A light-weight segmentation method for multi-state Brownian trajectories

Ismail El Korde,[1] Jason M. Lewis,[1] Erik Clarkson,[2] Tommy Dam,[1,3]
Peter Jönsson,[1] Tobias Ambjörnsson,[2] and Joakim Stenhammar[1,*]

[1] *Division of Physical Chemistry, Lund University, Sweden*
[2] *Centre for Environmental and Climate Science, Lund University, Sweden*
[3] *Sir William Dunn School of Pathology, University of Oxford, UK*

(Dated: June 5, 2025)

Single-particle tracking methods have emerged as a crucial tool for the characterization of dynamical and diffusive processes in a range of biological and synthetic systems. Here, we propose a simple and light-weight yet accurate method for the segmentation of multi-state Brownian trajectories based on an optimised Gaussian filtering of the displacement time series combined with an automated fitting to a Gaussian mixture model. We demonstrate our method on synthetic, 2-state Brownian trajectories and show that our method provides high levels of accuracy in terms of segmentation and the estimation of self-diffusion coefficients for reasonably well-separated values of the diffusion coefficients. Compared to methods based on deep learning or hidden Markov models our method imposes a much lower computational load, making it suitable for fast and accurate online processing of single-particle trajectories from microscopy images.

INTRODUCTION

The revolution in optical microscopy techniques over the last three decades have significantly enhanced the ability to investigate biological systems at the molecular level. Among these, single-molecule fluorescence microscopy [1–3] in conjunction with single-particle tracking (SPT) algorithms [4–10] has revolutionised the study of molecular dynamics in living cells by providing unprecedented spatiotemporal resolution. By enabling the tracking of individual molecules, these methods have yielded crucial insights into intracellular transport, protein-protein interactions, and the fundamental mechanisms that govern cellular processes [1, 7, 11–16], as well as fundamental knowledge about diffusive processes [17, 18].

A fundamental challenge in SPT analysis is the characterisation of multiple diffusive states based on single-particle trajectories, which can often exhibit significant heterogeneity due to complex interactions with the cellular environment [16, 19–22]. The ability to accurately analyse single-particle diffusive processes therefore provides a route to understanding intercellular processes such as protein-ligand binding and cellular signaling [23–25]. The traditional approach for analysing diffusive processes in homogeneous systems involves the calculation of the mean-square displacement (MSD), which, for molecules undergoing simple Brownian motion characterised by a single diffusion coefficient, serves as a straightforward method for inferring molecular transport properties [26–28]. For biological systems with heterogeneous diffusion behavior, wherein molecules transition between multiple diffusive states due to transient interactions with cellular structures, MSD-based analysis methods are however less straightforward, since a temporally and spatially resolved sampling of the MSD is very statistically challenging [20, 29].

To address this limitation, more sophisticated methods have been developed to achieve parameter estimation and segmentation of multi-state diffusive trajectories. One widely used class of methods is based on so-called *hidden Markov models* (HMMs), whereby trajectories are segmented based on probabilistic modelling to find the maximum-likelihood (or Bayesian) estimate of the diffusion coefficients $\boldsymbol{D_i}$ and transition probabilities p_{ij} of the different states [30–39]. Traditionally, HMM models require as input a predefined number of diffusive states [32], although more recent HMM-based methods overcome this requirement, instead inferring the most probable number of diffusive states directly from the data [30, 33–35, 38]. Although significantly more accurate than MSD-based methods, inferring diffusive parameters using HMM methods can become computationally demanding for complex systems with many diffusive states. In addition to these statistical methods, several well-performing deterministic methods for change-point detection between diffusive states have been developed [40–44], together forming a large toolbox of segmentation algorithms for multi-state particle trajectories.

More recently, methods based on machine learning (ML) [45, 46] and, more specifically, deep learning [47–53] have emerged as a versatile and accurate tool to handle the analysis of single-molecule dynamics, including extraction of trajectories from (potentially noisy) microscopy images, segmentation of multi-state Brownian trajectories, and the analysis of "anomalous", non-Brownian diffusion processes, where the MSD grows super- or sublinearly with time [47, 48, 50, 51]. While these data-driven approaches often demonstrate impressive accuracy in analysing complex diffusion behaviors, they are typically computationally heavy and require large training datasets. Furthermore, they typically lack physical interpretability due to the complex nature of the ML architecture, making it difficult for the user to assess the quality of the parameter extraction.

In this study, we propose an alternative, simple and light-weight approach for the segmentation of multi-state Brownian trajectories. Our method uses a simple, one-dimensional Gaussian filter that acts on the displacement time series to achieve efficient noise reduction. The resulting displacement distribution is fitted to a Gaussian mixture model in order to adjust the filter width in an automated manner to optimally classify the different diffusive states. In spite of its relative simplicity and low computational cost, we show that, for reasonably well-separated and long-lived diffusive states, the methodology is able to achieve accurate segmentation while not relying on the large amounts of (synthetic or experimental) training data required for ML-based methods. Compared to these methods, our approach is furthermore physically transparent, in that the automated fitting and filter optimisation process can be monitored to evaluate the quality of the segmentation. In contrast to HMM-based approaches, segmentation of the trajectory is done *before* parameter estimation, meaning that diffusion constants and lifetimes can be estimated using the well-established methods for single-state trajectories, as we will demonstrate below. In summary, our method adds to the growing toolbox of multi-state SPT algorithms, providing a simple and light-weight tool for the segmentation of single-particle trajectories and the estimation of self-diffusion coefficients and state lifetimes.

MODEL AND METHODS

Throughout this paper, we will consider the trajectory of a single Brownian particle which can adopt a number of discrete diffusive states i, each characterised by a self-diffusion coefficient D_i. The different states can, for example, model the binding and unbinding of a ligand to a receptor or the clustering between two protein molecules in a lipid bilayer, leading to a sudden decrease in diffusivity. The transition between states is described as a Poisson process with transition rates λ_{ij}, representing the probability per unit time of transitioning from state i to state j. For simplicity of presentation, our analysis will be restricted to the case of 2-state Brownian diffusion in 2 spatial dimensions, but the method can be generalised to higher-dimensional trajectories and/or higher numbers of states in a straightforward manner.

Data generation

Synthetic, 2-state Brownian trajectories are generated by solving the discretised, overdamped Langevin equation using Euler's method:

$$\boldsymbol{r}(t_{\alpha+1}) = \boldsymbol{r}(t_\alpha) + \sqrt{2D_i\delta t}\boldsymbol{\omega}(t_\alpha), \tag{1}$$

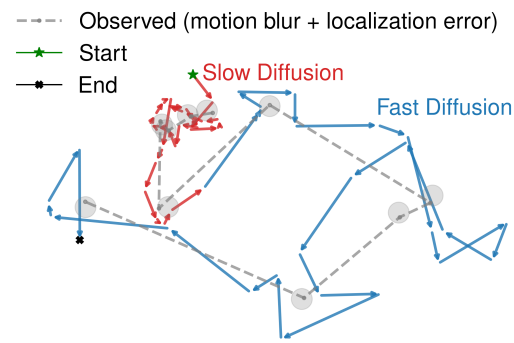

FIG. 1. Example of a 2-state Brownian trajectory: The true trajectory is represented by color-coded arrows, while the observed trajectory, shown as a dashed gray line with circular markers, is affected by motion blur ($n = 5$), where the path is averaged over non-overlapping time intervals. Localisation error is simulated by adding Gaussian noise to the blurred positions, represented by circles around each point, reflecting the uncertainty in position measurement.

where $t_\alpha = \alpha\delta t$ is the discretised time, δt is the time step used to generate the trajectory (not to be confused with the microscope acquisition time, Δt, which we introduce below), and $\boldsymbol{\omega}$ is a vector with the dimension of the system where each component is drawn from a normal distribution with zero mean and unit variance.

