

Single Quantum Dot Tracking Based on Perceptual Grouping Using Minimal Paths in a Spatiotemporal Volume

Stéphane Bonneau, Maxime Dahan, and Laurent D. Cohen, *Senior Member, IEEE*

Abstract—Semiconductor quantum dots (QDs) are new fluorescent probes with great promise for ultrasensitive biological imaging. When detected at the single-molecule level, QD-tagged molecules can be observed and tracked in the membrane of live cells over unprecedented durations. The motion of these individual molecules, recorded in sequences of fluorescence images, can reveal aspects of the dynamics of cellular processes that remain hidden in conventional ensemble imaging. Due to QD complex optical properties, such as fluorescence intermittency, the quantitative analysis of these sequences is, however, challenging and requires advanced algorithms. We present here a novel approach, which, instead of a frame by frame analysis, is based on perceptual grouping in a spatiotemporal volume. By applying a detection process based on an image fluorescence model, we first obtain an unstructured set of points. Individual molecular trajectories are then considered as minimal paths in a Riemannian metric derived from the fluorescence image stack. These paths are computed with a variant of the fast marching method and few parameters are required. We demonstrate the ability of our algorithm to track intermittent objects both in sequences of synthetic data and in experimental measurements obtained with individual QD-tagged receptors in the membrane of live neurons. While developed for tracking QDs, this method can, however, be used with any fluorescent probes.

Index Terms—Active contours, cellular imaging, energy minimization, group marching, minimal paths, perceptual grouping, quantum dot, single-molecule tracking (SMT).

I. INTRODUCTION

IMAGING techniques have become central in the development of cell biology since they offer a tool of choice to decipher the complex mechanisms controlling the organization of live cells [1], [2]. The ability to observe biological processes in real time is essential to determine their spatial and temporal dynamics. Important mechanisms such as molecular diffusion, gene expression, protein-protein interactions or cell signaling can now be detected optically [2], [3]. The need to quantitatively analyze them has, thus, spurred a large effort which involves techniques from cell biology, biochemistry, physics, or computer science.

Manuscript received December 6, 2004; revised May 19, 2005. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Robert F. Murphy.

S. Bonneau and L. D. Cohen are with CEREMADE, Université Paris-Dauphine, 75775 Paris cedex 16, France (e-mail: bonneau@ceremade.dauphine.fr; cohen@ceremade.dauphine.fr).

M. Dahan is with the Laboratoire Kastler Brossel, Département de Physique, Ecole Normale Supérieure, 75231 Paris cedex 05, France (e-mail: maxime.dahan@lkb.ens.fr).

Digital Object Identifier 10.1109/TIP.2005.852794

Several technologies have contributed to rapid advances in the field of biological imaging: i) the advent of genetically encoded markers, such as green fluorescent protein (GFP) and its variants, has truly revolutionized the biochemical methods required to selectively label molecules or cellular compartments and record their properties in living systems [4]; ii) following progress in opto-electronics, optical instruments have significantly gained in sensitivity and dynamics; and iii) the acquisition and processing of large amount of data have been made possible by the ever enhanced capabilities of computers.

Recent achievements in nanosciences have also been predicted to make a decisive impact on biological detection [5]. New nanomaterials are progressively entering the world of cell biology and their enhanced physical (optical, electrical, or mechanical) properties offer great promise for advanced imaging. Among these nanomaterials, semiconductor quantum dots (QDs) are certainly the most prominent example [6], [7]. Their interest for biological imaging stems from a combination of photophysical properties. First, they possess both a large absorption and a narrow emission spectrum, and are, thus, ideally suited for multicolor detection. Second, their photostability, far superior to conventional fluorophores (fluorescent proteins or organic dyes), overcomes the traditional limitation of photodegradation in fluorescent labeling. Finally, the synthesis and engineering of nanoparticles with different semiconductor materials and structures have expanded the range of possible emission wavelengths to the red and infrared regions, more appropriate for imaging in tissues or animals [8]. The potential of QDs as biological probes has now been demonstrated in many cases and their applications, in live cells or organisms, are rapidly multiplying [6], [7].

Among the variety of experimental techniques aiming at measuring molecular dynamics in live cells [3], single-particle tracking (SPT) is certainly one of the most sensitive. In this approach, a marker is specifically attached to a protein of interest whose lateral motion is then recorded with high spatial (about 10 nm) and temporal resolution (in the millisecond range). Such trajectories contain a wealth of information on the local organization of the membrane. In particular, they may potentially reveal transient interactions or temporary confinement that remain hidden with conventional imaging.

SPT was initially developed using large labels such as micron-sized latex beads or 40-nm gold nanoparticles [9]. Progress in single-molecule biophysics have recently allowed their replacement by smaller fluorescent probes (GFPs, fluorophores, or QDs) that enable investigations at a truly molecular




Fig. 1. (a) Time trace of the fluorescence intensity for a single QD. (b) Two-dimensional (2-D) fluorescence image. The image has been saturated for clarity, as indicated by the inverted gray scale map on the right. The circles correspond to the result of the detection stage. (c) Illustration of the peak selection on a fluorescence image profile. The two spot candidates (located along the same line x_1) pass the detection test.

scale. For single-molecule tracking (SMT), QDs are probably the most favorable probe since they combine a relatively small size (5–15 nm) with a remarkable brightness and a superior photostability, allowing long-term acquisition with a good signal-to-noise ratio (SNR). Using these nanoprobe, individual glycine receptors have for instance been tracked in the membrane of live neurons for durations over 20 min [10].

While very promising for the study of cell dynamics, the development of SMT raises, however, new challenges in terms of image processing. In these experiments, the motion of several tagged molecules is simultaneously recorded and needs to be quantitatively analyzed. The image sequences can be composed of up to thousands of frames and contain a number of molecules which varies from a couple to hundreds, depending on experimental conditions and protocols. This number may also fluctuate in consecutive frames since molecules can enter and exit the field of view. Using QDs as fluorescent probes raises additional difficulties. Due to complex physical processes, the QD fluorescence is intermittent, i.e., its emission intensity randomly alternates between bright and dark periods [see Fig. 1(a)] [11]. It means that tagged-molecules might temporarily disappear from the field of view. A robust processing tool is, therefore, required to account for all the specificities of SMT measurements and automatically extract relevant information from huge amounts of data.

Common approaches to motion tracking in dynamic fluorescence imaging analyze sequentially each frame of the image stack in a two-stage procedure [12]–[15]. For each frame, fluorescent spots are first detected and localized. These spots are then associated with trajectories obtained from the previous frames. In SMT experiments, the number of photons emitted by the probe is intrinsically limited and images can be noisy. Therefore, the detection stage is generally based on a noise removal process, such as anisotropic diffusion [12], wavelet thresholding [13] or matched filtering based on a model of the optical system [14], [15]. Since the probes are much smaller than their emission wavelength, the fluorescence spots correspond to diffraction-limited patterns. Neither shape nor intensity are, therefore, appropriate to identify the same particle over multiple frames. So, the only information available for the

association stage is location. Finding corresponding points in an image sequence solely with location information is referred to as the motion correspondence problem. This problem has been addressed with numerous methods, either in a statistical [16]–[18] or heuristic framework [19]–[21]. These approaches critically rely on the detection stage and can fail when ambiguities occur due to events like fusion of objects, missing detections, false detections or appearance/disappearance of multiple targets.

So far, no definitive approach has been established to simultaneously track multiple fluorescent spots which might transiently disappear. In this paper, we present a robust method to fulfill this challenging task. Our method does not use the traditional frame by frame approach, but considers the fluorescence image stack as a single three-dimensional (3-D) spatiotemporal volume. In this representation, molecular trajectories are viewed as 3-D curves which we consider as minimal paths in an image-dependent metric and seek to retrieve using an energy minimization technique [22]. First, fluorescent spots are detected in each frame of the sequence, using a model of the diffraction-limited pattern. The only required parameter is a threshold defined as the minimal SNR allowed for detection. This step gives rise to a set of points in the spatiotemporal volume. These points are not randomly scattered but exhibit significant amount of organization. The problem of motion correspondence is then mapped into a perceptual grouping one, given an unstructured set of points and the image sequence. The underlying idea is to use the fluorescence image to guide the grouping process, taking into account the fluorescence signal which exists but is too low to have been detected. We sequentially find pairs of points (among the set of detected spots) that have to be linked and the paths that join them, such that these paths are geodesics in a Riemannian metric computed from the 3-D image. Our approach naturally allows the inference of the fluorescent probe trajectory even when the fluorescence signal is below the detection threshold. It does not require an a priori knowledge of the number of molecules to track and accept going in or out from the field of view. It incorporates a model of the molecule dynamics in order to prevent unlikely matchings and reduce the problem complexity. Finally, an additional stage is applied to process the set of minimal paths,

in order to get true molecular trajectories with a localization accuracy of a few nanometers.

