processes generally involve nanometric objects (e.g., virus complexes, bacteria, receptors, mRNA, vesicles) that are observed with very high detail over extended periods of time thanks to numerous recent advances in fluorescent probes, labeling techniques, and microscopy systems [47]. When imaged through microscopes, these fluorescently labeled objects appear as compact and bright objects termed *particles* in the following. Typical datasets consist of several hundred 2D or 3D stacks of time-lapse digital images, each of which contains hundreds or thousands of particles whose trajectories need to be established to infer information about the underlying molecular mechanisms. In that context, multiple particle tracking (MPT) is of major relevance as it provides accurate measurements of single particle dynamics at the appropriate spatial and temporal scales [1], [10], [29], [37]. MPT is plagued, however, by several complicating factors. Proper identification of particle positions is the first [44],

[41]: Variable or low fluorescence intensity, structured backgrounds, electronic noise, momentary aggregation, and coalescence are just examples of factors impeding the detection of particles. The overall unreliability of particle detection consequently yields corrupted measurements: Some particles are missed by the detection procedure, coalesced particles are detected as one, some artifacts are wrongly considered as objects of interest. Determining the correct number of existing targets is yet another major issue: The number of particles is unknown beforehand since many objects may enter and leave the field of view during the acquisition process, while some particles can physically disappear and appear.

The characteristics of biological systems make the tracking stage per se both a challenging and an original problem. By contrast with many computer vision tracking tasks, the target profile in images is generally not a relevant feature for identifying particles through frames due to their similar small and isotropic shape and the random nature of the photon counting process. Thus, one can rely only on the particle spatial position, and object motion needs to be extensively exploited to resolve particle correspondence. As compared to the usual targets in video tracking problems, moreover, nanometric objects usually exhibit very erratic complex motion dynamics that make particle movements hard to predict. This issue is exacerbated by the low acquisition frame rates, as compared to the mechanism dynamics, that are typical of nanometer scale systems because of the extended period of time that is required to collect sufficient amount of photons from such small fluorescent objects.

Faced with all the above issues and because visual interpretation of the data proves to be both cumbersome and limited in practice, the bioimaging community is craving automated quantification methods able to process large-scale dynamic imaging data in a robust and systematic manner.

• N. Chenouard and J.-C. Olivo-Marin are with the Quantitative Image Analysis Unit, Institut Pasteur, 25 rue du docteur Roux, 75724 Paris Cedex 15, France, and CNRS, URA 2582, Paris, France.

E-mail: niclas.chenouard@gmail.com.
 • I. Bloch is with Telecom ParisTech, CNRS LTCI, 46 rue Barrault, F-75634 Paris Cedex 13, France.

Manuscript received 24 Jan. 2012; revised 25 Oct. 2012; accepted 3 Apr. 2013; published online 17 May 2013.

Recommended for acceptance by T. Cootes.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2012-01-0059.

Digital Object Identifier no. 10.1109/TPAMI.2013.97.

Early methods, typically relying on linking nearest detections from frame to frame [11], do not comply with the requirements of modern time-lapse imaging systems though, and new paradigms are required for modern microscopy [38], [36], [32]. In this paper, we present a novel Bayesian framework for MPT along with a dedicated tracking algorithm that bring original solutions to this challenge.

1.2 A Unified Bayesian Framework for Object Tracking in Biology

Although a few attempts to use Bayesian paradigms for MPT have been made in the past [18], [40], [42], the framework was only partly exploited until now for bioimaging. Bayesian tracking is a reference paradigm for complex multitarget tracking systems (e.g., radar tracking [7], computer vision [15]) because it allows accounting for statistics of the sensors, target motion characteristics, as well as the random inaccuracies of the tracking models [2]. This flexibility has led us to propose in this paper a unified Bayesian model for tracking multiple targets in microscopy image sequences, from which the optimization of different statistical criteria will result in specific tracking algorithms. Since target dynamics is a major feature for discriminating particles, we first build an accurate statistical motion model for nanometric objects that is inspired from biophysics. It consists of a hidden Markov model with multiple states whose parameters can be directly linked to the physical properties of the environment and of the targets. We also define a target existence model for biological particles to account for the possible appearance of new targets and the disappearance of some others. By doing so, track creation and termination mechanisms are seamlessly embedded into the statistical tracking procedure. Contrary to previous works by other groups [20] or ourselves [18], we do not consider here the case of target splitting and merging in an explicit manner, but rather consider it as part of a postprocessing step. We acknowledge that this could be seen as a disadvantage over previously cited works, especially so with [20], where fusion and fission of particles are integrated in the model. We, however, preferred in this work to concentrate on the problem of non or weakly interacting particles and defer a global solution incorporating split-and-merge events explicitly within our scheme to future work. Finally, the characteristics of the acquisition system are also included in the model through proper probabilistic density functions (pdf) for the false detections and the detection rate of particles. As a result, the framework is flexible and exhaustive enough to be adapted to a wealth of different biological cases.

1.3 Multiple Hypothesis Tracking for Cluttered Data

We define the track construction process as the maximization of the tracks likelihood in the proposed framework. By doing so, the resulting solution automatically accounts for the various aspects of the tracking problem (target existence, particle motion, etc.). The maximum likelihood (ML) estimator is preferred in our context because it helps in preserving track identity (better than the minimum mean squared error for instance) which is of first importance for biological studies. The large-scale nature of the problem,

however, makes the solution intractable in the general case and hampers the use of the most advanced estimation tools that were demonstrated on significantly smaller problems [24], [46]. We thus rewrite the optimization problem as an iterative multiframe tracking procedure, similar to the multiple hypothesis tracking algorithm (MHT) [33], [7], but using the proposed model. To track particles in extremely low signal-to-noise ratio (SNR) conditions, we present an original approach that contrasts with state-of-the-art techniques for MPT. We propose to use the MHT in a novel approach: We use very sensitive detection settings and exhaustively construct and maintain tracking hypotheses, many of which result from clutter, before deciphering correct ones when maximizing the tracks likelihood. The robustness of the proposed MHT to extremely cluttered data results from both the accurate statistical model of the tracking problem and from the integration of information from multiple future frames. We investigate the performance of this technique in two tracking scenarios (synthetic and real data) and show that it significantly outperforms existing tracking techniques for bioimaging in very low SNR conditions. The capability to recover large numbers of low intensity particle trajectories despite poor imaging conditions opens the way to new types of biological studies and is promising for the development of advanced Bayesian tracking methods for bioimaging.

1.4 Outline of the Paper

In Section 2, we present a short survey of existing MPT techniques in biological imaging, and we sketch several Bayesian tracking paradigms. We then propose in Section 3 a unified Bayesian framework for tracking multiple particles in microscopy images. A specialized MHT is then introduced in Section 4. In Section 5, we present a comprehensive study of the performance of the proposed MHT with highly cluttered images. Finally, we discuss our main results and future directions of research.

2 RELATED WORK

2.1 Non-Bayesian Methods for Target Tracking in Biology

While particle detection in microscope images is a fairly well understood problem and little hope is left for major improvement (we refer to [41] for a comparison of latest methods for this task), the failure of nearest-neighbor methods for particle tracking has led to developments in several directions. For instance, in the same vein as early MPT works, a distance-based multiframe tracking technique has been proposed in [38]. The tracking problem is expressed over multiple frames as the minimization of the sum of all track lengths. However, as shown in Section 5, this algorithm does not accommodate highly cluttered conditions and important information such as target dynamics is not accounted for. An alternative approach to this traditional MPT approach is to alleviate the preliminary frame-to-frame detection stage by simultaneously detecting and tracking particles so as to exploit further the temporal information and the smoothness of trajectories. As an example, in [36], the tracking task is formulated as a minimal path problem when considering the 2D image

sequence as a 3D spatiotemporal volume. Such an approach allows extracting the trajectory of an extremely dim target. However, multiple targets can be dealt with only when they are well apart, which is seldom the case in biology. Along the same line, a minimal path technique has been proposed in [9] to fill the gaps between partial track segments obtained by nearest detection linking. As a result, blinking targets were robustly tracked in 2D image sequences. An alternative when particles use a predefined path is to use a lower dimensional representation of the dataset, called kymograph, in combination with a line-extraction method for identifying particle trails [32].

While analyzing time-lapse sequences as a whole constitutes a step further in the improvement of MPT, these methods suffer from a lack of generality that makes them hard to use for biological applications. For example, the fundamental assumption, if available, pertaining the use of the kymograph representation is realistic only for a limited number of biological mechanisms. On the other side, leaving aside the computational aspects for large-scale 3D problems, a main issue of the minimal-path approach is the requirement for targets to produce long and smooth curves in the spatiotemporal volume. This is the case only when particles have a slow motion as compared to the acquisition rate, with almost overlapping positions in subsequent frames, by contrast with the majority of cell biology experiments. Moreover, prior knowledge on the system, particle motion characteristics, and physical data are not used by this kind of techniques.

2.2 Bayesian Particle Tracking

2.2.1 Preliminaries and Notations

In discrete time, each target is generally described at time $k \geq 1$ by a vector $\mathbf{x}(k) \in \mathbb{R}^n$ of n state variables representing its coordinates and characteristic features. The sequence of states from time 1 to k is denoted by $\mathbf{x}^k = \{\mathbf{x}(t)\}_{t=1\dots k}$. We assume that the state of a target depends on the previous states through the *state evolution equation*:

$$\mathbf{x}(k+1) = f_k(\mathbf{x}^k) + \mathbf{w}(k+1). \quad (1)$$

The vector $\mathbf{w}(k+1) \in \mathbb{R}^n$ is a random and additive process noise at discrete time $k+1$, while the vector-valued function f_k characterizes the deterministic part in the evolution from states \mathbf{x}^k to $\mathbf{x}(k+1)$ and is possibly time varying. The sequence of measurements associated with the state sequence is $\mathbf{z}^k = \{\mathbf{z}(t)\}_{t=1\dots k}$ and, in the standard model, the current measurement is related to the state through the *measurement equation*:

$$\mathbf{z}(k) = h_k(\mathbf{x}(k)) + \mathbf{v}(k).$$

The function $h_k: \mathbb{R}^n \rightarrow \mathbb{R}^m$ links the measured features to the state variables, while the noise $\mathbf{v}(k) \in \mathbb{R}^m$ is a random counterpart modeling uncontrolled fluctuations of the measurement process. Many standard techniques exist for the estimation of \mathbf{x}^k from \mathbf{z}^k and the computation of the related pdfs, the most common ones being the Kalman filter (KF) and its extensions [3], [34], and sequential Monte Carlo (SMC) methods (see [17] and references therein).