At each time step the particle changes its diffusion coefficient from D_1 (resp. D_2) to D_2 (resp. D_1) with a predefined transition rate λ_{12} (resp. λ_{21}) with the transition probability [31]

$$p_{12} = \frac{\lambda_{12}}{\lambda_{12}+\lambda_{21}}\left[1 - \exp\Big(-(\lambda_{12}+\lambda_{21})\delta t\Big)\right], \tag{2}$$

where $p_{11} = 1 - p_{12}$ and p_{21} is obtained by interchanging the indices 1 and 2 in Eq. (2). The average lifetimes are thus given by $\bar{\tau}_1 = 1/\lambda_{12}$ and $\bar{\tau}_2 = 1/\lambda_{21}$.

To mimic the blur induced by the camera's shutter count, following Ref. [54], we introduce the acquisition time $\Delta t \equiv n\delta t$, where the integer n is the number of discrete time steps over which the image is acquired. We assume constant illumination throughout the acquisition time Δt, meaning that all positions contribute equally to the average position measured over the time window. To mimic this process, we first generate a Brownian trajectory of length T by solving Eq. (1), which we then divide into T/n non-overlapping segments of length n, separated by Δt. The blurred trajectory is constructed by taking the arithmetic mean of the particle's position within each segment. In addition, we mimic the errors introduced by limited spatial resolution and the subsequent tracking process by adding a Brownian noise $\boldsymbol{\xi}$ of zero mean and variance σ^2 [37]. Taken together, the noisy trajectory $\bar{\boldsymbol{r}}$ is thus obtained as

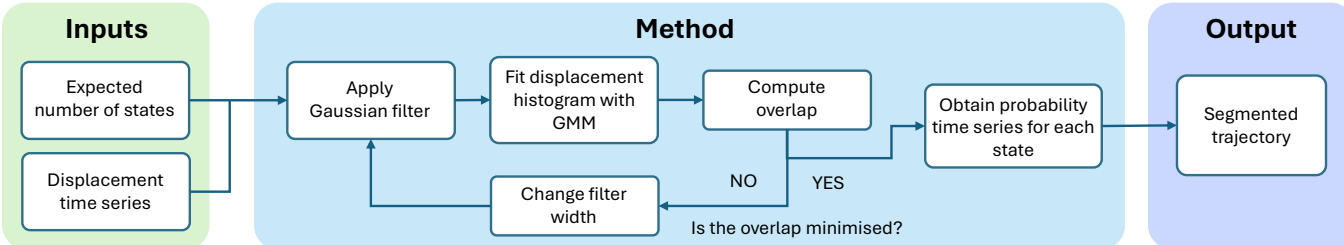

FIG. 2. Schematic of the different steps in the method: Given a set of trajectories and an expected number of states, the overlap between the displacement distributions is minimised using an optimised Gaussian filter acting on the displacement time series. The output consists of a segmentation of the given trajectory, inferring the current state of the particle at each given time point.

$$\bar{\boldsymbol{r}}(\bar{t}_\alpha) = \frac{1}{n}\sum_{m=0}^{n-1} \boldsymbol{r}(\bar{t}_\alpha + m\delta t) + \boldsymbol{\xi}(\bar{t}_\alpha), \tag{3}$$

where $\bar{t}_\alpha \equiv \alpha\Delta t, \alpha = 1, \ldots, T$ now represents time measured in units of the camera acquisition time. This procedure introduces both temporal and spatial correlations into the trajectory, which must be considered when analysing the particle's dynamics, such as determining the diffusion constant [54] or transition probabilities. An example of such a trajectory is presented in Fig. 1, illustrating the motion before and after the incorporation of motion blur and localization error. In the following, we will use Δt as our basic time unit when presenting our data, unless stated otherwise.

Processing & Classification

The primary goal of the method is to infer the state of the particle at each time step based on its positional time series and the expected number of states. As schematically depicted in Fig. 2, we start by computing the displacement time series from the trajectory $\bar{\boldsymbol{r}}(t)$, as generated from Eq. (3), by computing the Euclidean norm of the distance Δr between two successive positions of the particle:

$$\Delta r(\bar{t}_{\alpha+1}) = |\bar{\boldsymbol{r}}(\bar{t}_{\alpha+1}) - \bar{\boldsymbol{r}}(\bar{t}_\alpha)|. \tag{4}$$

Fig. 3a shows an example of such a noisy time series, clearly showing a large spread across the two states. Since the displacements Δx and Δy in each dimension are normally distributed with mean $\mu_i = 0$ and variance $\sigma_i^2 \propto 2D_i$, the displacement probability distribution $P_i(\Delta r)$ in state i of $\Delta r = \sqrt{\Delta x^2 + \Delta y^2}$ is Rayleigh distributed, *i.e.*

$$P_i(\Delta r) = \frac{\Delta r}{\sigma_i^2} \exp\left(-\frac{\Delta r^2}{2\sigma_i^2}\right), \tag{5}$$

yielding a 2-state displacement distribution given by

$$P(\Delta r) = \pi_1 P_1(\Delta r) + \pi_2 P_2(\Delta r), \tag{6}$$

with the weights $\pi_1 = \lambda_{21}/(\lambda_{12}+\lambda_{21})$ and $\pi_2 = \lambda_{12}/(\lambda_{12}+\lambda_{21})$ such that $\pi_1 + \pi_2 = 1$, as shown in Fig. 3b.

From Figs. 3a and b, we observe that, although large displacements are more frequent in the state with the higher diffusion coefficient, small displacements are frequent in both states, leading to considerable overlap between the two probability densities. D_1 and D_2 can then be obtained by fitting the Rayleigh mixture model (RMM) in Eq. (6) to the measured distribution [16, 35]. However, the large overlap between the two probability densities makes a direct inference of the state at a given time point from the raw data challenging, and therefore an RMM cannot be used directly to detect transitions or infer values of $\bar{\tau}_i$. To reduce this overlap and enable an explicit segmentation of the trajectory into a time series of the state of a given particle, we convolute the displacement time series $\Delta r(t)$ with a discrete Gaussian filter β of width f to obtain the *filtered* displacement time series $\widehat{\Delta r}(t; f)$:

$$\widehat{\Delta r}(\bar{t}_\alpha; f) = \frac{\sum_{m=0}^{T} \Delta r(\bar{t}_m)\beta(\bar{t}_\alpha - \bar{t}_m; f)}{\sum_{m=0}^{T} \beta(\bar{t}_\alpha - \bar{t}_m; f)}, \tag{7}$$

where $\beta(t; f) = \exp[-t^2/(2f^2)]$. Here, we also tacitly assume that all state lifetimes $\bar{\tau}_i$ are longer than, or at least comparable to, the filter width f to avoid significant blurring of the state transitions.

For $f \ll \Delta t$, β approaches a delta function, leading to a filtered displacement time series identical to the unfiltered one. Conversely, for $f \gg \Delta t$, β yields an effectively constant displacement time series. For intermediate values, the temporal correlations introduced by the filtering process will reduce the standard deviation of $P_i(\Delta r)$, leading to a decrease in the overlap between the individual probability distributions P_1 and P_2. This suggests