Our approach has been primarily developed to account for fluorescence intermittency in single-QD tracking measurements. It is, however, not limited to this case and can be used to analyze the motion of other punctual objects (other fluorescent markers, small GFP-tagged organelles, etc.), more particularly when their signal exhibit large intensity fluctuations. When tracking individual objects in live cells, these fluctuations can originate in the photophysical properties of the probes as well in displacements along the axial optical direction.

The remainder of this paper is organized as follows. In Section II, we present the detection stage. In Section III, we discuss the motion correspondence problem and describe our new approach based on minimal paths in order to link the fluorescent spots which have been detected. We also discuss the design of a practical implementation of this method. In Section IV, we explain how true molecular trajectories are extracted from the set of reconstructed tracks. Finally, in Sections V and VI, we present tracking results obtained from synthetic and real-world data, respectively.

In the following, we consider the fluorescence image stack as a single 3-D spatiotemporal volume $I(x, y, z)$ with spatial coordinates (x, y) and where z is the frame index. The size of this volume will be noted $N_x \times N_y \times N_z$. For the sake of clarity, the frame $I(\cdot, \cdot, z)$ is also denoted as $I_{|z}$.

II. DETECTION OF FLUORESCENT SPOTS

In this section, we propose an approach to automatically detect fluorescent spots. Our method is based on a model of the fluorescence image and allows the localization of spots with subpixel accuracy. The detection stage is sequentially applied to each of the N_z 2-D images which constitute the whole sequence.

A. Fluorescence Image Model

The size of a QD being much smaller than its wavelength λ , it acts as a punctual source of light. The response of the optical system (the point spread function, PSF) can, therefore, be approximated by an isotropic 2-D gaussian function with a standard deviation $\sigma_{\text{PSF}} = R_A/3$, where $R_A \approx 1.22\lambda/2N_A$ denotes the Airy radius and N_A the numerical aperture. Such a PSF can be discretized on a small support of size $M_x \times M_y$, where M_x and M_y are proportional to σ_{PSF} . In our simulations and experiments (detailed in Sections V and VI), the size of a (square) pixel is $p_w = 216.7$ nm, $N_A = 1.45$ and $\lambda = 605$ nm, so that $\sigma_{\text{PSF}} \approx 84.84$ nm ≈ 0.39 pixels and $M_x = M_y = 5$.

Neglecting the motion of the QDs during acquisition, a planar fluorescence image is a snapshot of the sample at a given time and results from the convolution of the PSF with the spatial distribution of point sources [see Fig. 1(b)].

The intensity of each pixel in the image corresponds to the number of detected photons added to a constant baseline offset B_L . We assume that fluorescence images are affected by the superposition of two independent noises: shot noise and dark noise [23]. Shot noise, due to the quantum nature of light, is Poisson

distributed with a variance equal to the number of detected photons. Dark noise refers to the noise caused by background photons, CCD readout noise and dark current. We model it by an additive zero-mean gaussian distribution with a variance σ_N^2 independent of the number of detected photons. The uncertainty in the intensity $I_{|z}(x, y)$ is given by

$$\text{Var}[I_{|z}(x, y)] = I_{|z}(x, y) - B_L + \sigma_N^2.$$

B. Estimation of the Background and σ_N

Let us assume that most of the pixels of $I_{|z}$ corresponds to the image background. It is not a strong assumption if the concentration of tagged molecules is moderate. One can assume that the intensity distribution of the image is the mixture of two gaussian probability density functions (PDFs): a narrow one, which describes the intensity distribution of the background, and a broader one, which describes the distribution of photons emitted by the QDs.

An estimate of the background mean value in $I_{|z}$ can be obtained by applying the well-known EM algorithm [24] to the intensity distribution of an averaged version of $I_{|z}$. This estimate, noted $\widehat{B}_L(z)$, is given by the mean of the PDF with the smallest mean and the biggest weight among the two gaussian PDFs obtained after convergence of the EM algorithm. The averaged image, noted $I_{|z}^m$, results from the convolution of $I_{|z}$ with a mask taking a constant value $1/(M_x M_y)$ on its support of size $M_x \times M_y$. Note that $\widehat{B}_L(z)$ decreases over z due to the residual autofluorescence and asymptotically tends to the baseline offset B_L . If the illumination field is homogeneous over the field of view, $\widehat{B}_L(z)$ is an accurate estimate of the background for the whole plane image $I_{|z}$.

Let $I_{|z}^\sigma(x, y)$ be the standard deviation of $I_{|z}(x, y)$ in a neighborhood of size $M_x \times M_y$. Therefore, $I_{|z}^\sigma$ is the square root of the convolution product between $(I_{|z} - I_{|z}^m)^2$ and a mask taking a constant value $1/[(M_x - 1)(M_y - 1)]$ on its support of size $M_x \times M_y$. Here again, one can assume that the PDF for grey levels in $I_{|z}^\sigma$ is the mixture of two gaussian PDFs. Let $\widehat{\sigma}_N(z)$ be the mean of the gaussian PDF with the biggest weight among the two PDFs obtained by applying the EM algorithm to the intensity distribution of $I_{|z}^\sigma$. The value of $\widehat{\sigma}_N(z)$ is expected to be an estimate of σ_N in the image $I_{|z}$.

C. Template Matching and Correlation Peak Selection

The fluorescent spots to detect present two convenient properties: They are isotropic and have a known spread. For such template matching, one of the most efficient method is the normalized cross-correlation. In our case, the template is the PSF, i.e., an isotropic 2-D gaussian function with standard deviation σ_{PSF} , discretized on a support of size $M_x \times M_y$. Let us note that G is this template, G^m is its intensity mean over the support, and G^σ is its intensity standard deviation. The normalized cross-correlation image $\gamma_{|z}$ is computed in the following manner (here, \star denotes the correlation operator):

$$\gamma_{|z} = \frac{(I_{|z} - I_{|z}^m) \star (G - G^m)}{I_{|z}^\sigma G^\sigma}.$$

Note that the computation of $I_{|z}^m$, $I_{|z}^\sigma$ and $\gamma_{|z}$ is really fast using the fast Fourier transform. The function $\gamma_{|z}$ is a similarity measure between the template G and the fluorescence image $I_{|z}$. To obtain a set of candidate spot locations, we apply to $\gamma_{|z}$ a local maximum detection in a 3×3 neighborhood.

Some correlation maxima are due to noise in the fluorescence image, so one needs to select the most reliable ones among the set of candidate spot locations. Let (x_m, y_m) be the coordinates of one of the correlation maxima. We can estimate the uncertainty in the intensity $I_{|z}(x_m, y_m)$ by

$$\text{Var}[I_{|z}(x_m, y_m)] \approx I_{|z}(x_m, y_m) - \widehat{B}_L(z) + \widehat{\sigma}_N^2(z).$$

Let $\Delta I_{|z}(x_m, y_m) = d_T \sqrt{\text{Var}[I_{|z}(x_m, y_m)]}$ be a confidence interval where d_T is a threshold corresponding to the minimal SNR allowed for detection. One decides that a spot is located at (x_m, y_m) in the frame $I_{|z}$ if [see Fig. 1(c)]

$$I_{|z}(x_m, y_m) > \widehat{B}_L(z) + \Delta I_{|z}(x_m, y_m).$$

D. Spot Location Refinement

In order to rapidly get a subpixel estimate of the location of each detected spot, we apply a three point estimator in the x - and y - directions. Let (x_m, y_m) be the pixel coordinates of a spot detected in the frame $I_{|z}$. The subpixel coordinate of that spot along the x axis, denoted x_0 , is estimated according to the following relation, based on the assumption that the spot profile is gaussian

$$x_0 = x_m + \frac{J_{(-1,0)} - J_{(1,0)}}{2J_{(-1,0)} - 4J_{(0,0)} + 2J_{(1,0)}}$$

where $J_{(i,j)} = \log[I_{|z}(x_m + i, y_m + j) - \widehat{B}_L(z)]$. The same process is performed in the orthogonal direction, in order to get a subpixel coordinate y_0 . The accuracy of this three point estimator, evaluated from Monte Carlo simulations, is about 0.1 pixel (about 20 nm) along each axis for a SNR greater than 10.