2.2.2 Bayesian Tracking in Biology

In biology, the Bayesian approach was pioneered with the algorithm proposed in [18] and used for virus tracking [1], which relies on a traditional two-step scheme: 1) particle detection, 2) association of detections with tracks. The point-correspondence problem between subsequent frames is solved with a nearest-neighbor linking algorithm that uses the probability of one given detection originating from an active track as the association cost. The probability of association is provided by a statistical description of particle motion relying on random switches between several basic kinetic models. The pdf estimation is efficiently achieved thanks to the interacting multiple models filter (IMM) [8]. Missed detections are dealt with by allowing tracks to associate with virtual detections for a number of frames before being terminated when the maximal number of subsequent virtual detections is exceeded. New tracks are automatically created from the remaining detections at the end of the association stage. Along the same line, a frame-to-frame tracking algorithm with probabilistic costs for associations was proposed in [20]. A mixture of motion models is used for particles, and linking detections and tracks at a given time point is viewed as a global optimization problem. Moreover, a specific postprocessing for linking partial track segments is proposed. The ability of the technique to track some membrane receptors in microscope images of cells has been demonstrated [20].

Both works in [18] and [20] are appealing because the improved motion models they use better account for particle dynamics. Also, the frame-to-frame detection linking procedures comply with the computational requirements of large-scale applications and specific postprocessing strategies have been developed to account for target splitting and merging events. However, the use of primitive heuristics for track termination and creation, and the lack of an accurate model to handle false and missing detections, hamper their use for very low SNR applications (as shown in Section 5).

Different methods relying on dedicated SMC methods have been proposed in [40] and [42]. The key idea is to use a joint state-space for all the targets and approximate the related multimodal pdfs using a set of carefully chosen random samples from this space. Thanks to this technique, several biological objects (membrane receptors, vesicles, microtubule tips) have been successfully tracked over image sequences of fluorescence microscopy [40]. Large-scale applications are, however, not tractable in this framework, closely spaced targets may lead to some track coalescence issues, and target existence is dealt with using external heuristics.

2.2.3 Bayesian Methods for Multiple Target Tracking

A large body of work is dedicated to Bayesian multiple target tracking in clutter, in particular for radar tracking and surveillance applications (see [7] and references therein). Numerous specific methods have been proposed to deal with nonlinear measurement systems, multiple sensors, maneuvers, and so on. However, the focus of many of these methods is less relevant for biological object tracking due to the differences that exist both between the acquisition

systems and the targets' nature. For instance, aerospace vehicles generally exhibit smooth and predictable trajectories, while particle motion is very erratic and unpredictable. This lack of adaptation is illustrated in Section 5.2 by the poor performances of a multiframe algorithm using a smoothness-promoting model [14] that is typical of surveillance applications.

The aim of the tracking procedure itself may also significantly differ depending on the field of applications. For instance, considerable effort has been dedicated recently to designing tracking methods for which preserving the target label is not required since for some aerospace applications the accurate assessment of objects' presence and trajectory is sufficient data. In particular, algorithms relying on the optimal subpattern assignment (OSPA) metric [39] have been developed over the past few years: set joint probabilistic data association (Set JPDA) in [43], minimum mean OSPA estimator in [19]. By contrast, in biology, preserving track identity is usually of first importance because of the downstream analysis that is performed to assess single target dynamics and to characterize biological mechanisms.

In the literature, general and powerful estimation methods such as the probabilistic hypothesis density (PHD) filter and its derivatives have recently become popular methods for multiple target state estimation [24] because of their claimed improved state estimation accuracy over standard methods [26].

However, at the time this work was done, the computational cost of these promising techniques has limited their use to tracking only few targets, despite some efforts to improve their efficiency [46], [25]. To tackle the large-scale dimensions of the particle tracking problem without compromising on the quality of the solution, we propose in this paper to adapt the MHT algorithm to the crucial features of the system. The MHT algorithm is a point-correspondence technique that is multiframe in the sense that several future frames are accounted for before associating detections with tracks at the current time-point [33]. This allows robust tracking as future frames information improves the significance of the tracking scores. In the sequel, we focus on the *track-oriented* implementation of the MHT [21], [5]. In this framework, a tracking hypothesis Θ^k up to time k is defined as a set of m tracks $\Theta^k = \{\theta_j^k\}_{j=1..m}$, each one consisting of some subsequent measurements and estimated states: $\theta_j^k = \{\mathbf{z}_j^k, \mathbf{x}_j^k\}$. The putative extensions of each track from frame k to $k+d$ are represented by trees of associations with measurements, as shown in Fig. 1 for a prototypical example. Some *virtual* detections can also be incorporated to cope for missed detections. The crucial point is that the set of global tracking hypotheses up to time $k+d$ when Θ^{k-1} is fixed, which we denote $\Omega_{\Theta^{k-1}}^{k:k+d}$, can be obtained by combining these independent trees. The technique thus saves considerable memory, for the size of the track trees is much smaller than the size of $\Omega_{\Theta^{k-1}}^{k:k+d}$. As a consequence, pruning is not required anymore and the optimal tracking solution may be derived. The main issue is to select the best combination of branches from different trees to form the optimal hypothesis (see [7] for a review). To build approximate solutions to this problem, efficient




Fig. 1. Construction of a tree of potential tracks in two subsequent frames for a single track. Real measurements are represented by a circle and virtual measurements by a square. The leaves correspond to the set of potential tracks that originate from the track at time $k-1$.

techniques have been developed based on tree exploration [21], [6] and multidimensional assignment [31]. For biological applications, however, computational performance is often less a concern than reproducibility and quality of results. Indeed, by contrast with some other fields, data analysis is usually performed offline and typical computation times are small compared to the other steps of biological experiments that are on scales of days, and sometimes of weeks or months. We will thus adapt a linear programming approach [27] to achieve optimal solutions for our tracking model.

3 A BAYESIAN MODEL FOR PARTICLE TRACKING IN CLUTTER

In this section, we propose a Bayesian framework for MPT in dense clutter. We introduce key concepts such as particle existence and realistic models of motion for nanometer-scale objects in cellular environments. It is demonstrated in Section 4 that this comprehensive model allows one to design extremely robust tracking algorithms by exploiting the statistical quantities described hereafter.

Our model aims at providing an accurate description of the conditional probability $\mathcal{L}(\Theta^K) = P\{\Theta^K|Z^K\}$ that relates a candidate set of tracks Θ^K to the set of measurements Z^K in an image sequence of length K . To do so, we exploit the fact that tracks are independent, as can be seen by rewriting $\mathcal{L}(\Theta^K)$ as

$$\mathcal{L}(\Theta^K) \propto \prod_{k=1..K} p\{Z_0(k)\} \cdot \prod_{\theta_j \in \Theta^K} p(\mathbf{x}_j^K, \mathbf{z}_j^K), \quad (2)$$

where $Z_0(k)$ is the set of detections that remain nonassociated at time k . As stated above, split and merge events are not considered explicitly; we assume, however, that short split and merge events will be detected by the MHT as missing detections/occlusions, while split and merge events of similar duration or longer than the sliding time window will not, thus requiring a postprocessing step as in [18]. We propose next a model for computing each track joint probability: $\chi_j^K = p(\mathbf{x}_j^K, \mathbf{z}_j^K)$.

3.1 Embedded Target Existence Model

We derive here a model of particle existence that is inspired by the work of Li and Li [22], [23] on target perceivability and relates to the framework of particle existence in [45].

One main complicating factor for target existence estimation is its dependence on the associations between tracks and detections that should lead to jointly estimating the existence and associations simultaneously for all the targets. Different techniques, such as a JPDA-like method and an SMC approach, have been proposed in [45] to deal with this issue. The problem is naturally alleviated in the MHT framework we propose in the sequel because we estimate target perceivability independently for all the putative tracks that are enumerated by the tree construction algorithm. It is worth noting that the integration of a model of target existence in a MHT framework has been previously proposed by Musicki and Evans in [28] in a different applicative context, but multiple putative tracks (tree branches in Fig. 1) were integrated to obtain the probability of track existence.

The perceivability of a target is defined as its capability of generating measurements [22]. For particles in general, we define a two states model of perceivability:

- a particle is *perceivable* (state s^1) if it can be detected,
- a particle is *nonperceivable* (state s^0) if it physically does not exist anymore or cannot be detected in images (e.g., fluorescence extinction, degradation).

In biology, some particles may also exhibit a blinking intensity pattern (e.g., quantum dots). In this case, we can define a third perceivability pattern s^2 that is characterized by a missed detection due a temporary fluorescence extinction. However, this state is transient only, with the possibility of returning to a normal intensity mode with a high probability (according to some statistical law characterizing the fluorophore). For the sake of space, we will consider in this paper the two states case only, which is more common. Calculations are essentially the same for the three states model.