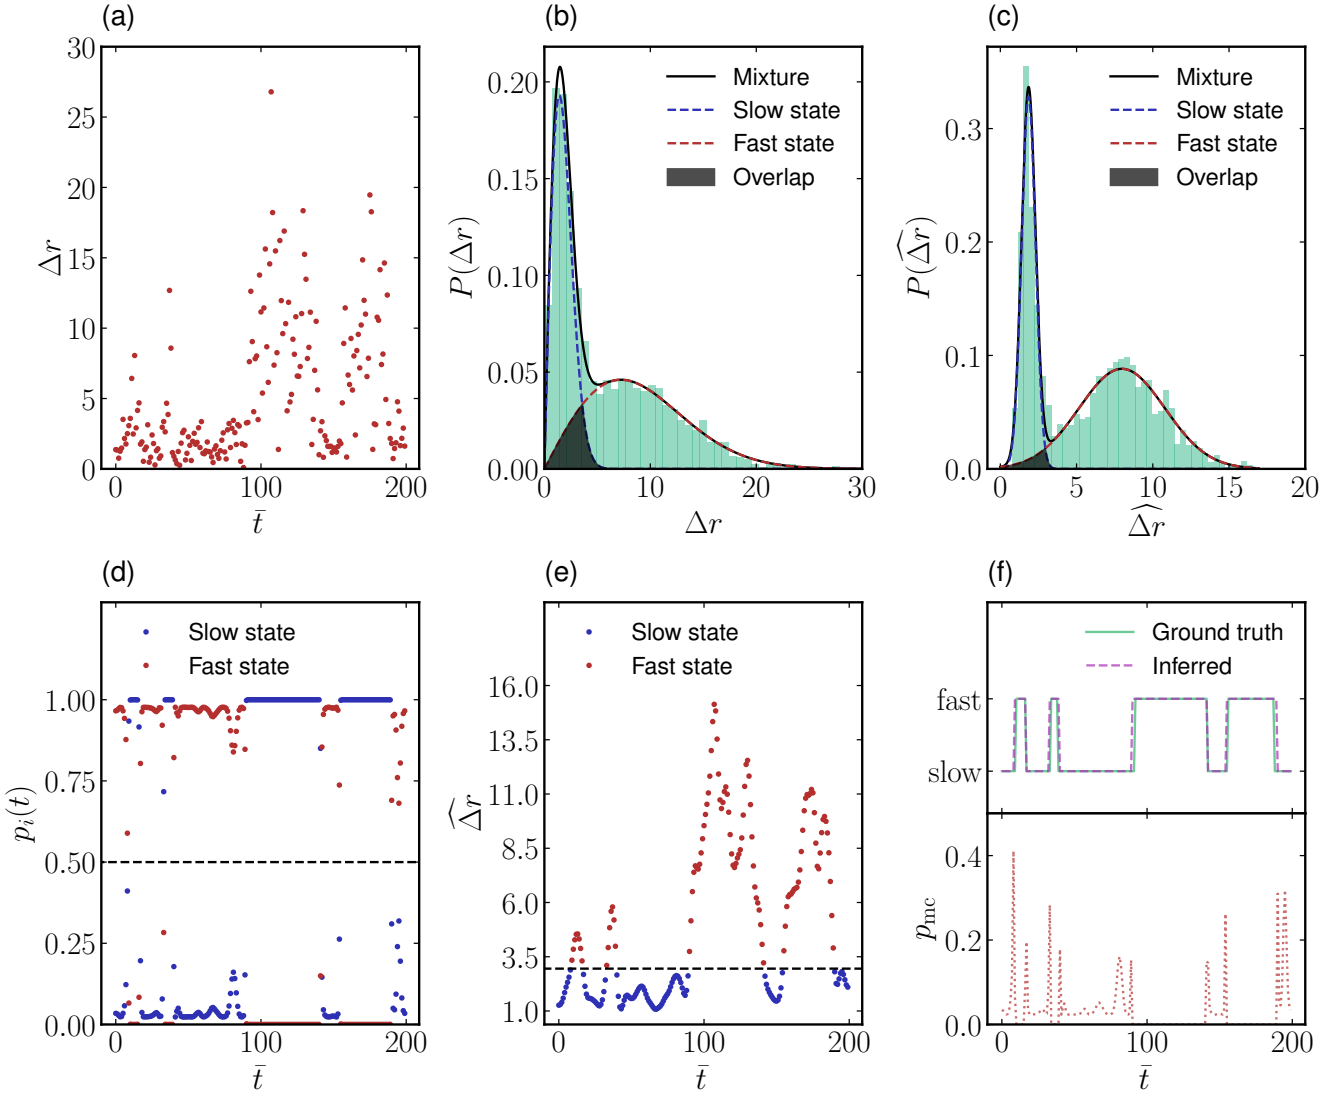

FIG. 3. (a) Displacement time series of a 2-state Brownian particle, with $\bar{\tau}_1 = \bar{\tau}_2 = 20$, and $D_2/D_1 = 5$. (b) Displacement distribution of the N_t unfiltered trajectories fitted to a Rayleigh mixture according to Eq. (6). The overlap between the displacement distributions of individual states is shaded. (c) Displacement distribution of the same N_t trajectories, but now after filtering using the optimal filter width f^* fitted to a sum of two Gaussians, according to Eq. (8). (d) Inferred time series of the probability p_i of belonging to the diffusive state i, extracted from the filtered displacement distribution. The black dashed line indicates the threshold, corresponding to $p_i = 0.5$. (e) The displacement time series in (a) after filtering using the optimal filter width f^*. The dashed line corresponds to the optimised threshold between states 1 and 2, inferred by the intersection point between the two fitted Gaussians in (c). (f) Comparison between the ground truth and inferred state timeseries of the trajectory (top panel) and the inferred probability of misclassification, $p_{\mathrm{mc}} \equiv \min(p_1, p_2)$ (bottom panel, *c.f.* Fig. 3d).

that the overlap θ_{12} between P_1 and P_2, which we use as an effective cost function, is minimised for some optimal value f^* of the filter width. For values of f in the neighborhood of f^*, each point in the filtered displacement time series will therefore be influenced by only a small number of adjacent points in time, and the correlations introduced by the filtering are weak. In this case, the filtered sequence has enough independence that the central limit theorem applies, and $P(\widehat{\Delta r})$ can be accurately described by a sum of two Gaussian distributions $\mathcal{N}(\mu, \sigma)$, according to

$$P(\widehat{\Delta r}) \approx \widehat{\pi}_1 \mathcal{N}(\widehat{\Delta r}|\widehat{\mu}_1, \widehat{\sigma}_1) + \widehat{\pi}_2 \mathcal{N}(\widehat{\Delta r}|\widehat{\mu}_2, \widehat{\sigma}_2), \qquad (8)$$

with $\widehat{\pi}_i$ denoting the weights of the filtered displacement distribution for each state after the filtering process, and $\widehat{\mu}_i$, $\widehat{\sigma}_i$ the corresponding means and standard deviations. This decomposition allows us to use Gaussian mixture models (GMMs) for fitting the composite distribution. A GMM is a form of soft clustering that models data as a mixture of Gaussian distributions, which, unlike hard

clustering methods that assign each data point to a single cluster without accounting for any uncertainty in cluster membership, compute the probability p_i that a given data point (or displacement) belongs to each component (state) i in the mixture. The parameters of the model — the means, covariances, and weights — are estimated using the expectation-maximization (EM) algorithm, which iteratively updates the probabilities of belonging to a given state and adjusts the model parameters to obtain an optimal fit [55, 56]; in our code, we use the GMM implementation provided by the `GaussianMixture` class in the `scikit-learn` library [57]. GMMs are employed to fit the distribution $P(\widehat{\Delta r})$ and to provide a pointwise estimate of the likelihood p_i that a given displacement originates from a specific diffusive state i. This fitting procedure further enables an approximation of the overlap θ_{12} between the two distributions, expressed as

$$\theta_{12}(f) \approx \int_{-\infty}^{c} \widehat{\pi}_1 \mathcal{N}(\widehat{\Delta r}(f)|\widehat{\mu}_1, \widehat{\sigma}_1)\, d\widehat{\Delta r} + \int_{c}^{+\infty} \widehat{\pi}_2 \mathcal{N}(\widehat{\Delta r}(f)|\widehat{\mu}_2, \widehat{\sigma}_2)\, d\widehat{\Delta r}, \tag{9}$$

where $c \in [\widehat{\mu}_1, \widehat{\mu}_2]$ is the intersection point of the two weighted distributions, corresponding to the point where the probabilities of each of the two states are equal. For systems with multiple states, this generalises to the pairwise sum of the overlap, $\theta(f) = \sum_{i>j} \theta_{ij}(f)$. Fig. 3c shows the filtered version of the displacement distribution in Fig. 3b, which visually demonstrates the radically increased separability between the two states.