III. GROUPING SPOTS TO FORM TRAJECTORIES

Applying the detection algorithm to each frame of the stack leads to a set of points in the 3-D spatiotemporal volume [see Fig. 4(b)]. In this section, we present a method which aims at clustering this cloud of points into a set of tracks, so that each track describes the motion of one fluorescent probe.

A. Definition

A track is a 3-D curve described by a list of points (x, y, z) defined in the 2-D+T space and ordered by increasing z coordinates. For a given point (x, y, z) of a track, the spatial coordinates x and y can take real values (spatial position is given with subpixel accuracy), but the temporal coordinate z can only take an integer value corresponding to a frame index. The first point of a track (i.e., the point with the minimal z coordinate) is called *tail*, and the last one *head*. A track is defined in each successive frame between the frame of its *tail* and the frame of its *head*.

Our goal is to find tracks defined over time range as large as possible. For this, we first obtain short tracks as sets of spots

detected in consecutive frames. We then group and complete these tracks to form wholes trajectories of molecules.

B. Tracking Constraints

A *uniqueness constraint* is adopted: each detected spot is associated to no more than one track and each track is associated to exactly one spot detected in a given frame.

In order to build a track, a first step is to link together neighboring spots detected in consecutive frames. In case the trajectory is not continuously visible, a spot could also be linked to another one in a distant frame, providing it is not too far spatially from the first one. We now detail a spatiotemporal likelihood criterion based on a physical model of molecule dynamics, in order to determine whether two spots belong to the same trajectory.

Early models of the plasma membrane (notably the fluid mosaic model [9]) postulated that proteins, homogeneously distributed within the membrane, move by free diffusion in a lipid bilayer. Although SMT experiments have revealed a large diversity in the motion of individual membrane molecules, the most general assumption is still that protein dynamics is a planar Brownian motion. In a homogeneous 2-D system, the probability distribution \mathcal{P} that a single-Brownian molecule moves by a distance r in a time period τ is

$$\mathcal{P}(r, \tau, D) = \left(\frac{1}{4\pi D\tau} \right) \exp\left(-\frac{r^2}{4D\tau} \right)$$

where D is the diffusion coefficient of the molecule. Let Ψ be the probability that the maximal distance covered by this molecule in Δ_z frame is r_Ψ

$$\Psi = \int_0^{r_\Psi} \mathcal{P}(r, t_s \Delta_z, D) 2\pi r dr$$

where t_s is the temporal sampling period. If the value of Ψ is a priori fixed, then $r_\Psi = c_\Psi \sqrt{D\Delta_z}$ where D is expressed in pixel²/frame and $c_\Psi = \sqrt{4|\log(1 - \Psi)|}$ is a constant. To each detected spot, we can, therefore, associate a paraboloidic search region [see Fig. 2(a)] defined by every point (x, y, z) satisfying

$$\sqrt{(x - x_s)^2 + (y - y_s)^2} \leq c_\Psi \sqrt{D|z - z_s|} \quad (1)$$

where (x_s, y_s, z_s) are the spot coordinates, c_Ψ is a constant solely dependent of the probability Ψ (which is assumed to be given) and D is an estimate of the diffusion coefficient of the molecule represented by the spot. In the association process, the *likelihood constraint* states that two spots can only be linked if at least one of them belongs to the search region associated to the other.

C. Trajectories Completed by Minimal Paths

The small size of QDs and large fluctuations of their intensities preclude any association of detected spots based on an appearance identification. Matching points over multiple frames solely using position information is known to be a NP-hard problem. To overcome this difficult problem, we propose to use the low fluorescence signal, too weak to have been detected, to guide the association process.

Rather than adopting the usual frame by frame approach, we consider the fluorescence image stack as a single 3-D spatiotem-

poral volume. In this volume, the tracks we want to retrieve are 3-D curves. To construct these curves, we iteratively find pairs of points (among the set of detected spots) that have to be linked and the paths that join them, such that these paths are geodesics in a Riemannian metric computed from the 3-D image I .

Our grouping method is divided into two main steps. First, we sequentially match spots detected in successive frames using, at each iteration, a nearest-neighbor criterion in the sense of a weighted distance. This step leads to a set of *partial tracks*,¹ a partial track being only composed of points detected in consecutive frames. Due to missing detections and blinking of QDs, several partial tracks defined on disjoint temporal intervals might correspond to the trajectory of the same molecule. Then, partial tracks are linked with minimal paths to constitute *completed tracks*. This second step can be seen as an extension to visual tracking of the more general approach aiming at completing curves with minimal paths [25], [26]. However, many original aspects have been included in our method. For instance, the construction of minimal paths takes into account information from both image features and tracking constraints. Moreover, each time a minimal path is added to a trajectory, image information is removed along the path in order to avoid trajectories to merge.

D. Image Dependent Metric

The motion correspondence problem is formulated as the search for a set of minimal paths in a metric space that depends on features of the 3-D image. Consequently, a relevant metric has to be defined. As explained below (see Section III-F), a minimal path between two given points is a curve joining them, along which the integral of a potential P is minimal. Since we wish the minimal path to correspond to the temporal trace of fluorescent particle, the path should follow high values of the fluorescence image. Therefore, we can simply define the potential P as the inverse of the fluorescence image I

$$P(x, y, z) = \frac{1}{I(x, y, z) + 1}.$$

Thus, P is a 3-D volume of size $N_x \times N_y \times N_z$ taking low values where the fluorescence signal is high. Other functions of I can also be considered in order to get a potential enhancing bright spots with respect to the background or vessel-like structures in the 3-D image [27]. In this paper, however, we limit ourselves to the simple expression above.

E. Construction of Partial Tracks

Pairs of spots detected in two consecutive frames and satisfying the likelihood constraint are sequentially matched to constitute a set of partial tracks. At each iteration the pair of spots with the minimal association cost is chosen, this cost being given by the integral of the potential P along the line segment separating the two spots.

To apply the likelihood constraint, the radius of a circular search area, which corresponds to the maximal distance covered by a molecule between two frames, is computed from the probability Ψ , which is an input parameter, and from an estimate of the diffusion coefficient. If the partial track is sufficiently long, the diffusion coefficient D assigned to the spot which is the *tail*

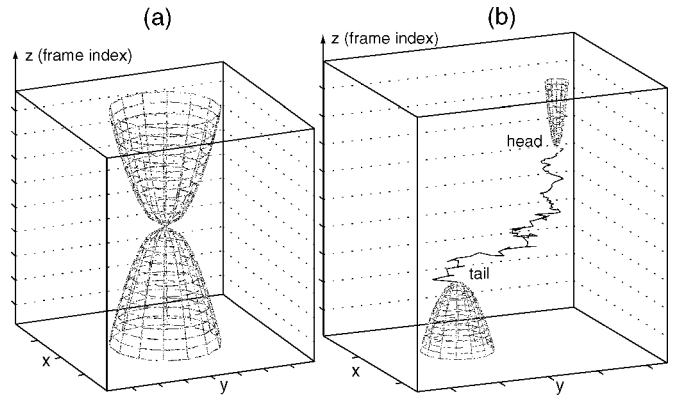


Fig. 2. (a) For each spot, a paraboloidic search region is defined in the spatiotemporal volume. (b) If molecule dynamics vary over time, the spatial spread of search regions associated to the *tail* and the *head* of a single-partial track may highly differ, reflecting the dynamical change.

(or the *head*) of a partial track can be easily estimated by using a windowed version of the mean-square displacement function (see the Appendix for details). It, thus, enables the algorithm to dynamically adapt the search regions in order to account for sudden variations in the molecule dynamics [see Fig. 2(b)]. Otherwise, an arbitrary value, D_{ini} , is assigned to the spot.

The proposed algorithm is given in Table I. It automatically stops when every pair of spots satisfying the likelihood constraint has been linked. Fig. 4(c) shows some typical results obtained with this approach.

The advantages of this algorithm are threefold. First, the association is done globally in the whole spatiotemporal volume. Therefore, it is more robust than a frame by frame method which only use the information obtained from the previous frames. Second, it is much less sensitive to false detections since it removes isolated spurious points among the set of detected spots. Finally, it is a fast algorithm, easy to implement.