The perceivability state sequence is modeled as a first-order Markov chain whereby the current perceivability state $s^i(k)$ depends on the previous perceivability state $s^i(k-1)$ but not on the earlier state sequence. We also define fixed transition probabilities between the subsequent states i and j , noted π_{ji} . The Markov assumption yields the computation of the probability of being in the state s^i at time k as the probability of transition from every state at time $k-1$ to $s^i(k)$:

$$p(s^i(k)) = \sum_{j \in \{0,1\}} \pi_{ji} p(s^j(k-1)).$$

We include the target perceivability concept in the Bayesian framework for tracking by writing the joint probability of a given target's state and measurement up to time k as follows:

$$\chi^k = \sum_{i \in \{0,1\}} p(\mathbf{x}^k, \mathbf{z}^k, s^i(k)). \quad (3)$$

The target perceivability is thus explicitly accounted for in the probability computation. Using the perceivability concept in the estimation of the tracks likelihood favors perceivable tracks and thus increases the robustness of the tracking procedure to spurious measurements. Indeed, the association of a track with a false detection lowers its probability of perceivability and a dummy target is

expected to have a low probability of perceivability. We further factorize the track joint probability (3) as

$$\begin{aligned} \chi^k &= p(\mathbf{x}^{k-1}, \mathbf{z}^{k-1}) \\ &\cdot \sum_{i \in \{0,1\}} [p(\mathbf{x}(k), \mathbf{z}(k) | s^i(k), \mathbf{x}^{k-1}, \mathbf{z}^{k-1}) \\ &\cdot p(s^i(k) | \mathbf{x}^{k-1}, \mathbf{z}^{k-1})], \\ &= \chi^{k-1} \sum_{i \in \{0,1\}} \xi^i(k) \lambda^i(k | k-1), \end{aligned} \quad (4)$$

which shows that χ^k can be evaluated by updating χ^{k-1} with the current state and detection, hence making it compatible with an iterative computation procedure. In (4), the update factor for each perceivability state is composed of two terms:

1. $\xi^i(k) \triangleq p(\mathbf{x}(k), \mathbf{z}(k) | s^i(k), \mathbf{x}^{k-1}, \mathbf{z}^{k-1})$ is the joint probability of the target state and measurement conditioned by the perceivability state $s^i(k)$ and the past measurements and states. The computation of $\xi^i(k)$ is detailed in Section 3.2.
2. $\lambda^i(k | k-1) \triangleq p(s^i(k) | \mathbf{x}^{k-1}, \mathbf{z}^{k-1})$ is the probability of the perceivability state $s^i(k)$ when only past measurements and states are known. It is thus termed the *predicted probability* of the state $s^i(k)$.

$$\begin{aligned} \lambda^i(k | k-1) &= \sum_{j \in \{0,1\}} p(s^i(k) | s^j(k-1), \mathbf{x}^{k-1}, \mathbf{z}^{k-1}) \\ &\cdot p(s^j(k-1) | \mathbf{x}^{k-1}, \mathbf{z}^{k-1}), \end{aligned} \quad (5)$$

where $p(s^i(k) | s^j(k-1), \mathbf{x}^{k-1}, \mathbf{z}^{k-1}) = \pi_{ji}$ from the Markov property of the perceivability chain and the conditional independence of the perceivability state on the target past states and measurements. The term $\lambda^j(k-1 | k-1) \triangleq p(s^j(k-1) | \mathbf{x}^{k-1}, \mathbf{z}^{k-1})$ is the probability of the perceivability state s^j at time $k-1$ when the measurement and states are known up to this time. It is thus named the *corrected probability* of the state s^j at time $k-1$. Still using Bayes' rule, we express $\lambda^j(k-1 | k-1)$ as

$$\lambda^j(k-1 | k-1) = \frac{\xi^j(k-1) \chi^{k-2}}{\chi^{k-1}} \lambda^j(k-1 | k-2), \quad (6)$$

which shows that the corrected probability of $s^j(k-1)$ can be computed as an update of its predicted probability of perceivability $\lambda^j(k-1 | k-2)$ with the additional data found at time $k-1$. Using these notations, we rewrite the predicted probability of perceivability in (5) as

$$\lambda^i(k | k-1) = \sum_{j \in \{0,1\}} \pi_{ji} \lambda^j(k-1 | k-1). \quad (7)$$

By combining (4) with (7) and (6), we design a sequential process for the estimation of χ^k : It consists of alternating the estimation of χ^r and that of $\{\lambda^i(r | r-1)\}_{i \in \{0,1\}}$ with $r = 1 \dots k$. Algorithm 1 describes the complete estimation scheme that is composed of three iterated main steps: 1) The probability of each perceivability state is first predicted, 2) a new measurement and state are then considered, which lead to the update of the joint

TABLE 1
Computation of $\xi^i(k)$ as a Function of the State $s^i(k)$ and the Type of the Associated Measurement $\mathbf{z}(k)$

Detection type	Perceivable target	Non-perceivable target
Real	$\xi^1(k) = P_D P_G g(\mathbf{z}(k) \mathbf{x}^k, \mathbf{z}^{k-1}) f(\mathbf{x}(k) \mathbf{x}^{k-1}, \mathbf{z}^{k-1})$	$\xi^0(k) = P_{FD}(\mathbf{z}(k), k)$
Virtual	$\xi^1(k) = 1 - P_D P_G$	$\xi^0(k) = 1$

probability of the sequence of states and measurements, and 3) the perceivability probabilities are updated accordingly. Each loop appends a new measurement and a new state until the whole track is processed, hence making the probability estimation compatible with tracking procedures, such as the MHT, which link detections sequentially over time. The recursive property of the estimation algorithm is also very desirable since all the previous data do not have to be reprocessed whenever a new measurement is received.

Algorithm 1. Sequential computation of states and measurements joint probability including perceivability states estimation.

Require: $z^K, \chi^1, \{\lambda^i(1|1)\}_{i \in \{0,1\}}$
for $k = 2$ **to** K **do**
 for $i = 0$ **to** 1 **do**
 $\lambda^i(k|k-1) \leftarrow \sum_{j \in \{0,1\}} \pi_{ji} \lambda^j(k-1|k-1)$
 {Probability of perceivability prediction}
 end for
 for $i = 0$ **to** 1 **do**
 Compute $\xi^i(k)$ (see Section 3.2)
 end for
 $\chi^k \leftarrow \chi^{k-1} \sum_{i \in \{0,1\}} \xi^i(k) \lambda^i(k|k-1)$
 {Update of the measurements and states joint probability}
 for $i = 0$ **to** 1 **do**
 $\lambda^i(k|k) \leftarrow \frac{\xi^i(k) \chi^{k-1}}{\chi^k} \lambda^i(k|k-1)$
 {Update of the perceivability probability}
 end for
end for
return χ^K

3.2 Measurement and State Joint Probability

The probability of the current measurement and state conditioned by the perceivability state can take four expressions, depending on the target perceivability state and on whether the detection is real or virtual. In the perceivable target case, we write

$$\xi^1(k) = p(\mathbf{z}(k)|s^1(k), \mathbf{x}^k, \mathbf{z}^{k-1}) p(\mathbf{x}(k)|s^1(k), \mathbf{x}^{k-1}, \mathbf{z}^{k-1}),$$

where $p(\mathbf{z}(k)|s^1(k), \mathbf{x}^k, \mathbf{z}^{k-1})$ is the measurement probability and depends on the type of the detection $\mathbf{z}(k)$. Indeed, if $\mathbf{z}(k)$ corresponds to a real detection, it is then assumed to originate from the target intensity, which implies that the target has been detected and that the detection has fallen into the track search gate. By contrast, if the detection is virtual, we consider either that the target has not been detected or that it has left the search gate. Hence, if the detection is real, the conditional detection probability is computed as: $p(\mathbf{z}(k)|s^1(k), \mathbf{x}^k, \mathbf{z}^{k-1}) = P_D P_G g(\mathbf{z}(k)|\mathbf{x}^k, \mathbf{z}^{k-1})$, where P_D is

the probability of detecting the target when it is perceivable, P_G is the probability that the corresponding detection falls into the search gate. We denote by $f(\mathbf{x}(k)|\mathbf{x}^{k-1}, \mathbf{z}^{k-1}) \triangleq p(\mathbf{x}(k)|s^1(k), \mathbf{x}^{k-1}, \mathbf{z}^{k-1})$ the pdf of state evolution when the target is perceivable and $g(\mathbf{z}(k)|\mathbf{x}^k, \mathbf{z}^{k-1})$ the conditional pdf of the measurement at time k , yielding

$$\xi^1(k) = \begin{cases} P_D P_G g(\mathbf{z}(k)|\mathbf{x}^k, \mathbf{z}^{k-1}) f(\mathbf{x}(k)|\mathbf{x}^{k-1}, \mathbf{z}^{k-1}) & \text{if } \mathbf{z}(k) \text{ is real,} \\ 1 - P_D P_G, & \text{otherwise,} \end{cases}$$

for the joint state and measurement probability. Standard estimation techniques for Bayesian tracking can be used (e.g., KF, SMC) in this case to compute $g(\mathbf{z}(k)|\mathbf{x}^k, \mathbf{z}^{k-1}) f(\mathbf{x}(k)|\mathbf{x}^{k-1}, \mathbf{z}^{k-1})$ since the target is implicitly assumed perceivable by the standard models. Appropriate transition and measurement models are required by estimation techniques to compute $\xi^1(k)$. In Section 3.4, such a model is given for biological scenarios.

A white noise model is generally assumed for digital biological microscopy. Hence, approximating the probability of the set of false detections in (2) as the product of independent events $P\{Z_0(t)\} = \prod_{\mathbf{z} \in Z_0(t)} P_{FD}(\mathbf{z}, t)$ is suited to most particle detection procedures. The probability that a detection is wrongly decided at position $P_{FD}(\mathbf{z}, t)$ depends on parameters such as the acquisition system, the biological sample fluorescence, the detection procedure settings, and may thus vary in space and time. On the other hand, we note that when a target is not perceivable with probability one ($\lambda^1(k|k) = 0$), the associated detection should be considered as a false detection, for a nonperceivable target does not generate any measurement by definition. To seamlessly integrate the perceivability model (4) in (2), we thus set $\xi^0(k) = P_{FD}(\mathbf{z}(k), k)$ if the detection $\mathbf{z}(k)$ is real. By contrast, we take $\xi^0(k) = 1$ to discard the effect on (2) of the association of a virtual measurement with a nonexistent target.

The different values for the joint probability $\xi^i(k)$ as a function of the target perceivability state and the detection type are summarized in Table 1.