After finding f^*, we obtain the time series representing the relative probability of each data point belonging to either of the two diffusive states, as illustrated in Fig. 3d. By assigning each filtered displacement to the state associated with the highest probability, we retrieve a segmented time series of the filtered displacements (Fig. 3e). Additionally, the pointwise probability values serve as quantitative indicators of the confidence associated with the state classification at each time point. In Fig. 3f, we show the resulting, inferred time series of the two diffusive states, together with the "ground truth" transition points used when generating the trajectory. In spite of the relatively large overlap between the unfiltered distributions (Fig. 3b), the reconstructed time series in Fig. 3f shows that all transitions are correctly detected within a small time window, indicating a successful segmentation of the trajectory. Here, we have also show the time series of $p_{\mathrm{mc}} \equiv \min(p_1, p_2)$ inferred from Fig. 3d, which measures the *a priori* probability of misclassification at a given time point. The possibility to temporally monitor this measure provides a fast estimate of the quality of the data and the filtering. In the following section, we will quantify the accuracy of this segmentation and its dependence on the parameters of the problem.

RESULTS & DISCUSSION

Segmentation

In the following, we will demonstrate our method on 2-state Brownian trajectories generated as described in the previous section. Our main focus is on the discriminability between two states, quantified by the ratio of the two diffusion coefficients, $\tilde{D} \equiv D_{\mathrm{fast}}/D_{\mathrm{slow}}$ and the shortest of the two lifetimes, $\tilde{\tau} \equiv \min(\bar{\tau}_1, \bar{\tau}_2)$, measured in units of the shutter time Δt. The effect of the trajectory length, T, will also be examined, as well as the motion blur due to the limited shutter speed, as emulated by the averaging in Eq. (3), and the rescaled tracking error $\tilde{\sigma} = \sigma/\sqrt{2D_{\mathrm{slow}}\Delta t}$.

We however start in the absence of blurring and tracking error, corresponding to $\tilde{\sigma} = 0$ and $n = 1$ in Eq. (3), and investigate the accuracy of the segmentation for various values of $\tilde{D}$ and $\tilde{\tau}$. We define the accuracy as the proportion of time steps at which the inferred diffusive state matches the ground truth:

$$\text{Accuracy} = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbb{1}\left[\hat{s}_t = s_t\right], \tag{10}$$

where $\hat{s}_t$ is the inferred diffusive state, s_t is the ground truth state at time t, and $\mathbb{1}[\cdot]$ is the indicator function, which equals 1 if the prediction is correct and 0 otherwise. Parameter space was explored by generating $N_t = 100$ trajectories of length $T = 1000$ for each pair $(\tilde{D}, \tilde{\tau})$. These parameters are important as $\tilde{\tau}$ puts an upper bound on the range of possible filter size values before the filtering process starts to blur the two states, while $\tilde{D}$ dictates the ideal filter size needed to minimise the overlap between the two states in the case of very large $\tilde{\tau}$.

Fig. 4a shows that our method gives a high precision segmentation (> 90% accuracy) for a large range of ($\tilde{D}$, $\tilde{\tau}$), although becomes inaccurate for a small separation between the diffusion coefficients ($\tilde{D} \leq 3$) or fast transition times ($\tilde{\tau} \leq 20$) compared to the camera resolution. A trade-off between $\tilde{D}$ and $\tilde{\tau}$ is observed, as small values of $\tilde{\tau}$ require a larger value of $\tilde{D}$ and vice versa in order to effectively minimise the cost function and obtain an accurate segmentation. Figures 4b and 4c highlight the underlying optimisation algorithm behind the segmentation method, respectively showing the overlap θ_{12} and the accuracy functions of the filter width f. First of all, we note that θ_{12} is a convex function across the whole domain of interest, allowing for a well-posed optimisation problem. Secondly, there is a strong anticorrelation between the overlap and the segmentation accuracy for all parameter values shown, which supports our use of overlap as a cost function. Although the maximum in accuracy is not perfectly aligned with the minimum in θ_{12}, their colocalisation is strong enough to ensure a