F. Background on Minimal Paths

Before explaining how partial tracks are linked, we give a succinct description of prior works related to minimal paths. Energy minimization techniques have been applied to a variety of problems in image processing and computer vision. Since the original work on snakes [28], they have notably been used for boundary detection. In brief, an active contour model, or snake, is a curve that deforms its shape in order to minimize, along itself, an energy combining an internal part (which smooths the curve) and an external part (which guides the curve toward some particular features of the image). Given an initial curve, the minimization of the energy often leads to local minima. To be less sensitive to local minima, Cohen and Kimmel [29] have proposed to modify this energy and, provided that two endpoints of the curve are initially given, to find the global minimum of

$$E(C) = \int_C \{w + P(C(s))\} ds = \int_C \tilde{P}(C(s)) ds \quad (2)$$

where s represents the arc-length parameter, w is a regularization term² and P a positive potential which takes lower values

²It has been shown in [29] that increasing w decreases the curvature magnitude of the geodesics minimizing (2). Thus, the constant w controls the smoothness of the minimal path and can be considered as a regularization term.

¹Both partial and completed tracks are defined as tracks (see Section III-A).

TABLE I
 ALGORITHM TO CONSTRUCT PARTIAL TRACKS

Definition:

- S_i^z refers to the i^{th} spot detected in the frame of index z . Coordinates of this spot are noted (x_i^z, y_i^z, z) .
- $M^{w,z}$ is the weighted distance matrix between spots detected in successive frames of index z and $(z+1)$. Therefore, $M_{i,j}^{w,z}$ is estimated by integrating the potential P along a straight raypath going from (x_i^z, y_i^z, z) to $(x_j^{z+1}, y_j^{z+1}, z+1)$.
- $M^{e,z}$ is the planar Euclidean distance matrix between spots detected in successive frames z and $(z+1)$, such that:

$$M_{i,j}^{e,z} = \sqrt{(x_i^z - x_j^{z+1})^2 + (y_i^z - y_j^{z+1})^2}.$$

- R^z is a one-dimensionnal array such that R_i^z is the radius of the circular search aera associated to the spot S_i^z .
- L is the set of pairs of spots satisfying the likelihood constraint. Thus, L is the list of triplets (i, j, z) defined by:

$$L = \left\{ (i, j, z) \text{ such as } \{ M_{i,j}^{e,z} \leq R_i^z \text{ or } M_{i,j}^{e,z} \leq R_j^{z+1} \} \right\}.$$

Initialization:

- For each frame index z , compute $M^{w,z}$, $M^{e,z}$ and set $R^z = c_{\Psi} \sqrt{D_{\text{ini}}}$, where $c_{\Psi} = \sqrt{4|\log(1 - \Psi)|}$.
- Create the list L .

Loop:

While L is not empty, do:

- Find (i^*, j^*, z^*) such that $M_{i^*,j^*}^{w,z^*} = \min_{(i,j,z) \in L} \{ M_{i,j}^{w,z} \}$.
- Set $M_{i^*,j^*}^{e,z^*} = \infty$ and $M_{j^*,i^*}^{e,z^*} = \infty$. Update L .
- The spots $S_{i^*}^{z^*}$ and $S_{j^*}^{z^*+1}$ have to be linked. Three cases can occur: **(1)** None of the two spots to link has already been assigned to a partial track. Then, a new partial track is created with these two spots. **(2)** Only one of the two spots to link has already been assigned to a partial track. Then, the other spot is added to the partial track. **(3)** The two spots to link have been assigned to different partial tracks. Then, the two partial tracks are merged.
- In the previous step, each time the cases **(2)** or **(3)** occur, the radii of the circular search aera associated to the tail and the head of the partial track are updated by estimating the diffusion coefficient of the track (see the Appendix).

near desired features in a 2-D image. This energy $E(C)$ is similar to the one of geodesics [30].

The solution of this minimization problem is obtained through the computation of the minimal action map \mathcal{U} , which is the minimal energy integrated along a path between the starting point p_0 and any given point p of the plane

$$\mathcal{U}(p) = \inf_{C \in \mathcal{A}_{p_0,p}} E(C) = \inf_{C \in \mathcal{A}_{p_0,p}} \left\{ \int_C \tilde{P}(C(s)) ds \right\}$$

where $\mathcal{A}_{p_0,p}$ is the set of all paths between p_0 and p . Since we assumed $P > 0$, the action map \mathcal{U} has a convex like behavior and only one local minimum: the starting point p_0 . The geodesic path connecting p_0 and a point p_1 is easily found by a backpropagation from p_1 to p_0 , solving

$$\frac{dC(s)}{ds}(s) = -\nabla \mathcal{U} \text{ with } C(0) = p_1 \text{ and } C(\ell_C) = p_0$$

where ℓ_C is the length of the curve C . This backpropagation is a simple gradient descent on \mathcal{U} from p_1 until p_0 is reached.

Noting that the minimal action map satisfies the Eikonal equation

$$\|\nabla \mathcal{U}\| = \tilde{P} \text{ and } \mathcal{U}(p_0) = 0 \quad (3)$$




Fig. 3. (a) Potential P . Low values are displayed in a dark-gray tone. We seek a minimal path joining the two points defined in the top right and the bottom left of the image. (b) Minimal action map \mathcal{U} obtained by simultaneously propagating two fronts until a shock occurs. A backpropagation on \mathcal{U} from the meeting point, which is a saddle-point of \mathcal{U} , gives the minimal path joining the two sources.

the map \mathcal{U} can be efficiently computed with the *fast marching method* (FMM) [22], [29], [31]. This algorithm is particularly appropriate to get the minimal path between two points because the computation of the map \mathcal{U} in the whole grid is not required. Indeed, the minimal path between p_0 and p_1 can be computed using the FMM by simultaneously expanding wavefronts from p_0 and p_1 until a shock occurs between the two fronts [25], [26]. The minimal path passes by the point where the two fronts meet, which is a saddle-point of \mathcal{U} . Therefore, a backpropagation on \mathcal{U} from this meeting point to p_0 and p_1 provides the searched path (see Fig. 3).

This property has been used in [25] to reconstruct a set of significant curves given an unstructured set of points. A similar approach is adopted here to find the partial tracks that have to be matched and the paths that link them.

G. Completion of Partial Tracks

Preprocessing: Once the partial tracks have been obtained, we remove them from the potential [see Fig. 4(d)] in order to prevent minimal paths from following the same fluorescence traces. The potential $\tilde{P} = w + P$ is modified along every partial track in the following way. If no fluorescence signal had been recorded at (x, y, z) , $\tilde{P}_{x,y,z}$ would have taken a value close to

$$\tilde{P}_{\text{mod}}(z) = w + \frac{1}{\widehat{B}_L(z) + 1}$$

where $\widehat{B}_L(z)$ is an estimate of the background in the image $I|_z$. For each point (x, y, z) belonging to a partial track, we attribute to the potential the value $\tilde{P}_{\text{mod}}(z)$ in the vicinity of (x, y, z) . This vicinity corresponds to the size of a fluorescent spot in a planar image, so the potential modification is made in a neighborhood of size $M_x \times M_y \times 1$ around (x, y, z) .

Tracking Constraints: To obtain completed tracks, we sequentially find the minimal path joining the *tail* of a partial track to the *head* of another one. The set of *tails* and *heads* of partial




Fig. 4. Illustration of the association process. The motion of four Brownian particles has been simulated along a single line of a synthetic image stack. Every displayed image is a sectioning view of the stack showing the xz plane where this motion occurs. (a) Fluorescence image showing the fluorescence trace of the four particles. (b) Cloud of points issued from the detection stage. (c) Partial tracks obtained by clustering the cloud of points issued from the detection stage. (d) Potential P , defined as the inverse of the fluorescence image, after removal of signal along each partial track. (e)–(g) Minimal action map \mathcal{U} obtained at different iterations. (h) Tracks resulting from the completion of partial tracks with minimal paths. The trajectories of moving particles are correctly found and completed where no fluorescence signal is present.

tracks constitute a set of sources from which fronts are simultaneously propagated, according to the Eikonal equation (3). Two kinds of expanding wavefronts are distinguished [see Fig. 4(g)].

- From the *tail* of a partial track, a front is propagated in the downward direction (in the direction of decreasing z).
- From the *head* of a partial track, a front is propagated in the upward direction (in the direction of increasing z).

If we want the likelihood constraint to be satisfied, the range of influence³ of each source must necessarily be the search region associated with the source. Consequently, each front propagation is restricted to a paraboloidic domain defined by (1). Furthermore, the uniqueness constraint implies that the *tail* (or the *head*) of a partial track should not be linked to more than one extremity of another partial track.