3.3 Perceivability Chain Settings

3.3.1 Initial Probability of Perceivability

To initiate the recursive estimation procedure described in Algorithm 1 for a target appearing at frame k , it is necessary to provide $\{\lambda^i(k|k)\}_{i \in \{0,1\}}$. We base our design on the model of false detections and on the prior target appearance events, which is the more commonly available prior information. In a given frame k , a new target is assumed to appear at the position \mathbf{z} with probability $P_{NT}(\mathbf{z}, k)$, depending on position and time. In biology, such flexibility is important as the biological environment and various

biological and chemical mechanisms play a crucial role in the appearance of new particles in the field of view. As an example, the proximity to the Golgi apparatus, the protein packaging site for most eukaryotic cells, is a determinant factor of protein appearance events. Also, the rate of successful particle detection may decrease with time due to fluorescence extinction processes.

When the probability of generating a false detection $P_{FD}(\mathbf{z}, k)$ is very small compared to the probability that a new target appears, a detection at location \mathbf{z} that is nonassigned to a track must almost surely originate from a perceivable target. Hence, intuitively, $\lambda^1(k|k)$ should be close to one in this case. Conversely, if the probability of target appearance is much smaller than the probability of false detection $P_{FD}(\mathbf{z}, k) \gg P_{NT}(\mathbf{z}, k)$, such a detection must originate from noise with a probability close to one.

These two requirements lead us to set the initial values of the probability of perceivability as

$$\lambda^0(k|k) = \frac{P_{FD}(\mathbf{z}, k)}{P_{FD}(\mathbf{z}, k) + P_{NT}(\mathbf{z}, k)},$$

$$\lambda^1(k|k) = \frac{P_{NT}(\mathbf{z}, k)}{P_{FD}(\mathbf{z}, k) + P_{NT}(\mathbf{z}, k)},$$

and we compute the first state and measurement probability as: $\chi^k = \sum_{i \in \{0,1\}} \xi^i(k) \lambda^i(k|k)$. The estimation loop defined in Algorithm 1 is then applied to the remaining measurements and states.

3.3.2 State Transition Parameters

Discriminating a new particle from a nonperceivable target becoming perceivable again after a long time interval is barely feasible for fluorescent objects that are alike. We therefore assume that nonperceivable targets cannot become perceivable again, and we set $\pi_{01} = 0$ and $\pi_{00} = 1$. Linking partial tracks originating from the same particle can still be done in a postprocessing step, as done in [9] and [20]. The other state transition parameters are not known a priori for most biological applications, but fixing them beforehand is required.

The perceivability state model can be seen as a Bernoulli process: Between subsequent frames, one perceivable particle remains so with probability π_{11} and becomes nonperceivable with probability π_{10} . For a perceivable particle, the mean delay before switching to the state s^0 is therefore π_{10}^{-1} . We propose to fix the transition parameter π_{10} in such a way that the expected track length given by the perceivability model matches the prior expected track length: l_m . To do so, we set $\pi_{10} = l_m^{-1}$ and we generally fix $\pi_{01} = 0.05$. This value is consistent with many nanometer scale bioimaging systems. For instance, in [29], trajectories with an average length of 16.9 frames were measured. Using an approximation of π_{01} provides the desired behavior of the perceivability estimation procedure because it also relies on the observation model and not on the prior transition model solely. The remaining transition parameter, π_{11} , is simply fixed as: $1 - \pi_{10}$. An exhaustive study of the target perceivability probability as a function of the state transition parameters is provided in [13] and shows the consistency of the estimation process behavior with the parameters interpretation.

3.4 Particle Motion Model

We propose to adopt the conveyor-belt model discussed in Appendix A, which can be found in the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2013.97>, for the motion of particles in biological environments. We define a jump Markov state transition model between two primary motion models: random walk and directed transport, with state switch probabilities P_{on} and P_{off} . Fixing P_{on} and P_{off} and changing parameters of each primary model allows one to account for a variety of biological cases.

For the random walk model, we define the particle state as the vector of Cartesian coordinates: $\mathbf{x}(k) = (x(k), y(k), z(k))$ such that the conditional pdf $p(\mathbf{x}(k+1)|\mathbf{x}(k))$ is Gaussian with mean $\mathbf{x}(k)$. As a consequence, the state evolution equation (1) is linear with Gaussian process noise and the KF can be used for state estimation. The state transition and noise covariance matrices are defined such that the conditional pdf of the state $\mathbf{x}(k+1)$ corresponds to the distribution prescribed by the Brownian motion model: $\mathbf{F}_{\text{Brownian}} = \mathbf{I}_3$ and $\mathbf{Q}_{\text{Brownian}} = \sigma_b^2 \mathbf{I}_3$, with \mathbf{I}_3 the 3×3 identity matrix and $\sigma_b = 2D\delta t$ (see Appendix A, available in the online supplemental material).

We assume hereafter that the force field \mathbf{G} in the transport model is locally constant and varies slowly in space and time. In the absence of further information, we assume a white and zero-mean Gaussian model of force field fluctuations between two subsequent positions of a target. We thus use the following target state definition: $\mathbf{x}(k) = (x(k), y(k), z(k), d_x(k), d_y(k), d_z(k))$, where $d_x(k)$ is the expected displacement due to the force field in the horizontal direction ($\frac{\mathbf{a}}{m\gamma} \delta$ projected onto the x -axis, see Appendix A, available in the online supplemental material). As a result, we obtain a linear state evolution equation with

$$\mathbf{F}_{\text{belt}} = \begin{pmatrix} \mathbf{I}_3 & \mathbf{I}_3 \\ \emptyset & \mathbf{I}_3 \end{pmatrix}, \quad \mathbf{Q}_{\text{belt}} = \begin{pmatrix} \sigma_{\text{teth}}^2 \mathbf{I}_3 & \emptyset \\ \emptyset & \sigma_d^2 \mathbf{I}_3 \end{pmatrix}$$

that allows applying the KF. By including $\mathbf{d} = (d_x(k), d_y(k), d_z(k))$ in the target state (with $\sigma_d^2 \mathbf{I}_3$ its covariance matrix), the field fluctuations are evaluated by the estimation procedure based on the measured particle positions. The two KFs for the primary models of transport are combined in an IMM filter with P_{on} and P_{off} for the transition probabilities. Hence, efficient state estimation is achieved without compromising on its accuracy, which is crucial in our case because thousands of putative tracks may be built simultaneously by the MHT algorithms in large-scale applications.

4 LARGE-SCALE MULTIFRAME TRACKING

In this section, we present an MHT tracking algorithm that is suited to large biological data with dense clutter. Our contribution mainly consists of the seamless integration of the comprehensive probabilistic particle tracking framework introduced in Section 3 so as to obtain a self-contained and robust MHT algorithm. To solve the MHT problem exactly in reasonable time, we adapt a linear programming algorithm. As a result, our technique permits the representation of the

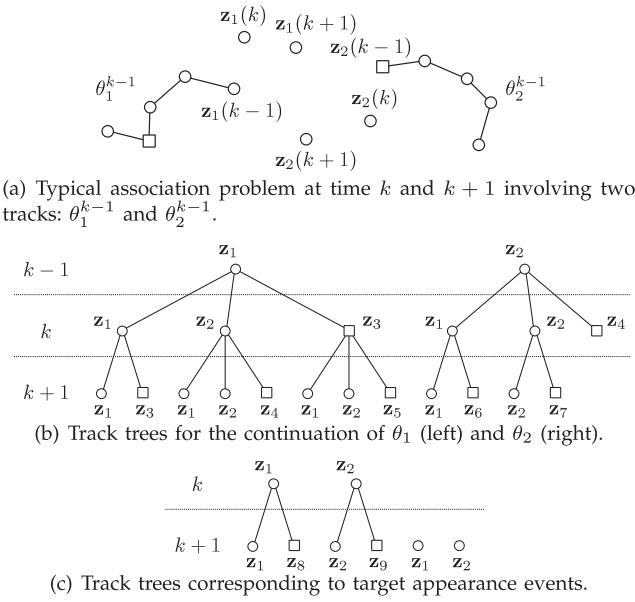


Fig. 2. Construction of an ensemble of track trees for an example association problem. Track continuation and new target appearance yield separate trees.

full solution space over multiple frames, which is beneficial under extremely poor imaging conditions.

4.1 Bayesian Multiframe Tracking with Trees of Associations

For an image sequence of length K , we aim at building the set of tracks Θ^{*K} with maximum likelihood: $\Theta^{*K} = \arg \max_{\Theta^K} \mathcal{L}(\Theta^K)$, where $\mathcal{L}(\Theta^K)$ has been defined in Section 3 for biological sciences, and Ω^K is the space of feasible tracks up to frame K . Since optimizing $\mathcal{L}(\Theta^K)$ over the whole solution space Ω^K is not a feasible option when K gets large, we adopt a sequential approach for tracking. Let us assume optimal tracks Θ^{*k-1} are available up to frame $k-1$. Then, the MHT approach is to extend the current tracking solution at frame k by solving an optimization problem that accounts for d future frames:

$$\Theta^{*k+d} = \arg \max_{\Theta^{k+d} \in \Omega_{\Theta^{*k-1}}^{k+d}} \mathcal{L}(\Theta^{k+d}). \quad (8)$$

As shown in Fig. 2, the potential extensions of an existing tracks set can be represented with trees of associations with measurements in the set of frames from time k to $k+d$. More formally, each track $\theta_j^{k-1} \in \Theta^{*k-1}$ can be associated either with a virtual measurement modeling its temporary disappearance or with a detection in $Z(k)$ that falls into its search gate. The potential associations of θ_j^{k-1} at frame k form a set $\Gamma_j^k = \{\theta_{j,z_i}^k\}_{i=1..p}$ of tracks where θ_{j,z_i}^k stands for the potential track built by associating the track θ_j^{k-1} with the measurement $z_i(k)$, which is either virtual or real. We recursively apply the association process for every set Γ_j^t with $t = k \dots k+d$, yielding the formation of potential tracks $\Gamma_j^{k:k+d} = \{\Gamma_j^t\}_{t=k..k+d}$ as the construction of trees of feasible associations. Nodes in $\Gamma_j^{k:k+d}$ without any link to a node at the next level of the tree are coined *leaves* and constitute the possible track segments for continuing the track θ_j^{k-1} . It is worth noting that a single measurement may yield multiple nodes in the track trees. We, however,

impose that different tracks do not share detections, hence resulting in numerous incompatibilities between leaves during the construction of tracking hypotheses.