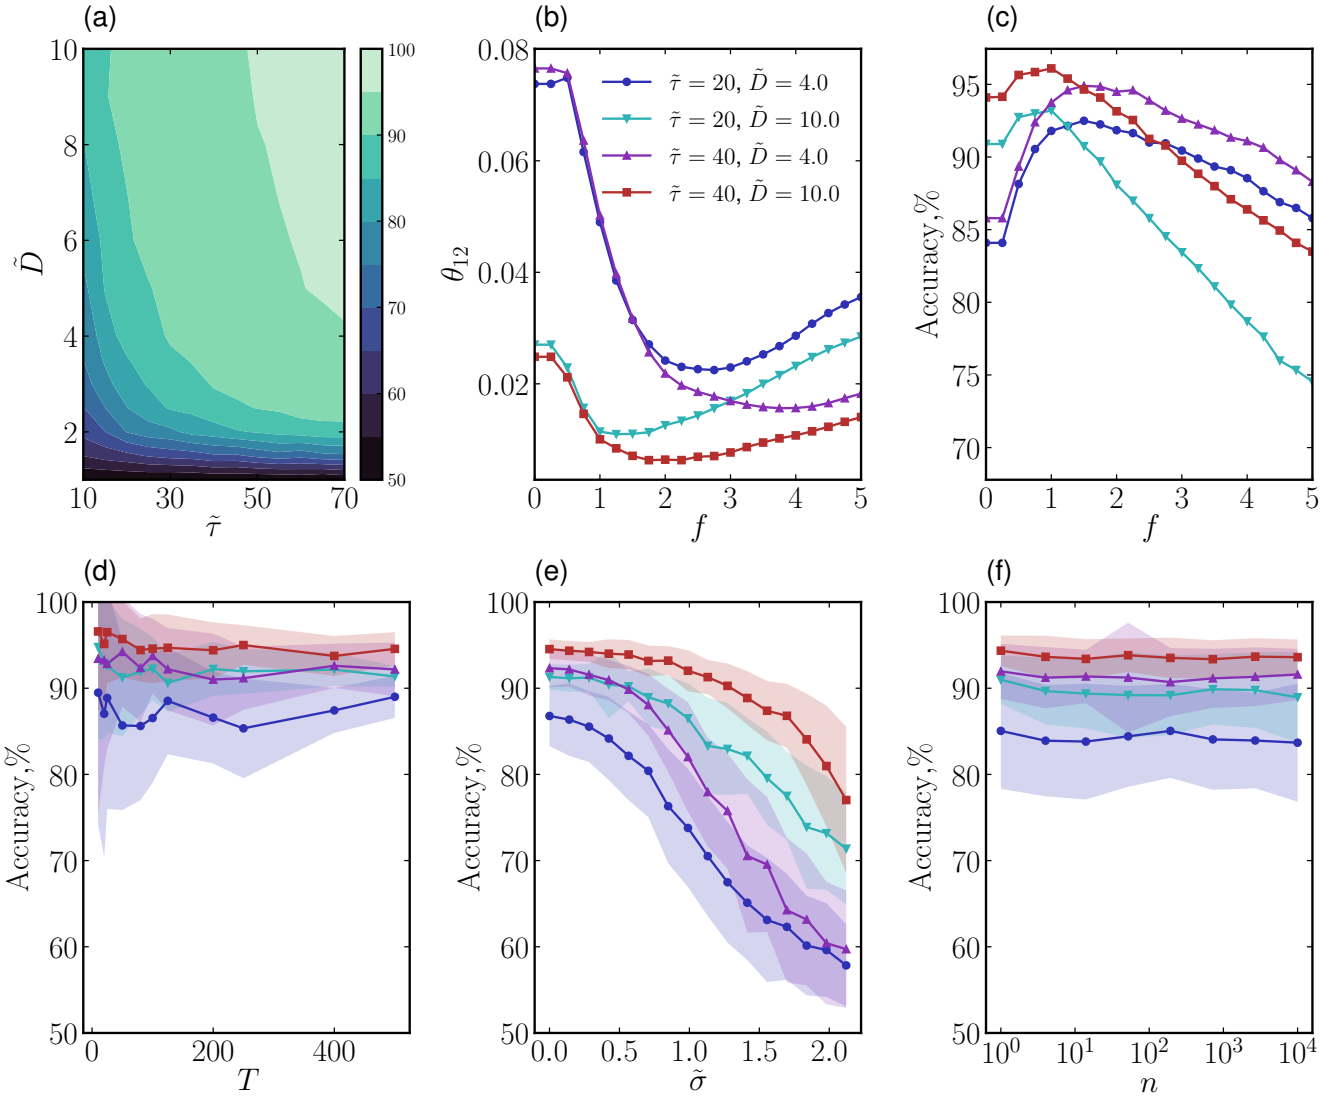

FIG. 4. (a) Colour plot showing the segmentation accuracy as a function of $\tilde{D}$ and $\tilde{\tau}$, obtained after averaging over $N_t = 100$ trajectories of length $T = 2000$. (b) The overlap between the displacement distribution of the two states θ_{12} and (c) the segmentation accuracy, both as a function of the filter width f for the same values of $\tilde{D}$ and $\tilde{\tau}$ as in the other panels. (d) The accuracy as a function of the trajectory length T for various $\tilde{D}$ and $\tilde{\tau}$, keeping fixed the total number of sampled displacements $N_t \times T = 20000$. (e) The accuracy as a function of the localisation error $\tilde{\sigma}$ for the same values of $\tilde{D}$ and $\tilde{\tau}$ as in panel (b). (f) The accuracy as a function of the intensity of motion blur, as mimicked by varying n in Eq. (3). In this case, the "ground truth" when determining the segmentation accuracy for the blurred trajectory was determined through a majority vote of all n points involved in the blurring process. Where applicable, shadings represent one standard deviation of the accuracy over $N_t = 100$ trajectories.

near-optimal performance in spite of the simplicity and fully automated nature of the optimisation method. The results further show that the optimal filter width is generally small, falling in the range $1 \leq f \leq 5$, meaning that it is generally well-separated from the smallest transition time $\bar{\tau}$ apart from the smallest values explored in Fig. 4a.

In Fig. 4d we furthermore demonstrate that the method is independent of the length of individual trajectories T. To verify this, we generate N_t trajectories of varying lengths T, keeping the total number of displacements $N_t \times T$ constant, fixed at 20000 displacements. We observe that the accuracy of the segmentation method remains consistent regardless of trajectory length. While this might seem surprising, it is a direct consequence of the probabilistic nature of the segmentation method, which only relies on sampling a probability distribution of displacements between adjacent time points as in Figs. 3b and c, a process which does not require long individual time series. Therefore, the only strict requirement on the time series length for successful segmentation is that

$T \gg f^*$. However, as we will see in the following section, a reliable estimate of the lifetimes $\bar{\tau}_i$ using the segmented trajectories impose more conservative conditions on the trajectory lengths.

We now add the effect of localisation error and motion blur by generating $N_t = 200$ Brownian trajectories with $T = 1000$, $\tilde{D} = 4.0$ and $\tilde{\tau} = 20$, to which we apply different intensities of both spatial localisation error, as controlled by $\tilde{\sigma}$, and motion blur, adjusted by changing n, the number of discretised steps over which the particle position in the original trajectory is averaged, when generating the trajectory according to Eq. (3). In Fig. 4e, we vary $\tilde{\sigma}$ from 0, corresponding to no localisation error, to 2.1, corresponding to a case where the localisation error is significantly larger than the motion of the particle in its slowest state. The results show that the segmentation accuracy is essentially unaffected by the localisation error for $\tilde{\sigma} < 1$. However, when $\tilde{\sigma} \geq 1$, *i.e.*, when the localisation error is of the order of the instantaneous displacement, the accuracy decays slowly, as expected. This is a consequence of the large noise introduced by the tracking error, which makes the overlap θ_{12} larger and thus the states become harder to distinguish from each other.

In Fig. 4f, we show the corresponding impact of the motion blur as quantified by the number of points n used to generate the rescaled trajectory in Eq. (3). Qualitatively, this yields a spurious correlation between adjacent displacements $\Delta r(\bar{t}_\alpha)$ and $\Delta r(\bar{t}_{\alpha+1})$ [54], which is generally challenging to handle using HMM-based methods [37]. For our segmentation algorithm, the main effect is a slight blurring of the unfiltered probability distribution in Fig. 3b, which however does not lead to any measurable decline in the accuracy of segmentation, as shown in Fig. 4f. Unlike the localisation error σ, a direct translation of the parameter n to experimentally relevant ones is however not straightforward, as the underlying, discretised Brownian process used for the averaging is in itself not a faithful description of the real-world, continuous-time dynamics in the experimental system. Therefore, a continuous-time trajectory with motion blur corresponds to the limit $n \to \infty$. For a full assessment of the impact of motion blur, a variation of the "blur factor", R, which describes the impact of the shutter function on the trajectory is therefore more appropriate, although also more technically complicated. The data generation in Eq. (3) gives equal weight to each position in the underlying trajectory, corresponding to a top-hat (or square-wave) shutter function that counts every photon during the acquisition time, yielding $R = 1/6$ [54].

Extraction of dynamical properties

Once the trajectories have been segmented, the diffusion coefficient for each state can be directly determined from the corresponding trajectory segments. Since the segmented trajectories correspond to *single-state* ones, the dynamical analysis is significantly simplified, allowing for simple extraction of model parameters such as diffusion coefficients and state lifetimes. Assuming isotropic, Brownian diffusion, the displacements in each dimension can be utilised to compute the diffusion coefficient by averaging over both dimensions according to [54]

$$\langle \Delta_i^2 \rangle = 2D_i \Delta t - 2(2D_i R \Delta t - \sigma^2) \tag{11}$$

where $\langle \Delta_i^2 \rangle$ is the mean-squared, unidimensional displacement averaged over the displacements corresponding to the ith state and the blur factor R equals 1/6 if all frames are weighted equally during the acquisition process [54], as in Eq. (3).

To verify the accuracy of the parameter estimation, in Fig. 5a and b we extract the diffusion coefficient from 200 segmented trajectories with $T = 5000$, $\tilde{D} = 10.0$, and $\tilde{\tau} = 20$ using Eq. (11). Fig. 5a shows that, even for low values of the localisation error, the faster of the two diffusion coefficients is systematically underestimated by $\sim 10\%$. This is due to an asymmetry in the misclassification, where by a larger part of the classification mistakes misclassifies displacements from the "slow" state as being "fast" than vice versa, resulting in a slight understimation of D_{fast}. On the other hand, the accuracy in determining D_i is suprisingly robust to the localisation error $\tilde{\sigma}$: even when this is in the region $\tilde{\sigma} \geq 1$ where classification errors start becoming significant (green curve in Fig. 4e), the impact on the predicted diffusion is negligible until $\tilde{\sigma} \geq 1.5$, although its standard deviation (shaded regions in Fig. 5a) increases steeply with $\tilde{\sigma}$, particularly for the slow state which gets the highest relative impact by the noise. This indicates that the additional classification errors introduced by localisation noise for moderate $\tilde{\sigma}$ is symmetric in the sense that it leads to "fast" states being misclassified as "slow" ones as often as the opposite case, creating an effective cancellation of errors when calculating the diffusion coefficients. Fig. 5b, which shows the effect of motion blur on the extraction of D_i, exhibits the same systematic underestimation of D_{fast}, although the accuracy is largely unaffected by the introduction of motion blur even as $n \to \infty$. Importantly, this does not imply that motion blur does not affect the displacement time series, but only illustrates that this effect can be efficiently corrected for using the additional term in Eq. (11), since the calculation of D_i is done on the single-state trajectories after segmentation.