Adding Minimal Paths: The condition for a *head* and a *tail* to be connected is that their fronts meet before other fronts. Fronts are propagated until a downward front and an upward front meet. A backpropagation on \mathcal{U} from the meeting point gives the minimal path which links the *tail* of a partial track to the closest (in the sense of a weighted distance) *head* of another partial track. The potential \tilde{P} is then corrected using \tilde{P}_{mod} along the minimal path that has just been found. As explained in the beginning of the section, it prevents two minimal paths which link two different pairs of partial tracks from following the same fluorescence trace. The two sources which have been joined are removed from the set of sources, and the minimal action map \mathcal{U} is updated by removing the two meeting fronts. The whole process (backpropagation on \mathcal{U} to get the minimal path,

³The *range of influence* of a source p_0 is the set of points of the solution domain which can be influenced by p_0 .

modification of \tilde{P} , updating of \mathcal{U} and propagation of fronts) is iterated until no shock between opposite fronts (i.e., backward and upward fronts) occurs anymore.

Solving the 3-D Eikonal Equation: Our approach requires to solve the Eikonal equation in a 3-D grid [32]. Classic finite difference schemes for the (3) tend to overshoot and are unstable. Rouy and Tourin [33] showed that the correct viscosity solution u for $\mathcal{U}_{x,y,z}$ is given by the following Godunov upwind difference scheme

$$\begin{aligned} & (\max\{u - \mathcal{U}_{x-1,y,z}, u - \mathcal{U}_{x+1,y,z}, 0\})^2 \\ & + (\max\{u - \mathcal{U}_{x,y-1,z}, u - \mathcal{U}_{x,y+1,z}, 0\})^2 \\ & + (\max\{u - \mathcal{U}_{x,y,z-1}, u - \mathcal{U}_{x,y,z+1}, 0\})^2 \\ & = \tilde{P}_{x,y,z}. \end{aligned} \quad (4)$$

The FMM is a single-pass algorithm solving the isotropic Eikonal equation in a consistent, accurate and efficient way [26], [29], [31], [32], [34]. In order to solve (3), grid points are selected in the order of increasing values of \mathcal{U} . The computational complexity is of order of $\mathcal{O}(N \log_2 N)$ for N grid points in the 3-D image, where the $\log_2 N$ factor comes from a heapsort algorithm.

We have privileged one of the variants of the FMM, called *group marching method* (GMM) [35]. Rather than sorting the solution in a narrow band to march forward a single grid point, the GMM advances a group of points in order to get a faster algorithm. To fix instability, at each iteration the points of the narrow band are updated twice, in two opposite orders. The heapsort algorithm is not required anymore, and the computational complexity drops to $\mathcal{O}(N)$. Note that computing a minimal path with the GMM is faster than the Dijkstra's approach [36], which is a $\mathcal{O}(N \log_2 N)$ method to obtain the shortest paths in a graph, and does not suffer from the metrication error inherent to graph algorithms on a discrete grid [29]. The original GMM is detailed in Table II.

Constraints on Front Propagation: Some improvements must be added to the original GMM to force each front to propagate upward or downward into a paraboloidic domain. For every grid point p at which \mathcal{U} has been computed, we have to know the source from which p depends. It is easy to keep track of the source during the front propagation by generating a source index map \mathcal{S} . When $\mathcal{U}(p)$ is updated with the (4), we simply set $\mathcal{S}(p) = k$, where k is the source index of the 6-connexity neighbor of p which is *Alive* and has the minimal value of \mathcal{U} . Then, to be given a \mathcal{U} value, each grid point must satisfy the *paraboloid constraint*, which is detailed in Table III.

To check if a downward front and an upward front meet at a given iteration of the GMM loop, we have to find a pair of neighbors that have different source indexes such that at least one of the two points was moved from the *Trial* to the *Alive* set at the previous iteration. Then, a backpropagation on \mathcal{U} from the meeting point gives the minimal path. The gradient descent is made with subpixel accuracy using a trilinear interpolation to estimate derivatives with a centred difference scheme. To force the z coordinate of the minimal path to take an integer value (corresponding to a frame index), we can either locally adapt the gradient descent step or reparameterize the minimal path using spline interpolation.

TABLE II
ALGORITHM FOR 3-D GROUP MARCHING

Definition:

- *Alive* set is the set of all grid points at which the action value \mathcal{U} has been computed and fixed.
- *Trial* set is the set of next grid points to be examined (i.e. points of the narrow band) for which an estimate of \mathcal{U} has been computed solving (4).
- *Far* set is the set of all other grid points, for which there is not yet an estimate for \mathcal{U} .

Initialization:

- *Alive* set is confined to the start point p_0 . Set $\mathcal{U}(p_0) = 0$.
- *Trial* set is confined to the 6-connexity neighbors p of p_0 . Set $\mathcal{U}(p) = \tilde{P}(p)$.
- *Far* is the set of all other grid points q . Set $\mathcal{U}(q) = \infty$.
- Set $\mathcal{U}_{\min} = \min(\tilde{P})$ and $\Delta_{\mathcal{U}} = \mathcal{U}_{\min}/\sqrt{3}$.

Loop:

While the *Trial* set is not empty, do:

- Set $\mathcal{U}_{\min} = \mathcal{U}_{\min} + \Delta_{\mathcal{U}}$.
- Scan the *Trial* set in the reverse order. For each *Trial* point p satisfying $\mathcal{U}(p) \leq \mathcal{U}_{\min}$:
 - * For each 6-connexity neighbor q of p which is not *Alive*, update $\mathcal{U}(q)$ solving (4).
- Scan the *Trial* set in the forward order. For each *Trial* point p satisfying $\mathcal{U}(p) \leq \mathcal{U}_{\min}$:
 - * For each 6-connexity neighbor q of p which is not *Alive*:
 - Update $\mathcal{U}(q)$ solving (4).
 - If q is *Far*, add it to the *Trial* set.
 - * Move p from the *Trial* to the *Alive* set.

TABLE III
PARABOLOID CONSTRAINT IN 3-D GROUP MARCHING

Let s be the source index of the 6-connexity neighbor of (x, y, z) which is *Alive* and has the minimal \mathcal{U} value. Let (x_s, y_s, z_s) be the coordinates of the source of index s , and D_s the diffusion coefficient associated to this source. When (x_s, y_s, z_s) is the *tail* of a partial track, (x, y, z) satisfies the paraboloid constraint if the two following conditions are satisfied:

$$\begin{cases} \sqrt{(x - x_s)^2 + (y - y_s)^2} & \leq c_{\Psi} \sqrt{D_s |z - z_s|} \\ z & < z_s \end{cases}$$

When (x_s, y_s, z_s) is the *head* of a partial track, (x, y, z) satisfies the paraboloid constraint if the two following conditions are satisfied:

$$\begin{cases} \sqrt{(x - x_s)^2 + (y - y_s)^2} & \leq c_{\Psi} \sqrt{D_s |z - z_s|} \\ z & > z_s \end{cases}$$

Illustration: Fig. 4 shows a representative example of our grouping process. This approach enables the detection and completion of trajectories of moving particles, even where the fluorescence signal is missing or too low to have been detected.

IV. GETTING TRUE MOLECULAR TRAJECTORIES

The algorithm performs an additional stage to analyze the completed tracks. There are two main reasons to know whether each point of a track corresponds to a true fluorescent probe or not, and, if so, to get a precise estimate of the particle position: i) although the detection and association stages give the position of fluorescent particles with a subpixel accuracy, a much better

precision can be achieved. A fluorescent particle can in fact be localized with an arbitrarily high accuracy if a sufficient number of photons has been collected [23]. ii) During the completion of partial tracks, the particle positions are extrapolated based on the properties of the minimal paths. They do not necessarily correspond to true positions of the probes, in particular during phases when the QD emission is completely off and not only below the detection threshold.

A set of potential particle positions is deduced from the set of completed tracks. If a fluorescence signal truly exists at position (x_m, y_m, z_m) , an accurate estimate of the position (x_0, y_0) of the probe in the 2-D image I_z can be obtained by fitting the experimental data in the vicinity of (x_m, y_m) with the following gaussian model:

$$g(x, y) = B + A \exp\left(-\frac{(x - x_0)^2 + (y - y_0)^2}{2s^2}\right) \quad (5)$$

where (A, B, s, x_0, y_0) are the model parameters. Therefore, we apply around each potential particle position (x_m, y_m) a nonlinear gaussian least-squares fit in a neighborhood of size $M_x \times M_y$. One decides that the fluorescent probe truly exists if the fit has converged. For the experimental results presented in Section VI, an isolated QD can, thus, be localized with an accuracy of 0.025–0.050 pixel (5–10 nm). Furthermore, this approach allows the discrimination of several fluorescent particles separated by a distance smaller than the resolution limit, using a mixture of gaussian models rather than (5) in the fit process. In this case, the localization accuracy depends on the true point-to-point distance, but is still below the resolution limit [14].