4.2 Automatic Track Creation and Termination

When a target disappears from the surveillance volume, its probability of perceivability generally decreases toward 0 since no sequence of likely associations can be found over the time window. As pointed out in Section 3.2, when $\lambda^1(k|k-1)$ is small, the track to detection association either participates in the tracking hypothesis likelihood as a false detection if the measurement is real ($\chi^k = \chi^{k-1} P_{FD}(\mathbf{z}(k), k)$), or does not influence χ^k if the measurement is virtual ($\chi^k = \chi^{k-1}$). Therefore, the tracking model automatically acts as if the target does not exist anymore. Continuing to extend the track corresponding to a nontarget track thus amounts to identifying some false detections, which could be done without maintaining the track. To save memory and limit the number of calculations, we thus fix a lower bound λ_t on the value of $\lambda^1(k|k-1)$ such that a track segment is stopped and labeled as *terminated* as soon as its predicted probability of perceivability falls beyond this value. In the following, we fix $\lambda_t = 10^{-5}$.

Terminating potential tracks results in leaves in $\Gamma_j^{k:k+d}$ that can end before the frame $k+d$. This is illustrated in Fig. 2b with the leaf formed by the association of the track θ_2 with the virtual detection $\mathbf{z}_4(k)$. When a terminated potential track is selected by the optimization procedure constructing $\Theta^{*k:k+d}$, the corresponding existing track is terminated and is not processed further. By doing so, the target disappearance detection is embedded within the global optimization procedure: Tracks are ended when their continuations do not provide any further improvement of the objective score function $\mathcal{L}(\Theta^{k:k+d})$. This approach contrasts with standard techniques that handle track termination thanks to procedures that are generally independent of the association stage.

Similarly to disappearance, target appearance detection is fully embedded within the optimization process. As shown in Fig. 2c, for each detection in $Z^{k:k+d}$ we create a tree of associations with the given detection as its root, thus allowing for new targets to appear. Systematically creating a new tree for each detection constitutes a fully automatic track creation process with no ad hoc decision procedure: The appearance of new targets events is detected by the association selection procedure. Indeed, selecting one leaf at time k from the tree originating from a detection $\mathbf{z}_i(k)$ results in the creation of a new track that aims at following the putative target from which $\mathbf{z}_i(k)$ originates. On the contrary, if no leaf is selected in this tree, the track creation is then rejected and $\mathbf{z}_i(k)$ is assumed to originate either from an already existing target or from clutter.

4.3 Optimization with Constrained Linear Programming

When track segments up to time $k-1$ are fixed to Θ^{*k-1} , the maximization problem (8) boils down to selecting the subset of track segments $\Theta^{*k:k+d} \subset \Gamma^{k:k+d}$, with $\Gamma^{k:k+d}$ the set of extended and newly created putative tracks, that maximize the partial cost function:

$$\tilde{\mathcal{L}}(\Theta^{k:k+d}) = \prod_{t=k}^{k+d} p(Z_0(t)) \cdot \prod_{\theta_j^{k:k+d} \in \Theta^{k:k+d}} \chi_j^{k:k+d}, \quad (9)$$

where the set of nonassociated detections $Z_0(t)$ at each frame t depends on the set of selected tracks. The partial score $\chi_j^{k:k+d}$ is computed as $\chi_j^{k:k+d} = \chi_j^{k+d} / \chi_j^{k-1}$ as track segments prior to time k are fixed.

The maximization of the form (9) is constrained by the compatibility between selected track segments and the requirement to continue any existing and nonterminated track, which makes it combinatorial. Following the approach in [27], we recast the problem as an optimization problem for which efficient algorithms can be used. We rewrite the maximization of (9) as the maximization of a linear program $\mathbf{b}^T \mathbf{s}$ where \mathbf{b} is a binary vector and \mathbf{s} is a score vector of the same size. \mathbf{b} is such that $\theta_j^{k:k+d}$ is selected in $\Gamma^{k:k+d}$ if $\mathbf{b}(j) = 1$, while $\mathbf{s}(j) = \log(\chi_j^{k:k+d})$. We introduce the concept of *dummy track* so as to set $\mathbf{b}^T \mathbf{s}$ equal to $\log(\tilde{\mathcal{L}}(\Theta^{k:k+d}))$ for a given set of tracks $\Theta^{k:k+d}$. A dummy track contains exactly one detection $\mathbf{z}(t)$ with $k \leq t \leq k+d$. Selecting this track means that the detection is not associated with any track segment in $\Theta^{k:k+d}$. In accordance with the probabilistic model introduced in Section 3, we extend \mathbf{s} with values for dummy tracks that depend on the type of detection: $\log(P_{FD}(\mathbf{z}(t), t))$ if $\mathbf{z}(t)$ is real, and 0 if $\mathbf{z}(t)$ is virtual (in order not to penalize the rejection of virtual measurements).

The maximization of $\mathbf{b}^T \mathbf{s}$ needs to be constrained to make the optimal solution feasible and unique (e.g., \mathbf{b} empty is a trivial optimal solution of the nonconstrained problem as all elements of \mathbf{s} are negative). One faces two types of constraints: 1) Each measurement in $Z^{k:k+d}$ has to be selected by at most one track, and 2) each nonterminated track in Θ^{*k-1} needs to be extended with a track segment. For each tracking problem, we build an assignment matrix A with binary elements that establish the correspondence between track segments and measurements. To include constraints on the extension of existing tracks, we also account for measurements at time $k-1$ associated with the tracks in Θ^{*k-1} . For instance, for the tracking problem in Fig. 2, we build a 31×15 association matrix:

$$A = \begin{array}{c} \text{track} \\ \text{extension} \\ \text{creation} \\ \text{dummy} \end{array} \begin{array}{c} \overbrace{\begin{array}{cccccccc} \mathbf{z}_j(k-1) & \mathbf{z}_j(k) & \mathbf{z}_j(k+1) \end{array}} \\ \left(\begin{array}{cccccccccccccccc} 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ & & & & & & \vdots & & & & & & & & \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ & & & & & & \vdots & & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ & & & & & & \vdots & & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right) \end{array}$$

Each measurement in $Z^{k:k+d}$ corresponds to a column in A , while rows represent track segments in $\Gamma^{k:k+d}$ and dummy tracks. Thanks to the use of dummy tracks for measurements, except at time $k-1$, we are able to write our two

classes of constraints as a homogeneous set of M (the number of columns of A) linear equality constraints:

$$\mathcal{S} = \{\mathbf{b}^T A_j = 1 \text{ for } j = 1 \dots M\}.$$

We also create the set \mathcal{B} of N (number of track segments) binary constraints for the vector \mathbf{b} :

$$\mathcal{B} = \{\mathbf{b}(j) = 1 \text{ or } 0 \text{ for } j = 1 \dots N\}.$$

As a result, we rewrite the tracking problem as a constrained linear program P :

$$(P) : \mathbf{b}^* = \arg \max_{\mathbf{b} \in \mathbb{R}^N} \mathbf{b}^T \mathbf{s} \text{ such that } \mathcal{B} \text{ and } \mathcal{S}.$$

As constraints \mathcal{B} are nonconvex, relaxing P is not a valid option to retrieve directly the optimal solution \mathbf{b}^* . Instead, we use a specialized *branch-and-bound* technique for integer linear programs [12]. The algorithm consists of using the *simplex* algorithm for solving a series of linear convex programs: For each problem, a subset of elements of \mathbf{b} is fixed to either 0 or 1, and the other ones are relaxed to the real interval $[0, 1]$ (which correspond to convex constraints). These relaxed problems provide upper bounds on the solution. The procedure is stopped when the best relaxed solution for an appropriate set of subproblems coincides with a binary solution. The algorithm is guaranteed to converge to the optimal solution of P when using a suitable series of subproblems.


In practice, block-diagonal assignment matrices are frequently observed since large tracking problems may be composed of spatially disjoint tracking subproblems. We thus first cluster track trees such that groups of trees do not share any detection. Each subproblem is then solved independently in a parallel manner. For it, we use the *lp_solve* open-source library that provides efficient algorithms [4]. As a result, we show in Section 5 that the computational cost is kept low although the tracking problems we have processed are large.

4.4 Tree Maintenance


After Θ^{*k+d} is selected, we update track trees to preserve their compatibility with the associations in $\Theta^*(k)$: The branches directly linked to root nodes at time $k-1$ that do not correspond to $\Theta^*(k)$ are pruned. Detecting target appearance/disappearance events is also achieved at this stage depending on whether terminated track segments and new trees with a root node at time k are selected. These principles are illustrated in Fig. 3 for the example problem of Fig. 2. Then, the best tracking hypothesis Θ^{*k-1} is updated with the tracks from the selected solution $\Theta^*(k)$, and the root nodes of trees are shifted to the time k since no more ambiguities and incompatibilities remain at this time. The whole tracking procedure is then repeated for the next frame: Potential tracks are formed, trees are clustered, the best tracking hypothesis is selected, and the trees are updated.

4.5 Scalability

We have conducted an experimental study of the practical scalability of the proposed MHT algorithm. First, the stability over time of the computational cost, a crucial



(a) Branches forming Θ^{*k+1} are bold face with nodes indicated by a star.



(b) Maintained trees after association selection. A new track with root node $z_1(k)$ is decided. Level k is now fixed.

Fig. 3. Example of trees maintenance operations for the tracking problem in Fig. 2.

feature for long time sequence processing, was investigated. To do so, we have measured the number of clusters, trees, and hypotheses over the 482 frames-length sequence of microscopy images presented in Section 5.2. The results, summarized in Fig. 4, show that the tracking problem reaches a steady size in terms of tree-based representation after only few frames. This translates into a steady computational cost for each frame of the sequence. In a second experiment, we have taken regions of interest of increasing size in the same dataset and measured the size of the tree-based representation of the tracking problem in each case. Results in Fig. 4 show that the number of leafs grows in a linear fashion only with respect to the size of the region of interest. This is thanks to the structure of the tracking problem that can be divided in clusters. Moreover, experiments show that extra targets are handled in new clusters that interact with the other ones only in a limited extent. The computational cost is therefore increased just by the extra amount of processing that is required for the new clusters. Together, these experiments demonstrate the practicality of the proposed tracking technique for large-scale data.