Extracting the lifetimes of the two states is generally a much more demanding task, regardless of the segmentation method, as it intrinsically requires trajectories that are long enough to cover at least two transitions between the two states [50]. Similar to previous methods [58], we estimate the average lifetime $\bar{\tau}_i$ by assuming a Poissonian behaviour and sampling the lifetime τ_i^{m} of segmented trajectories, followed by fitting of the resulting histogram to

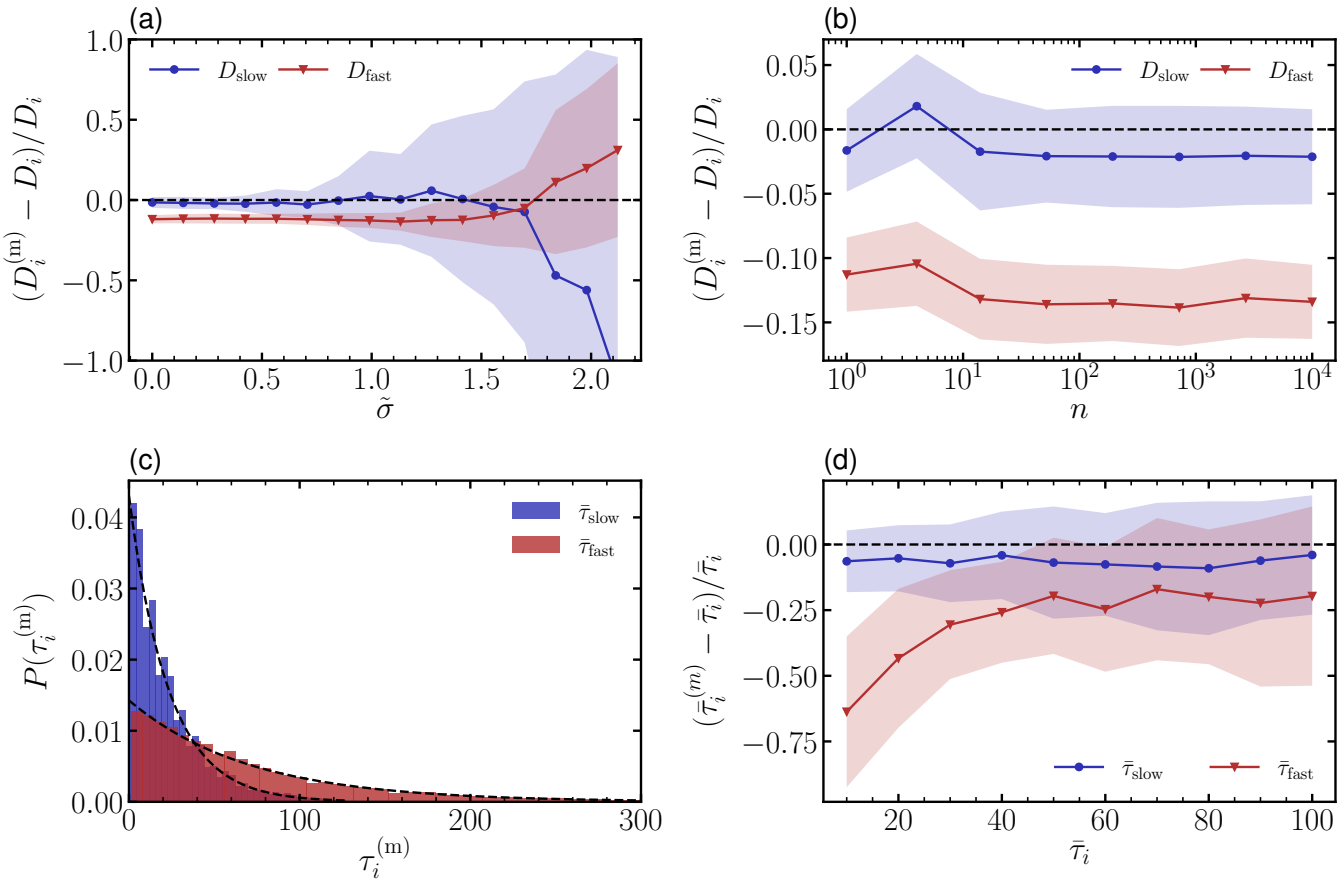

FIG. 5. (ab) Mean relative error in the inferred diffusion coefficient $D_i^{(\mathrm{m})}$, as a function of (a) the rescaled localisation error $\tilde{\sigma}$ and (b) the blurring strength n, in both cases using $\tilde{D} = 10$ and $\bar{\tau}_1 = \bar{\tau}_2 = 20$. (c) Probability distribution of the measured lifetime $\tau_i^{(\mathrm{m})}$. The distributions are fitted to an exponential (dashed lines), yielding estimated average lifetimes of $\bar{\tau}_1^{(\mathrm{m})} = 23.3$ and $\bar{\tau}_2^{(\mathrm{m})} = 70.1$, compared to the true values of $\bar{\tau}_1 = 20$ and $\bar{\tau}_2 = 60$, corresponding to a relative error of approximately 15%. (d) Relative mean error of the measured average lifetimes $\bar{\tau}_{\mathrm{fast}}^{(\mathrm{m})}$ and $\bar{\tau}_{\mathrm{slow}}^{(\mathrm{m})}$ as a function of the true value $\bar{\tau}_i$, using $\bar{\tau}_{\mathrm{fast}} = \bar{\tau}_{\mathrm{slow}}$. Where applicable, shading represents one standard deviation of the estimate measured over $N_t = 200$ trajectories.

an exponential distribution $P(\tau_i^{(\mathrm{m})}) = \exp(-\tau_i^{(\mathrm{m})}/\bar{\tau}_i)/\bar{\tau}_i$, as shown in Fig. 5c. Compared to the determination of diffusion coefficients, this extraction is sensitive to errors in the segmentation process, since the misclassification of a displacement within a state introduces a "virtual transition", leading to an underestimation of lifetimes, and primarily the lifetime $\bar{\tau}_{\mathrm{fast}}$ of the state with the faster diffusion. Nevertheless, Fig. 5d shows a relative error within ±25% for $\tau \geq 30$, consistent with the trend observed in Figure 4a, and indicating that the method can provide reasonably accurate lifetime estimates if using sufficiently numerous and long trajectories. This also illustrates another advantage compared to fitting the unfiltered distribution to a Rayleigh mixture model, which does not enable extraction of $\bar{\tau}$, but only of the ratio between the two lifetimes.

CONCLUSION

In this study, we have presented and verified a novel method for the segmentation of multi-state Brownian trajectories using an automated optimisation of a Gaussian filtering process. By minimising the overlap between the displacement distributions of different diffusive states, our method provides a robust and interpretable framework for trajectory segmentation in cases where the two states are (*i*) sufficiently separable in terms of diffusion coefficients ($D_{\mathrm{fast}}/D_{\mathrm{slow}} \geq 4$) and (*ii*) sufficiently long-lived compared to the camera resolution ($\tilde{\tau} \geq 10$). Our results show that the method performs with high accuracy under experimentally feasible conditions, while being easy to implement and computationally efficient: the segmentation of $N_t = 1000$ trajectories each containing 100 steps requires only approximately 15 seconds on a standard laptop.

The main strength of the proposed method therefore lies in it being accurate but simultaneously lightweight and transparent: the model does not require training, the trajectory segmentation is fast enough to be done on-the-fly, and the quality of the filtered probability distributions and corresponding fits can be directly evaluated visually or by goodness-of-fit estimates to assess the quality of the segmentation. It therefore also presents itself as a potentially attractive alternative for on-line preprocessing of experimental data, to obtain approximate parameters for more advanced methods. Furthermore, the fact that the segmentation step does not require any information about the dynamical parameters such as diffusion coefficients or lifetimes, these can be inferred from the segmented, single-state trajectories, avoiding the complication of parameter fitting to multi-state trajectories. The segmentation furthermore does not rely on an underlying Brownian diffusive process with mean-square displacements that grow linearly in time: Since the segmentation is only based on the Euclidean distance between adjacent points in the trajectory, it can be extended to anomalously diffusing trajectories as long as the diffusivities are reasonably well-separated, although the calculation of D_i using Eq. (11) would need to be generalised accordingly. Finally, unlike some more advanced methods that can infer an optimal number of states, our approach assumes *a priori* knowledge about the optimal number of states. A generalisation of our method to more than 2 states is however very straightforward, and preliminary results indicate similar accuracies for 3-state trajectories as long as the diffusivities are well separated. The method can also be extended towards the detection of the number of components in the Gaussian mixture, similar to previous ML- or HMM-based methods [30, 33–35, 38, 47–53], using extended probabilistic frameworks [55, 57].