V. VALIDATION OF THE ALGORITHM

In order to evaluate the efficiency of our tracking algorithm and validate its use, we first tested it on sequences of synthetic images. We describe here the method which we used to simulate the diffusive motion of blinking objects and the quantitative results that we obtained.

A. Model Used to Generate Synthetic Data

Sequences of synthetic images were obtained by simulating the motion of N_T punctual objects. Each particle was represented as a gaussian spot with a fixed width σ_{PSF} and a mean amplitude A . A noise-free image was obtained by integrating in each pixel the signal of all the gaussian spots. In accordance with our fluorescence image model (see Section II-A), a constant baseline offset B_L , a gaussian background noise (with standard deviation σ_N), and a Poisson-distributed photon noise were added to the intensity value of each pixel of each image. Between frames, the position of the center of each spot evolved according to an isotropic Brownian motion (with diffusion coefficient D). To avoid boundary problems, each image in the sequence had a size 80×80 pixels and was considered as a subimage within a larger frame with size 120×120 pixels in which the particles were allowed to diffuse. Each time one of them escaped by diffusion from the larger frame, a new object was created with a random position. In addition, we assumed that the particles were not interacting and had all the same diffusion coefficient D in a given sequence.

TABLE IV
VALIDATION RESULTS: DETECTION RATE R_d

SNR	5			10			15				
	N_q	10	20	30	10	20	30	10	20		
	.999	.991	.996	1	1	.997	1	1	.965	10^{-3}	0
	.999	.997	.997	1	.998	.998	.999	1	.999	10^{-2}	
	.997	.990	.994	.998	.998	.993	.997	.998	.996	10^{-1}	
	.999	.998	.997	1	.998	1	.999	1	1	10^{-3}	0.1
	.981	.998	.996	.999	.999	.997	.998	.999	.999	10^{-2}	
	.989	.987	.993	.996	.995	.996	.997	.996	.994	10^{-1}	
	.999	.999	.994	1	.999	.999	.999	.999	.999	10^{-3}	0.3
	.995	.991	.996	.999	.992	.999	.999	.999	.998	10^{-2}	
	.997	.992	.985	.984	.996	.997	.996	.997	.994	10^{-1}	
										D	f_{off}

We used a standard method to describe the fluorescence intermittency. Each probe was modeled as a two-level system randomly alternating between an “On” ($A > 0$) and an “Off” state ($A = 0$). The dynamics of the system was governed by k_{on} (resp. k_{off}) the rate at which it switched from the “Off” (resp. “On”) to the “On” (resp. “Off”) state. In this simple Poisson process, the durations of the “On” and “Off” periods follow exponential distributions. At a given time, each probe has a probability $f_{\text{off}} = k_{\text{off}}/(k_{\text{on}} + k_{\text{off}})$, which we term the *extinction parameter*, to be in an “Off” state.

In each sequence, the initial conditions were set by randomly selecting for each spot its position within the 120×120 pixel image and determining its fluorescence state (“On” or “Off”) by means of a binomial distribution with parameter f_{off} .

In our simulations, we chose the values of the different parameters in order to remain as close as possible to the experimental conditions. In the sequence of synthetic images, the pixel size p_w and the width σ_{PSF} took the values given in Section II-A. The temporal sampling period is t_s and the gaussian noise standard deviation σ_N were, respectively, equal to 75 ms and 5, their values in the experiments. We tested our algorithm by simulating sequences with the following values.

- SNR = $A/\sqrt{A + \sigma_N^2} = 5, 10$ or 15. The last value is a minimum for what is usually obtained in single-QD imaging.
- Mean number of punctual objects in the view field: $N_q = (80^2/120^2)N_T = 10, 20$ or 30, corresponding to the different densities of immunolabeling.
- Diffusion coefficient: $D = 10^{-1}, 10^{-2}$ or $10^{-3} \mu\text{m}^2/\text{s}$ ($1 \mu\text{m}^2/\text{s} = 1.59 \text{ pixel}^2/\text{frame}$). These values correspond to slow, intermediate and rapid diffusion in the plasma membrane [9], [10].
- Extinction parameter: $f_{\text{off}} = 0, 0.1$ or 0.3 (with a fixed value $k_{\text{on}} = 0.05 \text{ frame}^{-1}$), varying between no intermittency and a significant fraction of the time (30%) in the “Off” state.

For each set of the four parameters listed above, six sequences, each composed of 100 images, were simulated and analyzed.

B. Results

The whole tracking algorithm has been applied to the sequences of synthetic images. The values given to the few parameters required by the algorithm are the following:

- detection threshold: $d_T = 3$;
- initial diffusion coefficient: $D_{\text{ini}} = 1 \mu\text{m}^2/\text{s}$;

- probability of allowed displacements: $\Psi = 0.95$;
- regularization term: $w = 1$.

Both detection and tracking performances were quantitatively evaluated. The efficiency with which fluorescent spots are detected was evaluated using a *detection rate* R_d defined as the mean of the ratio (N_c/N_v), where N_v is the number of particles which are visible (in the “On” state) in a given frame and N_c is the number of points of reconstructed trajectories which belong to that frame. To analyze the performance of tracking, we use the *track-based error* E_t , introduced in [37], and defined as $E_t = 1 - (M_c/M_v)$, where M_v is the total number of reconstructed trajectories and M_c the number of reconstructed trajectories for which the start and end points correspond to the same true trajectory.

The results for the detection rate R_d and the track-based error E_t (mean and standard deviation) obtained from the analysis of synthetic data are summarized in Tables IV and V, respectively. The detection rate is very good, even when the SNR is low and the density of spots is high, showing the robustness and ability of our algorithm for detecting fluorescent spots in realistic experimental conditions.

The track-based error indicates that our algorithm is able to reconstruct individual molecular trajectories efficiently, even when the tracked objects vanish from the field of view for quite long periods. The algorithm begins to fail when either the number of spots is too high or the duration of “Off” states is too long. It can also fail, but to a lesser extent, when spots are very close. Note that solving the motion correspondence problem when the trajectories of two diffusive particles cross each other is particularly challenging, since no information about the direction of motion can be used to retrieve the trajectory of each particle.

Running time highly depends on the number of spots and the extinction probability, and varies from tens of seconds to 3 min per sequence (on a 1.8-GHz Athlon processor). Solving the Eikonal equation in a 3-D grid is undoubtedly the main computational expense of our tracking approach. So, the algorithm complexity is of the same order than the GMM.

VI. EXPERIMENTAL RESULTS

In this section, we apply our tracking algorithm to data obtained in experiments on the membrane dynamics of single-glycine receptors in live cultured neurons. The motivations for these experiments, as well as the methods and protocols which

TABLE V
VALIDATION RESULTS: TRACK-BASED ERROR E_t

SNR	5			10			15						
	N_q	10	20	30	10	20	30	10	20			30	
		.076 (.088)	.150 (.050)	.164 (.057)	.000 (.000)	.000 (.000)	.062 (.038)	.014 (.034)	.050 (.077)	.110 (.013)	10^{-3}	0	
		.107 (.071)	.130 (.066)	.195 (.061)	.054 (.085)	.116 (.035)	.136 (.098)	.075 (.086)	.090 (.070)	.090 (.041)	10^{-2}		
		.128 (.059)	.233 (.059)	.281 (.049)	.154 (.071)	.200 (.055)	.255 (.029)	.136 (.092)	.174 (.090)	.260 (.051)	10^{-1}		
		.133 (.082)	.172 (.071)	.136 (.023)	.050 (.041)	.059 (.047)	.091 (.065)	.013 (.031)	.050 (.033)	.027 (.031)	10^{-3}	0.1	
		.126 (.092)	.231 (.079)	.190 (.040)	.094 (.056)	.115 (.041)	.137 (.049)	.088 (.058)	.023 (.038)	.107 (.064)	10^{-2}		
		.168 (.091)	.276 (.021)	.371 (.060)	.220 (.075)	.222 (.069)	.238 (.063)	.093 (.120)	.275 (.088)	.378 (.077)	10^{-1}		
		.172 (.121)	.376 (.118)	.244 (.066)	.047 (.075)	.110 (.061)	.167 (.100)	.057 (.092)	.215 (.107)	.192 (.055)	10^{-3}	0.3	
		.139 (.129)	.275 (.080)	.299 (.072)	.086 (.095)	.178 (.044)	.270 (.150)	.055 (.053)	.151 (.029)	.213 (.110)	10^{-2}		
		.270 (.142)	.319 (.048)	.354 (.056)	.207 (.152)	.222 (.077)	.427 (.068)	.163 (.151)	.354 (.063)	.359 (.127)	10^{-1}		
												D	f_{off}