5 VALIDATION AND RESULTS

We present next a series of two experiments performed on both synthetic and real datasets with a number of state-of-the-art MPT techniques and with the proposed MHT. We have already mentioned that handling poor imaging conditions is a major challenge when tracking particles in biological environments. We have thus designed a benchmark for assessing the capability of algorithms to track particles in low SNR fluorescence images.

For the sake of reproducibility of the following results, our MHT algorithm has been implemented in Java as an open-source plug-in for the Icy software [16] and is available in the official online repository.




Fig. 4. Scalability of the MHT. Number of clusters (bottom line), trees (middle line), and leaves (top line) maintained by the MHT algorithm for a real sequence of microscopy images. Measured values are represented as a function of frame number for a region of size 400×400 pixels (left). We give the mean of the same measures for a region of increasing size (right).

5.1 Low-SNR Synthetic Images Processing


We have first investigated tracking performance for synthetic, yet realistic, 2D+T datasets, for which controlling conditions and assessing tracking performance is facilitated.

5.1.1 Benchmark Setup


Synthetic data have been generated using the procedure described in Appendix B, available in the online supplementary material, with the following settings:

- sequences of 50 frames, each one being a 256×256 pixels image with a background of intensity 50 that is corrupted by a white Gaussian noise of variance $\sigma^2 = 25$,
- 20 particles exhibiting a conveyor-belt motion in the first frame,
- 0.1 mean number of particle appearance events, and 0.01 probability of disappearance for each particle between two subsequent frames.

We show in Fig. 5a two examples of generated trajectories that illustrate the variability of the particle movements and the possibility for targets to appear and disappear over time.



(a) Time projection of trajectories (solid lines) generated for the benchmark, and of measurements provided by the wavelets particle detector. Left: PSNR = 4.77 dB, right: PSNR = 2.04 dB.



(b) A series of cropped images around some particles with different intensity amplitudes A .

Fig. 5. Synthetic data for tracking performance assessment. Sample images show that for PSNR < 4 dB discriminating particles from noise is very ambiguous. In that case, we show (top right) that most detected positions correspond to clutter.

TABLE 2
Detection Performances for the Benchmark Data

PSNR (dB)		7.78	6.02	4.77	3.01	2.04
Reference detections		9555	8404	8431	7884	8625
Gaussian fitting detector [20]	TPR	1.0	1.0	1.0	1.0	0.99
	RR	0.96	0.96	0.80	0.76	0.74
Wavelet-based detector [30]	TPR	0.77	0.72	0.71	0.37	0.37
	RR	0.99	0.99	0.99	0.98	0.94

To assess tracking performance in images of various quality, we have used five different values for the amplitude of the intensity profile of particles: $A = 30, 20, 15, 10, 8$, which correspond to the peak SNR (PSNR (in dB) = $10 \log_{10} \frac{A}{\sigma}$) values of 7.78, 6.02, 4.77, 3.01, and 2.04 dB, respectively. As shown in Fig. 5b, the image quality varies from bright particles to barely visible targets. We have generated 10 image sequences for each condition, for a total of 1,528 reference tracks.

We have compared the results obtained with three state-of-the-art techniques for MPT in microscope images: *Genovesio'06* [18], *Jaqaman'08* [20], *Sbalzarini'05* [38], and the proposed MHT. Two different detection techniques have been used to investigate the influence of the detection stage: The fitting of Gaussian mixture models algorithm used in [20] with conservative settings and a wavelet-based filtering method [30] with sensitive settings. The parameters for the Gaussian fitting and for the *Jaqaman'08* algorithm have been fixed to optimal values by the first author of the paper [20].

The detection performances for both methods are given in Table 2. Tracking performances are measured using the following complementary measures that are discussed in Appendix C, available in the online supplemental material: the OSPA distance for tracks [35], which is a metric accounting for both localization and track labeling errors,

and the tracks recovery rate (RR) and true-positive rate (TPR) that are sensitivity and robustness indicators, respectively. The Jaccard similarity coefficient (JSC) aggregates both the RR and TPR in a single index. The JSC, RR, and TPR coefficients take in the interval $[0, 1]$, the optimum being 1, while the OSPA is averaged over the sequence length and takes value 0 for a perfect match.

5.1.2 Results

We provide the complete set of results in Table 3. It is instructive to take a close look at results first when the image quality is the highest (PSNR ≥ 6.02 dB). Under these conditions, the Gaussian fitting detector provides accurate measurements with almost no false detections, which leads all tested algorithms, except *Sbalzarini'05*, to provide very similar results characterized by a very high robustness. No significant differences are present, which is not surprising as target density and appearance/disappearance events are not major issues for this benchmark. Moreover, the strategy of *Jaqaman'08* and *Genovesio'06* to cope with missed detections (i.e., linking partial segments) is effective because no false detections corrupt the measurements. On the other hand, *Sbalzarini'05* is penalized by its distance-based linking method that does not account for specific motion models, by contrast with other techniques that accommodate more complex dynamics. Under these conditions, using more sensitive detection settings (wavelet-based detection) is shown to decrease the tracking performance for all tracking algorithms but the MHT. Indeed, the TPR is decreased by the presence of false detections that these techniques may not handle properly. By contrast, track recovery is significantly improved for the MHT (RR ≥ 0.92) when using the ability of local maxima separation of the wavelet-based detector that helps resolving closely spaced targets that may yield a single measurement when using the detector in [20]. At the same time, the TPR of the MHT is not degraded by

TABLE 3
Tracking Performances (Median Values over 10 Trials) for Three MPT Techniques and the MHT

PSNR (dB)		Gaussian shape fitting detection [20]				Wavelet-based detection [30]			
		<i>Sbalzarini'05</i>	<i>Genovesio'06</i>	<i>Jaqaman'08</i>	MHT	<i>Sbalzarini'05</i>	<i>Genovesio'06</i>	<i>Jaqaman'08</i>	MHT
7.78	JSC	<u>0.75</u>	<u>0.81</u>	<u>0.81</u>	<u>0.80</u>	<u>0.31</u>	<u>0.73</u>	<u>0.48</u>	<u>0.90*</u>
	TPR	0.90	1	1	1	0.34	0.88	0.51	0.98
	RR	0.80	0.82	0.81	0.81	0.79	0.83	0.87	0.92
	OSPA	<u>6.11</u>	<u>2.67</u>	<u>2.59</u>	<u>2.43^{ns}</u>	<u>11.41</u>	<u>4.38</u>	<u>8.47</u>	<u>3.19</u>
6.02	JSC	<u>0.73</u>	<u>0.79</u>	<u>0.85</u>	<u>0.89</u>	<u>0.24</u>	<u>0.69</u>	<u>0.34</u>	<u>0.92*</u>
	TPR	0.92	0.94	0.96	1	0.27	0.88	0.36	0.98
	RR	0.78	0.84	0.85	0.89	0.74	0.75	0.82	0.94
	OSPA	<u>6.42</u>	<u>2.83</u>	<u>2.44</u>	<u>2.25^{ns}</u>	<u>13.23</u>	<u>5.11</u>	<u>10.85</u>	<u>2.42</u>
4.77	JSC	<u>0.61</u>	<u>0.62</u>	<u>0.64</u>	<u>0.68</u>	<u>0.22</u>	<u>0.56</u>	<u>0.30</u>	<u>0.89***</u>
	TPR	0.89	0.87	0.89	0.92	0.25	0.81	0.33	0.97
	RR	0.58	0.57	0.68	0.61	0.72	0.67	0.77	0.92
	OSPA	<u>7.33</u>	<u>4.73</u>	<u>5.10</u>	<u>4.53</u>	<u>15.49</u>	<u>7.36</u>	<u>11.64</u>	<u>2.92***</u>
3.01	JSC	<u>0.53</u>	<u>0.54</u>	<u>0.60</u>	<u>0.57</u>	<u>0.20</u>	<u>0.37</u>	<u>0.24</u>	<u>0.80**</u>
	TPR	0.89	0.89	0.89	0.90	0.23	0.67	0.27	0.95
	RR	0.58	0.57	0.66	0.61	0.66	0.43	0.67	0.84
	OSPA	<u>8.23</u>	<u>6.54</u>	<u>5.80</u>	<u>5.34</u>	<u>17.36</u>	<u>11.01</u>	<u>14.22</u>	<u>4.46**</u>
2.04	JSC	<u>0.42</u>	<u>0.44</u>	<u>0.57</u>	<u>0.49</u>	<u>0.05</u>	<u>0.07</u>	<u>0.04</u>	<u>0.73**</u>
	TPR	0.81	0.78	0.81	0.86	0.05	0.08	0.04	0.93
	RR	0.46	0.51	0.64	0.55	0.38	0.26	0.39	0.76
	OSPA	<u>9.01</u>	<u>7.30</u>	<u>6.07</u>	<u>6.65</u>	<u>21.86</u>	<u>16.01</u>	<u>22.10</u>	<u>4.86*</u>

For each SNR condition, irrespective of the detection methods, the best result is underlined. Statistical hypothesis testing with Kolmogorov-Smirnov test has been performed between the best (underlined) and second-best techniques: *** for p -value < 0.001 , ** for p -value < 0.01 , * for p -value < 0.05 , and ^{ns} for p -value ≥ 0.05 .

the spurious measurements as it allows automatically detecting them.