In summary, our presented method contributes to the growing pool of single-particle tracking methods by providing a computationally efficient and interpretable alternative to existing segmentation techniques for multi-state Brownian trajectories, forming another methodological tool in the toolbox for understanding the dynamics, molecular transport processes, and interactions in complex cellular systems.

AUTHOR CONTRIBUTIONS

IEK, JML, TD, EC, TA, PJ, and JS conceptualised the research. IEK, JML, and JS designed the model. IEK and JML conducted the software implementation and data analysis. IEK, JML, and JS wrote the paper with input from all the authors.

ACKNOWLEDGMENTS

Helpful discussions with Carlo Manzo and Giovanni Volpe are kindly acknowledged. This work was funded by the Knut and Alice Wallenberg Foundation (grant 2019.0079) and the Swedish Research Council (Grant 2019-03718).

* joakim.stenhammar@fkem1.lu.se

[1] W. E. Moerner and L. Kador, Optical detection and spectroscopy of single molecules in a solid, Phys. Rev. Lett. **62**, 2535 (1989).
[2] T. A. Klar, S. Jakobs, M. Dyba, A. Egner, and S. W. Hell, Fluorescence microscopy with diffraction resolution barrier broken by stimulated emission, Proc. Natl. Acad. Sci. USA **97**, 8206 (2000).
[3] W. E. Moerner and D. P. Fromm, Methods of single-molecule fluorescence spectroscopy and microscopy, Rev. Sci. Instrum. **74**, 3597 (2003).
[4] F. Simon, L. E. Weiss, and S. van Teeffelen, A guide to single-particle tracking, Nat. Rev. Meth. Primers **4**, 66 (2024).
[5] M. Orrit and J. Bernard, Single pentacene molecules detected by fluorescence excitation in a *p*-terphenyl crystal, Phys. Rev. Lett. **65**, 2716 (1990).
[6] G. Seisenberger, M. U. Ried, T. Endreß, H. Büning, M. Hallek, and C. Bräuchle, Real-time single-molecule imaging of the infection pathway of an adeno-associated virus, Science **294**, 1929 (2001).
[7] C. Manzo and M. F. Garcia-Parajo, A review of progress in single particle tracking: From methods to biophysical insights, Rep. Prog. Phys. **78**, 124601 (2015).
[8] H. Shen, L. J. Tauzin, R. Baiyasi, W. Wang, N. Moringo, B. Shuang, and C. F. Landes, Single particle tracking: From theory to biophysical applications, Chem. Rev. **117**, 7331 (2017).
[9] J.-Y. Tinevez, N. Perry, J. Schindelin, G. M. Hoopes, G. D. Reynolds, E. Laplantine, S. Y. Bednarek, S. L. Shorte, and K. W. Eliceiri, TrackMate: An open and extensible platform for single-particle tracking, Methods **115**, 80 (2017).
[10] M. Lelek, M. T. Gyparaki, G. Beliu, F. Schueder, J. Griffié, S. Manley, R. Jungmann, M. Sauer, M. Lakadamyali, and C. Zimmer, Single-molecule localization microscopy, Nat. Rev. Meth. Primers **1**, 39 (2021).
[11] M. de Brabander, R. Nuydens, A. Ishihara, B. Holifield, K. Jacobson, and H. Geerts, Lateral diffusion and retrograde movements of individual cell surface components on single motile cells observed with Nanovid microscopy., J. Cell. Biol. **112**, 111 (1991).
[12] N. Ruthardt, D. C. Lamb, and C. Bräuchle, Single-particle tracking as a quantitative microscopy-based approach to unravel cell entry mechanisms of viruses and pharmaceutical nanoparticles, Mol. Ther. **19**, 1199 (2011).
[13] J. A. Varela, C. Åberg, J. C. Simpson, and K. A. Dawson, Trajectory-based co-localization measures for nanoparticle-cell interaction studies, Small **11**, 2026 (2015).

[14] T. Germier, S. Kocanova, N. Walther, A. Bancaud, H. A. Shaban, H. Sellou, A. Z. Politi, J. Ellenberg, F. Gallardo, and K. Bystricky, Real-time imaging of a single gene reveals transcription-initiated local confinement, Biophys. J. **113**, 1383 (2017).

[15] W. M. Śmigiel, L. Mantovanelli, D. S. Linnik, M. Punter, J. Silberberg, L. Xiang, K. Xu, and B. Poolman, Protein diffusion in *Escherichia coli* cytoplasm scales with the mass of the complexes and is location dependent, Sci. Adv. **8**, eabo5387 (2022).

[16] E. Amselem, B. Broadwater, T. Hävermark, M. Johansson, and J. Elf, Real-time single-molecule 3D tracking in *E. coli* based on cross-entropy minimization, Nat. Commun. **14**, 1336 (2023).

[17] V. Sposini, D. Krapf, E. Marinari, R. Sunyer, F. Ritort, F. Taheri, C. Selhuber-Unkel, R. Benelli, M. Weiss, R. Metzler, and G. Oshanin, Towards a robust criterion of anomalous diffusion, Commun. Phys. **5**, 305 (2022).

[18] M. Woringer, I. Izeddin, C. Favard, and H. Berry, Anomalous subdiffusion in living cells: Bridging the gap between experiments and realistic models through collaborative challenges, Front. Phys. **8**, 134 (2020).

[19] R. Simson, E. D. Sheets, and K. Jacobson, Detection of temporary lateral confinement of membrane proteins using single-particle tracking analysis, Biophys. J. **69**, 989 (1995).

[20] S. Huet, E. Karatekin, V. S. Tran, I. Fanget, S. Cribier, and J.-P. Henry, Analysis of transient behavior in complex trajectories: Application to secretory vesicle dynamics, Biophys. J. **91**, 3542 (2006).

[21] T. Morisaki, W. G. Müller, N. Golob, D. Mazza, and J. G. McNally, Single-molecule analysis of transcription factor binding at transcription sites in live cells, Nat. Commun. **5**, 4456 (2014).

[22] J.-B. Masson, P. Dionne, C. Salvatico, M. Renner, C. G. Specht, A. Triller, and M. Dahan, Mapping the energy and diffusion landscapes of membrane proteins at the cell surface using high-density single-molecule imaging and Bayesian inference: Application to the multiscale dynamics of glycine receptors in the neuronal membrane, Biophys. J. **106**, 74 (2014).

[23] L. Veya, J. Piguet, and H. Vogel, Single molecule imaging deciphers the relation between mobility and signaling of a prototypical G protein-coupled receptor in living cells, J. Biol. Chem. **290**, 27723 (2015).

[24] V. T. Chang, R. A. Fernandes, K. A. Ganzinger, S. F. Lee, C. Siebold, J. McColl, P. Jönsson, M. Palayret, K. Harlos, C. H. Coles, E. Y. Jones, Y. Lui, E. Huang, R. J. C. Gilbert, D. Klenerman, A. R. Aricescu, and S. J. Davis, Initiation of T cell signaling by CD45 segregation at 'close contacts', Nat. Immunol. **17**, 574 (2016).

[25] M. Chouliara, V. Junghans, T. Dam, A. M. Santos, S. J. Davis, and P. Jönsson, Single-cell measurements of two-dimensional binding affinity across cell contacts, Biophys. J. **120**, 5032 (2021).