have been used, are discussed in more details in [10]. In brief, α_1 subunits of glycine receptors were sequentially labeled with primary antibodies (mab2b, Alexis Chemicals) targeting an extracellular epitope, biotinylated anti-mouse Fab fragments and streptavidin-coated QDs (QD605, QDots corp). The concentration of the reagents and the corresponding incubation times were adjusted such that very few receptors are labeled and individual QDs can be easily distinguished. The low surface density ensures that a single receptor is labeled by at most one QD. However, we can not exclude that a single QD is bound to more than one receptor, possibly due to the divalence of the primary antibody. Note that the methods used to label glycine receptors are fairly standard and can be easily adapted to track the motion of any transmembrane proteins having an extracellular epitope. Data were taken at room temperature with an epifluorescence inverted microscope (Olympus IX70) equipped with an oil-immersion objective ($60\times$, $N_A = 1.45$). Images were recorded with a low-noise back-illuminated CCD Camera (MicroMax EBFT512, Roper Scientific). The QDs were excited by a mercury lamp (excitation filter 525DF45, dichroic filter 560DRLP, Omega Filters) and their fluorescence detected with a band emission filter (595DF60 or 605WB20, Omega Filters).

Sequences of fluorescence images were acquired for glycine receptors diffusing in the dendritic membrane. In this part of the neuronal cell, far from the cell body, the autofluorescence is very low and the detection is not limited by background noise. In addition, we recorded the motion of receptors on dendrites with a thickness smaller than the axial resolution (about $1\ \mu\text{m}$) so that the diffusion of the tagged proteins along the optical axis did not lead to defocusing of the fluorescence spots. In these experiments, the difficulty in tracking the molecules originates essentially in the intensity fluctuations and the fluorescence intermittency of the probes.

The tracking algorithm has been applied to several sequences constituted of hundreds frames. These sequences and tracking results are available in our website.⁴ The values of the algorithm parameters were similar to the ones used for the validation (see Section V), except for the detection threshold d_T which was equal to 8.

Figs. 5 and 6 present some results obtained with small image stacks for a better clarity. A trajectory is displayed by linking




Fig. 5. Illustration of the proposed approach for the tracking of individual glycine receptors in the membrane of live neurons. The tracking method has been applied to a small stack of size $N_x \times N_y \times N_z$, where $N_x = 76$ pixels, $N_y = 37$ pixels, and $N_z = 100$ frames. (a) First frame of the fluorescence image stack, saturated for clarity. (b) Sectioning view of the image stack resulting from a maximal intensity projection (MIP) in the yz plane. Trajectories are surimposed to the fluorescence trace of QDs in motion. (c) Sectioning view of the image stack resulting from a MIP in the xz plane. Trajectories are surimposed to the fluorescence trace of QDs in motion. (d) Three-dimensional representation of the trajectories in a spatiotemporal volume.

points for which the gaussian fit has converged. Therefore, a segment joining points with distant z coordinates corresponds to an extinction of the QD (see Fig. 6). This illustrates the ability of our algorithm to track fluorescent probes which transiently disappear.

VII. CONCLUSION

In this paper, we present a novel approach for tracking point objects using minimal paths in a spatiotemporal volume. The algorithm is based on a relevant model of fluorescence image and molecular dynamics in membrane and requires few parameters. Solving the multiframe motion correspondence problem is particularly complex since the number of solutions to be examined is subject to combinatorial explosion. In order to overcome this difficult challenge in a reasonable computation time, we propose to use a greedy algorithm which is locally (but not globally) optimal. The robustness and efficiency of our perceptual

⁴Available: www.lkb.ens.fr/recherche/optetbio/image/index.htm.




Fig. 6. Three-dimensional representation in a spatiotemporal volume of glycine receptors trajectories. These trajectories are issued from the tracking of QDs in a stack of size $N_x \times N_y \times N_z$, where $N_x = 113$ pixels, $N_y = 97$ pixels, and $N_z = 300$ frames. A trajectory is displayed by linking points for which the gaussian fit has converged. The motion of individual proteins in the membrane is clearly observable over several tens of seconds, with a spatial resolution of a few nanometers. Despite the blinking of QDs, which is characterized by the sudden interruption of some trajectories (as the one indicated by the arrow), our approach enables an automatic reconstruction of each individual motion.

grouping approach has been illustrated by analyzing synthetic and real-world data.

Future studies will include a systematic comparison of results obtained with our method and other algorithms, both with synthetic and experimental data. We will also apply our method to tracking experiments for which the SNR is lower and the background noise becomes a limiting factor.

Although more specially developed in the context of QD tracking, our approach can, however, be applied with other fluorescent probes and should be useful for many experiments in the growing field of SMT experiments.

APPENDIX

ESTIMATION OF THE DIFFUSION COEFFICIENT

Consider a trajectory defined by the N_p successive positions $r_j = (x_j, y_j)$ of a particle for $j \in \{0, \dots, N_p - 1\}$. The second moment of the probability distribution of displacements $\mathcal{P}(r, \tau, D)$ is the mean square displacement (MSD) of this particle, denoted by $\rho(\tau)$

$$\rho(\tau) = \int_0^\infty \mathcal{P}(r, \tau, D) 2\pi r dr = 4D\tau.$$

This MSD can be estimated with a discrete time sequence

$$\hat{\rho}(n t_s) = \frac{1}{N_p - n} \sum_{j=0}^{N_p - n} \|r_{j+n} - r_j\|^2$$

where t_s is the temporal sampling period and $n \in \{0, \dots, N_p - 2\}$. Consequently, a least-squares fit to $\hat{\rho}$ gives an estimate of the diffusion coefficient D [9], [38]. Note that the accuracy of

this estimation can be improved using a weighted least-squares fit [39]. If a dynamical change can occur, one have better to estimate an instantaneous diffusion coefficient using a windowed version of $\hat{\rho}$. For instance, in the algorithm given in Table I, the diffusion coefficient associated to the tail of a partial track is obtained by fitting the MSD computed only with the first points of the partial track.

ACKNOWLEDGMENT

The authors would like to thank M.-V. Ehrensperger, S. Lévi, and A. Triller for providing them with experimental data and for fruitful discussions.