Under low SNR conditions (PSNR ≤ 4.77 dB), we observe a significant decrease of performance for all MPT methods whenever the Gaussian fitting technique is used for detecting particles. The RR is lowered because the conservative settings of the detector do not allow identifying a sufficient rate of target positions (≤ 80 percent of detected locations). As a result, in the limit case PSNR = 2.04 dB, no tested algorithm complies with the quality requirements of most biological applications. On the other hand, the wavelet-based detector provides a higher number of correct particle positions in low SNR conditions (≥ 94 percent of detected locations), although the measurements are corrupted by a high number of false-positive (FP) detections (63 percent of spurious measurements). Results show that no prior method is able to recover correct measurements from these corrupted datasets, hence resulting in poor tracking performances. These results contrast with those of the MHT. Indeed, the TPR for the MHT is consistently very high (≥ 0.93) despite the very low SNR conditions. It demonstrates that the proposed MHT identifies true and false detections within the dataset while tracking the targets. This unique ability is possible thanks to the accurate false detections model that the MHT seamlessly integrates into the association process. The robustness of the tracking decisions is further enhanced by the ability to account for a number of future frames. As a result, both the RR and TPR performances are kept very high, and the MHT is shown to provide high-quality results, regardless of the imaging conditions. The experiments therefore highlight that it is worth combining a very sensitive detection procedure with the MHT in the case when the detection rate is low and postprocessing methods to link partial track segments are not effective. This stems from the robustness of the MHT that allows recovering true tracks from the many cluttered hypotheses during the optimization of the MHT score.

5.2 Golgi Units Tracking in Microscope Images

We have processed an image sequence from a challenging biological assay to assess the performance of the different MPT methods under real conditions. The aim is to track several thousand Golgi units in Chinese hamster ovary cells over 481 frames (504×405 pixels each) that have been acquired thanks to a spinning-disk confocal microscope. Fluorescent Golgi units appear as bright spots in an uneven and noisy background. We measured PSNR values as low as 2.5 dB for lowest intensity particles, which makes handling false and missed detections a major issue. An expert has manually identified Golgi unit trajectories in a 73×64 pixels crop during 128 time steps, resulting in 124 individually labeled particles. We have then applied the proposed MHT, *Genovesio'06*, *Jaqaman'08*, and *Sbalzarini'05* algorithms with the same settings as above and the most appropriate detection method. For the MHT, four settings were investigated: $d = 2, 3, 4$, and 5. We have also included results for *Chetverikov'99* [14], a multiframe tracking algorithm (limited to $d = 2$) that uses a smooth motion model typical of surveillance applications, to investigate the direct usability of algorithms developed in other contexts than imaging of live cells.

TABLE 4
Performance Assessment of Golgi Units Tracking

	TP	FP	JSC	OSPA	time
Gaussian mixture-model fitting [20]					
<i>Sbalzarini'05</i> [38]	62	60	0.34	13.31	0.5 s
<i>Genovesio'06</i> [18]	32	45	0.19	15.15	0.4 s
<i>Jaqaman'08</i> [20]	57	16	0.41	14.10	6.7 s
<i>Chetverikov'99</i> [14]	31	19	0.22	14.11	0.5 s
Wavelet-based detection [30]					
MHT $d = 2$	49	11	0.36	8.13	1.9 s
MHT $d = 3$	96	16	0.69	3.67	4.9 s
MHT $d = 4$	105	5	0.81	3.61	7.2 s
MHT $d = 5$	106	5	0.82	3.52	11.8 s

We summarize tracking results obtained by the investigated methods in terms of true-positive (TP) and FP tracks, JSC, and mean OSPA in Table 4. We also give the computation time obtained with a Mac Pro Quad 2.66 GHz as a rough indication of the methods computational cost only, as the algorithms differ in language and optimization of the implementation. The results given in Table 4 are consistent with the tracking benchmark on synthetic data: The proposed MHT is the only method we have tested that is able to accurately process this cluttered dataset. The other methods are plagued by the low quality of the measurements. As a result, only a few correct tracks are built, while spurious tracks are numerous. The decrease of performance as compared to the synthetic scenario may be explained by the more compactly packed particles and the higher number of target appearance and disappearance events that make postprocessing techniques to link partial tracks less effectively for the real data. By contrast, the MHT is proven effective as soon as $d > 2$: The vast majority of tracks are correctly recovered, with only a very few spurious ones. Interestingly, increasing the value of d from 4 to 5 improves the tracks quality only marginally (0.09 OSPA improvement), which is in accordance with the short range of autocorrelation of diffusive motion types of biological particle. For those cases, the MHT can be used with a short depth ($d \leq 5$) without any significant loss of performance while reducing the computational cost. As a result, the computation time values we have measured for the MHT are in the same range as for other techniques when using such settings.

The results of *Chetverikov'99* are plagued by numerous cluttered tracks (TPR = 0.62), which highlights the lack of accurate model for spurious and missing detections. Moreover, the low rate of recovered tracks (RR = 0.25) combined with the extra series of experiments with the high SNR synthetic data of Section 5.1 (data not shown) indicate that its smoothness-promoting motion model hampers the recovery of full trajectories because it is not suited to the bioimaging applications investigated here. It thus illustrates the impracticality of using multiframe algorithms developed for other fields without major modifications.

Finally, the MHT procedure with a depth $d = 4$ was applied to the full sequence of images. It required 33 minutes to compute 9,850 trajectories from the 481 full-size frames of the sequence, which is fast in regard to the complexity and




Fig. 6. 9,850 trajectories of Golgi units built by the proposed MHT over 481 frames.

scale of the scene. Several thousand tracks lying in a small area of the surveillance volume are shown in Fig. 6. Results show the unique ability of the proposed algorithm to deal with high densities of targets in cluttered conditions and to take into account various types of movements.

6 DISCUSSION

In this paper, we have first proposed a Bayesian model that is suited to the problem of tracking multiple particles in microscopy image sequences. This model features an integrated target existence model as well as realistic motion models for nanometer-scale targets in cellular environments. We have established a physical interpretation of the parameters of these Markovian models so as to facilitate the adaptation of the framework to most of biological systems. Moreover, the various elements of the model are seamlessly integrated in the general framework so that every aspect of the problem is accounted for when deriving statistical quantities characterizing tracking solutions. Future improvements of the model will include application-specific features such as additional motion and imaging models, as well as integrated handling of split and merge events.

We have then described a new MHT algorithm for extracting object trajectories in image sequences by exploiting the proposed tracking model. The proposed MHT takes full advantage of the probabilistic framework for particle tracking, for a computational cost compatible with the large-scale nature of the problem. The solution space is exhaustively represented thanks to track trees, while the optimization is achieved with a linear program with nonconvex constraints. One major benefit of the technique is that all the tracking decisions are embedded in a single optimization problem, which ensures the consistency of the process.

As a result, we have shown with experiments on real and synthetic data that the MHT outperforms state-of-the-art techniques for particle tracking in low SNR microscopy images. The proposed algorithm is remarkably robust to clutter, which opens the way to new biological studies and thus may have a major impact in the bioimaging field. Overall, our work shows that bioimaging is a very promising domain for Bayesian methods and should

encourage one to revisit recent Bayesian tracking methods in this context.

ACKNOWLEDGMENTS

The authors are thankful to Prof. Karin Luby-Phelps for providing the Golgi image sequences and to Dr. Khuloud Jaqaman for having accepted to process the benchmark data with the methods presented in [20]. They are also thankful to the reviewers and the associate editor, whose constructive criticism and detailed suggestions have led to a much improved version of the paper. They are very much appreciative of the time and effort that the reviewers generously dedicated to help them strengthen the manuscript. This work was supported in part by the Institut Pasteur, the CNRS, and the Programme C'Nano of the Région IdF. Nicolas Chenouard received the PhD fellowship from the Programme C'Nano.

REFERENCES

- [1] N. Arhel, A. Genovesio, K.A. Kim, S. Miko, E. Perret, J.-C. Olivo-Marin, S. Shorte, and P. Charneau, "Quantitative Four-Dimensional Tracking of Cytoplasmic and Nuclear HIV-1 Complexes," *Nature Methods*, vol. 3, pp. 817-824, Oct. 2006.
- [2] Y. Bar-Shalom and W.D. Blair, *Multitarget-Multisensor Tracking Applications and Advances*, vol. 3. Artech House, Oct. 2000.
- [3] Y. Bar-Shalom, R.X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. Wiley-Blackwell, July 2001.
- [4] M. Berkelaar, K. Eikland, and P. Notebaert, *lp_solve version 5.5*, 2010.
- [5] S.S. Blackman, "Multiple Hypothesis Tracking for Multiple Target Tracking," *IEEE Trans. Aerospace Electronic Systems*, vol. 19, no. 1, pp. 5-18, Jan. 2004.
- [6] S.S. Blackman, R.J. Dempster, and R.W. Reed, "Demonstration of Multiple-Hypothesis Tracking (MHT) Practical Real-Time Implementation Feasibility," *Proc. SPIE*, vol. 4473, pp. 470-475, 2001.
- [7] S.S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [8] H.A.P. Blom and Y. Bar-Shalom, "The Interacting Multiple Model Algorithm for Systems with Markovian Switching Coefficients," *IEEE Trans. Automatic Control*, vol. 33, no. 8, pp. 780-783, Aug. 1988.
- [9] S. Bonneau, M. Dahan, and L.D. Cohen, "Single Quantum Dot Tracking Based on Perceptual Grouping Using Minimal Paths in a Spatiotemporal Volume," *IEEE Trans. Image Processing*, vol. 14, no. 9, pp. 1384-1395, Sept. 2005.
- [10] B. Brandenburg and X. Zhuang, "Virus Trafficking—Learning from Single-Virus Tracking," *Nature Rev. Microbiology*, vol. 5, no. 3, pp. 197-208, Mar. 2007.
- [11] M.K. Cheezum, W.F. Walker, and W.H. Guilford, "Quantitative Comparison of Algorithms for Tracking Single Fluorescent Particles," *Biophysical J.*, vol. 81, pp. 2378-2388, Oct. 2001.
- [12] D.-S. Chen, R.G. Batson, and Y. Dang, *Applied Integer Programming: Modeling and Solution*. Wiley & Sons, Jan. 2010.
- [13] N. Chenouard, "Advances in Probabilistic Particle Tracking for Biological Imaging," PhD thesis, Télécom ParisTech, Jan. 2010.
- [14] D. Chetverikov and J. Verestoi, "Feature Point Tracking for Incomplete Trajectories," *Computing*, vol. 62, pp. 321-338, 1999.
- [15] I.J. Cox and S.L. Hingorani, "An Efficient Implementation of Reid's Multiple Hypothesis Tracking Algorithm and Its Evaluation for the Purpose of Visual Tracking," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 2, pp. 138-150, Feb. 1996.
- [16] F. de Chaumont, S. Dallongeville, N. Chenouard, N. Herve, S. Pop, T. Provoost, V. Meas-Yedid, P. Pankajakshan, T. Lecomte, Y.L. Montagner, T. Lagache, A. Dufour, and J.-C. Olivo-Marin, "Icy: An Open Bioimage Informatics Platform for Extended Reproducible Research," *Nature Methods*, vol. 9, no. 7, pp. 690-696, 2012.
- [17] *Sequential Monte Carlo Methods in Practice*, A. Doucet, N. De Freitas, and N.J. Gordon, eds. Springer-Verlag, 2001.