[26] X. Michalet, Mean square displacement analysis of single-particle trajectories with localization error: Brownian motion in an isotropic medium, Phys. Rev. E **82**, 041914 (2010).

[27] X. Michalet and A. J. Berglund, Optimal diffusion coefficient estimation in single-particle tracking, Phys. Rev. E **85**, 061916 (2012).

[28] C. L. Vestergaard, P. C. Blainey, and H. Flyvbjerg, Optimal estimation of diffusion coefficients from single-particle trajectories, Phys. Rev. E **89**, 022726 (2014).

[29] P. J. Bosch, J. S. Kanger, and V. Subramaniam, Classification of dynamical diffusion states in single molecule tracking microscopy, Biophys. J. **107**, 588 (2014).

[30] S. A. McKinney, C. Joo, and T. Ha, Analysis of single-molecule FRET trajectories using hidden Markov modeling, Biophys. J. **91**, 1941 (2006).

[31] R. Das, C. W. Cairo, and D. Coombs, A hidden Markov model for single particle tracks quantifies dynamic interactions between LFA-1 and the actin cytoskeleton, PLoS Comput. Biol. **5**, e1000556 (2009).

[32] M. Ott, Y. Shai, and G. Haran, Single-particle tracking reveals switching of the HIV fusion peptide between two diffusive modes in membranes, J. Phys. Chem. B **117**, 13308 (2013).

[33] F. Persson, M. Lindén, C. Unoson, and J. Elf, Extracting intracellular diffusive states and transition rates from single-molecule tracking data, Nat. Meth. **10**, 265 (2013).

[34] N. Monnier, Z. Barry, H. Y. Park, K.-C. Su, Z. Katz, B. P. English, A. Dey, K. Pan, I. M. Cheeseman, R. H. Singer, and M. Bathe, Inferring transient particle transport dynamics in live cells, Nat. Meth. **12**, 838 (2015).

[35] Y. Matsuda, I. Hanasaki, R. Iwao, H. Yamaguchi, and T. Niimi, Estimation of diffusive states from single-particle trajectory in heterogeneous medium using machine-learning methods, Phys. Chem. Chem. Phys. **20**, 24099 (2018).

[36] A. Federbush, A. Moscovich, and Y. Bar-Sinai, Hidden Markov modeling of single-particle diffusion with stochastic tethering, Phys. Rev. E **109**, 034129 (2024).

[37] P. J. Slator, C. W. Cairo, and N. J. Burroughs, Detection of diffusion heterogeneity in single particle tracking trajectories using a hidden Markov model with measurement noise propagation, PLoS One **10**, e0140759 (2015).

[38] L. R. Baker, M. T. Gordon, B. P. Ziemba, V. Gershuny, J. J. Falke, and D. M. Bortz, Learning diffusion coefficients, kinetic parameters, and the number of underlying states from a multistate diffusion process: Robustness results and application to PDK1/PKCα dynamics, in *Handbook of Visual, Experimental and Computational Mathematics: Bridges through Data*, edited by B. Sriraman (Springer Nature Switzerland, Cham, 2025) pp. 1–36.

[39] E. Clarkson and T. Ambjörnsson, Bayesian and frequentist analyses of two-state single-molecule diffusion trajectories, J. Phys. A: Math. Theor. **58**, 195004 (2025).

[40] Y. Lanoiselée and D. S. Grebenkov, Unraveling intermittent features in single-particle trajectories by a local convex hull method, Phys. Rev. E **96**, 022144 (2017).

[41] S. Yin, N. Song, and H. Yang, Detection of velocity and diffusion coefficient change points in single-particle trajectories, Biophys. J. **115**, 217 (2018).

[42] A. R. Vega, S. A. Freeman, S. Grinstein, and K. Jaqaman, Multistep track segmentation and motion classification for transient mobility analysis, Biophys. J. **114**, 1018 (2018).

[43] A. Celisse, G. Marot, M. Pierre-Jean, and G. J. Rigaill, New efficient algorithms for multiple change-point detection with reproducing kernels, Comp. Stat. Data Anal. **128**, 200 (2018).

[44] S. Arlot, A. Celisse, and Z. Harchaoui, A kernel multiple change-point algorithm via model selection, J. Mach.

Learn. Res. **20**, 1 (2019).

[45] H. Seckler, J. Szwabinski, and R. Metzler, Machine-learning solutions for the analysis of single-particle diffusion trajectories, J. Phys. Chem. Lett. **14**, 7910 (2023).

[46] G. Muñoz-Gil, M. A. Garcia-March, C. Manzo, J. D. Martín-Guerrero, and M. Lewenstein, Single trajectory characterization via machine learning, New. J. Phys. **22**, 013010 (2020).

[47] M. Arts, I. Smal, M. W. Paul, C. Wyman, and E. Meijering, Particle mobility analysis using deep learning and the moment scaling spectrum, Sci. Rep. **9**, 17160 (2019).

[48] N. Granik, L. E. Weiss, E. Nehme, M. Levin, M. Chein, E. Perlson, Y. Roichman, and Y. Shechtman, Single-particle diffusion characterization by deep learning, Biophys. J. **117**, 185 (2019).

[49] J. Li, L. Zhang, A. Johnson-Buck, and N. G. Walter, Automatic classification and segmentation of single-molecule fluorescence time traces with deep learning, Nat. Commun. **11**, 5833 (2020).

[50] G. Muñoz-Gil, G. Volpe, M. A. Garcia-March, E. Aghion, A. Argun, C. B. Hong, T. Bland, S. Bo, J. A. Conejero, N. Firbas, Ó. Garibo-i-Orts, A. Gentili, Z. Huang, J.-H. Jeon, H. Kabbech, Y. Kim, P. Kowalek, D. Krapf, H. Loch-Olszewska, M. A. Lomholt, J.-B. Masson, P. G. Meyer, S. Park, B. Requena, I. Smal, T. Song, J. Szwabiński, S. Thapa, H. Verdier, G. Volpe, A. Widera, M. Lewenstein, R. Metzler, and C. Manzo, Objective comparison of methods to decode anomalous diffusion, Nat. Commun. **12**, 6253 (2021).

[51] B. Requena, S. Masó-Orriols, J. Bertran, M. Lewenstein, C. Manzo, and G. Muñoz-Gil, Inferring pointwise diffusion properties of single trajectories with deep learning, Biophys. J. **122**, 4360 (2023).

[52] J. Pineda, B. Midtvedt, H. Bachimanchi, S. Noé, D. Midtvedt, G. Volpe, and C. Manzo, Geometric deep learning reveals the spatiotemporal features of microscopic motion, Nature. Mach. Intell. **5**, 71 (2023).

[53] J. Kæstel-Hansen, M. de Sautu, A. Saminathan, G. Scanavachi, R. F. Bango Da Cunha Correia, A. Juma Nielsen, S. V. Bleshøy, K. Tsolakidis, W. Boomsma, T. Kirchhausen, and N. S. Hatzakis, Deep learning-assisted analysis of single-particle tracking for automated correlation between diffusion and function, Nat. Meth. **22**, 1091 (2025).

[54] A. J. Berglund, Statistics of camera-based single-particle tracking, Phys. Rev. E **82**, 011917 (2010).

[55] C. M. Bishop, *Pattern recognition and machine learning* (Springer, 2006).

[56] T. Hastie, *The elements of statistical learning: Data mining, inference, and prediction* (Springer, 2009).

[57] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, Scikit-learn: Machine learning in Python, J. Mach. Learn. Res. **12**, 2825 (2011).

[58] L. Schrangl, V. Mühlgrabner, R. Platzer, F. Kellner, J. Wieland, R. Obst, J. L. Toca-Herrera, J. B. Huppa, G. J. Schütz, and J. Göhring, Advanced quantification of receptor–ligand interaction lifetimes via single-molecule FRET microscopy, Biomolecules **14**, 1001 (2024).