REFERENCES

- [1] S. M. Hurlley and L. Helmut, "Special issue on biological imaging," *Science*, vol. 300, pp. 1–196, 2003.
- [2] Y. Shav-Tal, R. H. Singer, and X. Darzacq, "Imaging gene expression in single living cells," *Nature Rev. Mol. Cell Biol.*, vol. 5, pp. 855–862, 2004.
- [3] J. Lippincott-Schwartz and G. H. Patterson, "Development and use of fluorescent protein markers in living cells," *Science*, vol. 300, pp. 87–91, 2004.
- [4] J. Zhang, R. E. Campbell, A. Y. Ting, and R. Y. Tsien, "Creating new fluorescent probes for cell biology," *Nature Rev. Mol. Cell Biol.*, vol. 3, pp. 906–918, 2002.
- [5] A. P. Alivisatos, "The use of nanocrystals in biological detection," *Nature Biotechnol.*, vol. 22, pp. 47–52, 2004.
- [6] J. K. Jaiswal and S. M. Simon, "Potentials and pitfalls of fluorescent quantum dots for biological imaging," *Trends Cell Biol.*, vol. 14, pp. 497–504, 2004.
- [7] X. Michalet, F. F. Pinaud, L. A. Bentolila, J. M. Tsay, S. Doose, J. J. Li, G. Sundaresan, A. M. Wu, S. S. Gambhir, and S. Weiss, "Quantum dots for live cells, in vivo imaging, and diagnostics," *Science*, vol. 307, pp. 538–544, 2005.
- [8] S. Kim, Y. T. Lim, E. Soltesz, A. M. D. Grand, J. Lee, A. Nakayama, J. A. Parker, T. Mihaljevic, R. G. Laurence, D. M. Dor, L. H. Cohn, M. G. Bawendi, and J. V. Frangioni, "Near-infrared fluorescent type II quantum dots for sentinel lymph node mapping," *Nature Biotechnol.*, vol. 22, pp. 93–97, 2004.
- [9] M. J. Saxton and K. Jacobson, "Single particle tracking: applications to membrane dynamics," *Annu. Rev. Biophys. Biomol. Struct.*, vol. 26, pp. 373–399, 1997.
- [10] M. Dahan, S. Lévi, C. Luccardini, P. Rostaing, B. Riveau, and A. Triller, "Diffusion dynamics of glycine receptors revealed by single quantum dot tracking," *Science*, vol. 302, pp. 442–445, 2003.
- [11] M. Kuno, D. P. Fromm, H. F. Hamann, A. Gallagher, and D. J. Nesbitt, "Nonexponential blinking kinetics of single CdSe quantum dots: a universal power law behavior," *J. Chem. Phys.*, vol. 112, pp. 3117–3120, 2000.
- [12] W. Tvarusko, M. Bentele, T. Misteli, R. Rudolf, C. Kaether, D. L. Spector, H. H. Gerdes, and R. Eils, "Time-resolved analysis and visualization of dynamic processes in living cells," *Proc. Nat. Acad. Sci. USA*, vol. 96, pp. 7950–7955, 1999.
- [13] A. Genovesio, B. Zhang, and J. Olivo-Marin, "Tracking of multiple fluorescent biological objects in three dimensional video microscopy," presented at the Int. Conf. Image Processing, Barcelona, Spain, Sep. 2003.
- [14] D. Thomann, D. R. Rines, P. K. Sorger, and G. Danuser, "Automatic fluorescent tag detection in 3-D with super-resolution: application to the analysis of chromosome movement," *J. Microsc.*, vol. 208, pp. 49–64, 2002.
- [15] S. Bonneau, L. Cohen, and M. Dahan, "A multiple target approach for single quantum dot tracking," presented at the IEEE Int. Symp. Biomedical Imaging, Arlington, VA, Apr. 2004.
- [16] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Automat. Control*, vol. AC-24, no. 6, pp. 843–856, Dec., 1979.
- [17] I. J. Cox, "A review of statistical data association techniques for motion correspondence," *Int. J. Comput. Vis.*, vol. 10, pp. 53–66, 1993.
- [18] 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 Anal. Mach. Intell.*, vol. 18, no. 1, pp. 138–150, Jan. 1996.

- [19] I. K. Sethi and R. Jain, "Finding trajectories of feature points in a monocular image sequence," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-9, no. 1, pp. 53–73, Jan. 1987.
- [20] D. Chetverikov and J. Verestoy, "Feature point tracking for incomplete trajectories," *Computing*, vol. 6, pp. 321–338, 1999.
- [21] C. J. Veenman, M. J. T. Reinders, and E. Backer, "Resolving motion correspondence for densely moving points," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 1, pp. 54–72, Jan. 2001.
- [22] L. D. Cohen, "Minimal paths and fast marching methods for image analysis," in *Mathematical Models in Computer Vision: The Handbook*, N. Paragios, Y. Chen, and O. Faugeras, Eds. New York: Springer, 2005.
- [23] R. E. Thompson, D. R. Larson, and W. W. Webb, "Precise nanometer localization analysis for individual fluorescent probes," *Biophys. J.*, vol. 82, pp. 2775–2783, 2002.
- [24] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc. B*, vol. 39, pp. 1–38, 1977.
- [25] L. D. Cohen, "Multiple contour finding and perceptual grouping using minimal paths," *J. Math. Imag. Vis.*, vol. 14, pp. 225–236, 2001.
- [26] L. D. Cohen and T. Deschamps, "Grouping connected components using minimal path techniques," presented at the *IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, Kauai, HI, Dec. 2001.
- [27] Q. Lin, "Enhancement, detection, and visualization of 3-D volume data," Ph.D. dissertation, Dept. Elect. Eng., Linköping Univ., Linköping, Sweden, 2003.
- [28] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: active contour models," *Int. J. Comput. Vis.*, vol. 1, pp. 321–331, 1988.
- [29] L. D. Cohen and R. Kimmel, "Global minimum for active contour models: a minimal path approach," *Int. J. Comput. Vis.*, vol. 24, pp. 57–78, 1997.
- [30] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. J. Comput. Vis.*, vol. 22, pp. 61–79, 1997.
- [31] J. A. Sethian, "Fast marching methods," *SIAM Rev.*, vol. 41, pp. 199–235, 1999.
- [32] T. Deschamps and L. D. Cohen, "Fast extraction of minimal paths in 3-D images and applications to virtual endoscopy," *Med. Image Anal.*, vol. 5, pp. 281–299, 2001.
- [33] E. Rouy and A. Tourin, "A viscosity solution approach to shape from shading," *SIAM J. Numer. Anal.*, vol. 29, pp. 867–884, 1992.
- [34] J. A. Sethian, *Level Set Methods and Fast Marching Methods*. Cambridge, U.K.: Cambridge Univ. Press, 1999.
- [35] S. Kim, "An $\mathcal{O}(N)$ level set method for Eikonal equations," *SIAM J. Sci. Comput.*, vol. 22, pp. 2178–2193, 2001.
- [36] E. W. Dijkstra, "A note on two problems in connection with graphs," *Numer. Math.*, vol. 1, pp. 269–271, 1959.
- [37] J. Verestoy and D. Chetverikov, "Experimental comparative evaluation of feature point tracking algorithms," in *Performance Characterization in Computer Vision*, R. Klette, H. Stiehl, M. Viergever, and K. Vincken, Eds. Norwell, MA: Kluwer, 2000.
- [38] H. Quian, M. P. Sheetz, and E. L. Elson, "Single particle tracking: analysis of diffusion and flow in two-dimensional systems," *Biophys. J.*, vol. 60, pp. 910–921, 1991.
- [39] M. J. Saxton, "Single particle tracking: the distribution of diffusion coefficients," *Biophys. J.*, vol. 72, pp. 1744–1753, 1997.



Stéphane Bonneau received the M.E. degree in computer science from the Université de Technologie de Compiègne, Compiègne, France, and the M.S. degree in applied mathematics from the Ecole Normale Supérieure de Cachan, Cachan, France, both in 2003. He is currently pursuing the Ph.D. degree in applied mathematics at CEREMADE, Université Paris-Dauphine, Paris, France, within the context of a collaboration with the Laboratoire Kastler Brossel, Ecole Normale Supérieure.

His research interests are in variational methods and their applications to image analysis.



Maxime Dahan received the engineering degree from the Ecole Polytechnique, Paris, France, in 1993 and the M.S. and Ph.D. degrees in quantum physics from the Ecole Normale Supérieure, Paris, and the Université Paris 6, in 1994 and 1997, respectively.

From 1997 to 1999, he was a Postdoctoral Fellow at Lawrence Berkeley National Laboratory, Berkeley, CA, working with Dr. S. Weiss and Pr. D. Chemla. Since 2000, he has been a Research Scientist with CNRS and has been responsible for the Optics and Biology Group in the Laboratoire Kastler Brossel,

Ecole Normale Supérieure. His current research interests are focused on the optical detection of single molecules and its application to the study of biological processes. He also has a long-standing interest in the physical properties of semiconductor quantum dots.



Laurent D. Cohen (SM'96) received the M.S. and Ph.D. degrees in applied mathematics from the Université Paris 6, Paris, France, in 1983 and 1986, respectively.

He was a student at the Ecole Normale Supérieure, Paris, from 1981 to 1985. From 1985 to 1987, he was member of the Computer Graphics and Image Processing Group at Schlumberger Palo Alto Research, Palo Alto, CA, and Schlumberger Montrouge Research, Montouge, France, and remained Consultant there for a few years afterward. He joined

INRIA, France, in 1988, where he mainly worked with the Medical Image Understanding Group EPIDAURE. Since 1990, he has been a Research Scholar (Chargé then Directeur de Recherche) with CNRS in the Applied Mathematics and Image Processing group at CEREMADE, Université Paris-Dauphine, Paris. He has been member of program committee for boards for about 25 international conferences. His research interests and teaching at the university are applications of variational methods and partial differential equations to image processing and computer vision, like deformable models, minimal paths, surface reconstruction, image registration, image segmentation, and restoration.

Dr. Cohen received the CS 2002 Prize for Image and Signal Processing.