- [18] A. Genovesio, T. Liedl, V. Emiliani, W.J. Parak, M. Coppey-Moisan, and J.-C. Olivo-Marin, "Multiple Particle Tracking in 3-D+t Microscopy: Methods and Application to the Tracking of Endocytosed Quantum Dots," *IEEE Trans. Image Processing*, vol. 15, no. 5, pp. 1062-1070, May 2006.
- [19] M. Guerriero, L. Svensson, D. Svensson, and P. Willett, "Shooting Two Birds with Two Bullets: How to Find Minimum Mean OSPA Estimates," *Proc. Int'l Conf. Information Fusion*, pp. 1-8, July 2010.
- [20] K. Jaqaman, D. Loerke, M. Mettlen, H. Kuwata, S. Grinstein, S.L.L. Schmid, and G. Danuser, "Robust Single-Particle Tracking in Live-Cell Time-Lapse Sequences," *Nature Methods*, vol. 5, no. 8, pp. 695-702, July 2008.
- [21] T. Kurien, "Issues in the Design of Practical Multitarget Tracking Algorithms," *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, ed., Artech House, 1990.
- [22] N. Li and X.R. Li, "Target Perceivability and Its Applications," *IEEE Trans. Signal Processing*, vol. 49, no. 11, pp. 2588-2604, Nov. 2001.
- [23] N. Li and X.R. Li, "Tracker Design Based on Target Perceivability," *IEEE Trans. Aerospace Electronic Systems*, vol. 37, no. 1, pp. 214-225, Jan. 2001.
- [24] R.P.S. Mahler, "Multitarget Bayes Filtering via First-Order Multitarget Moments," *IEEE Trans. Aerospace Electronic Systems*, vol. 39, no. 4, pp. 1152-1178, Oct. 2003.
- [25] R.P.S. Mahler, "PHD Filters of Higher Order in Target Number," *IEEE Trans. Aerospace Electronic Systems*, vol. 43, no. 4, pp. 1523-1543, Oct. 2007.
- [26] R.P.S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
- [27] C. Morefield, "Application of 0-1 Integer Programming to Multitarget Tracking Problems," *IEEE Trans. Automatic Control*, vol. 22, no. 3, pp. 302-312, June 1977.
- [28] D. Musicki and R.J. Evans, "Target Existence Based MHT," *Proc. IEEE Conf. Decision and Control*, pp. 1228-1233, Dec. 2005.
- [29] E. Nicolas, N. Chenouard, J.-C. Olivo-Marin, and A. Guichet, "A Dual Role for Actin and Microtubule Cytoskeleton in the Transport of Golgi Units from the Nurse Cells to the Oocyte across Ring Canals," *Molecular Biology Cell*, vol. 20, no. 1, pp. 556-568, Jan. 2009.
- [30] J.-C. Olivo-Marin, "Extraction of Spots in Biological Images Using Multiscale Products," *Pattern Recognition*, vol. 35, no. 9, pp. 1989-1996, 2002.
- [31] A.B. Poore and A.J. Robertson, "A New Lagrangian Relaxation Based Algorithm for a Class of Multidimensional Assignment Problems," *Computational Optimization Applications*, vol. 8, no. 2, pp. 129-150, Sept. 1997.
- [32] V. Racine, M. Sachse, J. Salermo, V. Fraisier, A. Trubuil, and J.-B. Sibarita, "Visualization and Quantification of Vesicle Trafficking on a Three-Dimensional Cytoskeleton Network in Living Cells," *J. Microscopy*, vol. 225, no. 3, pp. 214-228, Mar. 2006.
- [33] D. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Trans. Automatic Control*, vol. 24, no. 6, pp. 843-854, Dec. 1979.
- [34] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter*. Artech House, 2004.
- [35] B. Ristic, B.-N. Vo, D. Clark, and B.-T. Vo, "A Metric for Performance Evaluation of Multi-Target Tracking Algorithms," *IEEE Trans. Signal Processing*, vol. 59, no. 7, pp. 3452-3457, July 2011.
- [36] D. Sage, F.R. Neumann, F. Hediger, S.M. Gasser, and M. Unser, "Automatic Tracking of Individual Fluorescence Particles: Application to the Study of Chromosome Dynamics," *IEEE Trans. Image Processing*, vol. 14, no. 9, pp. 1372-1383, Sept. 2005.
- [37] M.J. Saxton and K. Jacobson, "Single-Particle Tracking: Applications to Membrane Dynamics," *Ann. Rev. Biophysics Biomolecular Structure*, vol. 26, pp. 373-399, June 1997.
- [38] I.F. Sbalzarini and P. Koumoutsakos, "Feature Point Tracking and Trajectory Analysis for Video Imaging in Cell Biology," *J. Structural Biology*, vol. 151, pp. 182-195, Aug. 2005.
- [39] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A Consistent Metric for Performance Evaluation of Multi-Object Filters," *IEEE Trans. Signal Processing*, vol. 56, no. 8, pp. 3447-3457, Aug. 2008.
- [40] 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.
- [41] I. Smal, M. Loog, W.J. Niessen, and E. Meijering, "Quantitative Comparison of Spot Detection Methods in Fluorescence Microscopy," *IEEE Trans. Medical Imaging*, vol. 29, no. 2, pp. 282-301, Feb. 2010.
- [42] I. Smal, E. Meijering, K. Draegestein, N. Galjart, I. Grigoriev, A. Akhmanova, M.E. van Royen, A.B. Houtsmuller, and W.J. Niessen, "Multiple Object Tracking in Molecular Bioimaging by Rao-Blackwellized Marginal Particle Filtering," *Medical Image Analysis*, vol. 12, no. 6, pp. 764-777, Dec. 2008.
- [43] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, "Set JPDA Filter for Multitarget Tracking," *IEEE Trans. Signal Processing*, vol. 59, no. 10, pp. 4677-4691, Oct. 2011.
- [44] D. Thomann, D.R. Rines, P.K. Sorger, and G. Danuser, "Automatic Fluorescent Tag Detection in 3D with Super-Resolution: Application to the Analysis of Chromosome Movement," *J. Microscopy*, vol. 208, no. 1, pp. 49-64, Oct. 2002.
- [45] J. Vermaak, S. Maskell, and M. Briers, "A Unifying Framework for Multi-Target Tracking and Existence," *Proc. Int'l Conf. Information Fusion*, vol. 1, July 2005.
- [46] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic Implementations of the Cardinalized Probability Hypothesis Density Filter," *IEEE Trans. Signal Processing*, vol. 55, no. 7, pp. 3553-3567, July 2007.
- [47] C. Vonesch, F. Aguet, J.-L. Vonesch, and M. Unser, "The Colored Revolution of Bioimaging," *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 20-31, May 2006.



Nicolas Chenouard received the MS degree in engineering and bioinformatics from the Institut National des Sciences Appliquées Lyon, France, and the MSc degree in artificial intelligence and data mining from the University Pierre et Marie Curie (Paris VI), Paris, France, both in 2006, and the PhD degree in image and signal processing from the Institut Pasteur and Telecom ParisTech, Paris, in 2010. His doctoral work has focused on the development of Bayesian methods for multiple particle tracking in sequences of microscopy images of live biological samples. From 2010 to 2012, he was a postdoctoral fellow with the Biomedical Imaging Group and the Center for Biomedical Imaging, École Polytechnique Fédérale de Lausanne, Switzerland, working on the design of steerable wavelet frames in multiple dimensions and learning techniques for adaptive frame design. Since March 2012, he has been a postdoctoral fellow at the New York University School of Medicine, New York. His current research in neuroscience investigates synaptic mechanisms by using advanced microscopy and image processing techniques.



Isabelle Bloch graduated from the Ecole des Mines de Paris, France, in 1986. She received the master's degree from the University Paris 12 in 1987, the PhD degree from the Ecole Nationale Supérieure des Télécommunications (Telecom ParisTech), Paris, in 1990, and the Habilitation degree from the University Paris 5 in 1995. She is currently a professor in the Signal and Image Processing Department, Telecom ParisTech, in charge of the Image Processing

and Understanding Group. Her research interests include 3D image and object processing, computer vision, 3D and fuzzy mathematical morphology, information fusion, fuzzy set theory, structural, graph-based, and knowledge-based object recognition, spatial reasoning, and medical imaging. She is a member of the IEEE.



Jean-Christophe Olivo-Marin received the PhD and the HDR degrees in optics and signal processing from the Institut d'Optique Théorique et Appliquée, University of Paris-Orsay, France. He is the head of the Quantitative Image Analysis Unit, Institut Pasteur, and the chair of the Cell Biology and Infection Department. He was a cofounder of the Institut Pasteur Korea, Seoul, where he held a joint appointment as a chief technology officer from 2004 to 2005.

Previous to that, he was a staff scientist at the European Molecular Biology Laboratory, Heidelberg, from 1990 to 1998. His research interests are in image analysis of multidimensional microscopy images, computer vision and motion analysis for cellular dynamics, and in multidisciplinary approaches for biological imaging. He is a fellow of the IEEE, past chair of the IEEE SPS Bio Imaging and Signal Processing Technical Committee (BISP-TC), a senior area editor of the *IEEE Signal Processing Letters*, and a member of the Editorial Board of the journals *Medical Image Analysis* and *BMC Bioinformatics*. He was the general chair of the IEEE International Symposium on Biomedical Imaging in 2008.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**