

A Hidden Markov Model for Detecting Confinement in Single-Particle Tracking Trajectories

Paddy J. Slator^{1,2} and Nigel J. Burroughs^{3,*}

¹Centre for Medical Image Computing and Department of Computer Science, University College London, London, United Kingdom; ²Systems Biology Doctoral Training Centre, University of Warwick, Coventry, United Kingdom; and ³Mathematics Institute, University of Warwick, Coventry, United Kingdom

ABSTRACT State-of-the-art single-particle tracking (SPT) techniques can generate long trajectories with high temporal and spatial resolution. This offers the possibility of mechanistically interpreting particle movements and behavior in membranes. To this end, a number of statistical techniques have been developed that partition SPT trajectories into states with distinct diffusion signatures, allowing a statistical analysis of diffusion state dynamics and switching behavior. Here, we develop a confinement model, within a hidden Markov framework, that switches between phases of free diffusion and confinement in a harmonic potential well. By using a Markov chain Monte Carlo algorithm to fit this model, automated partitioning of individual SPT trajectories into these two phases is achieved, which allows us to analyze confinement events. We demonstrate the utility of this algorithm on a previously published interferometric scattering microscopy data set, in which gold-nanoparticle-tagged ganglioside GM1 lipids were tracked in model membranes. We performed a comprehensive analysis of confinement events, demonstrating that there is heterogeneity in the lifetime, shape, and size of events, with confinement size and shape being highly conserved within trajectories. Our observations suggest that heterogeneity in confinement events is caused by both individual nanoparticle characteristics and the binding-site environment. The individual nanoparticle heterogeneity ultimately limits the ability of interferometric scattering microscopy to resolve molecule dynamics to the order of the tag size; homogeneous tags could potentially allow the resolution to be taken below this limit by deconvolution methods. In a wider context, the presented harmonic potential well confinement model has the potential to detect and characterize a wide variety of biological phenomena, such as hop diffusion, receptor clustering, and lipid rafts.

INTRODUCTION

Single-particle tracking (SPT) experiments directly observe the motion of single molecules and hence offer a powerful method to analyze the membrane environment. For instance, detection and characterization of heterogeneous diffusion behaviors yields information on membrane structure (1,2). However, SPT methods require the molecule of interest to be tagged with a trackable label that is imaged over a number of time steps. A number of experimental design limitations constrain the amount of information that can be extracted from such data, including spatial accuracy, temporal resolution, and the tracking period. New technologies are capable of extending the trajectory length while retaining high sampling rates and high spatial resolution. For example, interferometric scattering microscopy (iSCAT)

can generate very long (50,000 step) trajectories, with high spatial (<2 nm) and temporal (up to 500 kHz) resolution (3–6). However, a fundamental problem that impacts interpretation is the effect of the tag itself (7). This is particularly relevant for iSCAT because the gold nanoparticle (AuNP) tags are 20–40 nm in diameter, whereas spatial resolution is estimated to be ~2 nm for a 20-nm AuNP (4,5); relative movements between the AuNP and the bound GM1 will thus convolve with the movement of the GM1. For example, an iSCAT study on model membranes demonstrated both Gaussian-like and ring-like confinement events, which was ascribed to transient multivalent binding of the tag (4). Thus, to extend this technique to *in vivo* experiments, there is a need to deconvolve the tag signature from the environment signal. Failure to achieve this separation means that interpretation of the high-resolution dynamics measured by these techniques may be limited to the order of the tag's size.

Analysis of SPT data is not straightforward primarily because of the stochastic nature of diffusion. This has led

Submitted February 27, 2018, and accepted for publication September 4, 2018.

*Correspondence: n.j.burroughs@warwick.ac.uk

Editor: Markus Deserno.

<https://doi.org/10.1016/j.bpj.2018.09.005>

© 2018 The Authors.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).



to the development of a range of statistical methods that detect deviations from Brownian motion, such as mean-square displacement (MSD) (8–13) and confinement (14–19) analyses. A new breed of methods model switching of the movement dynamics between various dynamic states (20–24), often within a hidden Markov chain framework (25–31). For high-resolution data, the latter techniques can utilize the high level of information present in the trajectory to extract detailed motion characteristics and potentially infer underlying biophysical mechanisms.

However, the majority of existing hidden Markov approaches only incorporate changes in the particle diffusivity and (or) drift. Such methods can only approximate confinement through a change in the particle's effective diffusion coefficient. To our knowledge, the only exception is the work of Bernstein et al. (31) who explicitly model confinement using a hidden Markov model (HMM) within a maximal likelihood framework. In this article, we develop an HMM harmonic potential well (HPW) confinement analysis method using a Bayesian approach. Specifically, the particle moves between two states hidden to the observer: free diffusion with (to be determined) diffusion coefficient D and confinement in an HPW (center location and well strength to be determined). We developed a Markov chain Monte Carlo (MCMC) algorithm to infer model parameters and hidden states from a single trajectory. We tested the algorithm on simulated data, then applied it to previously published experimental iSCAT trajectories of GM1 lipids diffusing in model membranes (4). Specifically, a (20 or 40 nm) AuNP was coated in cholera toxin B subunits (CTxBs) by streptavidin binding, with each CTxB then binding 5 GM1 molecules in the lipid membrane to form an AuNP/CTxB/GM1 complex. In trajectories of 20 nm AuNP/CTxB/GM1 diffusing in model membranes on a glass substrate, we detected clear periods of trapping in wells of a mean radius of 18 nm with a mean trapping time of 0.024 s. However, we also observed inherent heterogeneities in both AuNP/CTxB/GM1 particles and trapping sites, which ultimately affect trajectory characteristics.

This article is organized as follows. In **Methods**, we introduce the HPW confinement model and an associated inference (MCMC) algorithm. The full derivation of the MCMC algorithm is described in Note S1 of the **Supporting Materials and Methods**. In **Results**, we demonstrate accurate inference of model parameters and hidden states on simulated trajectories, then apply the algorithm to iSCAT trajectories of AuNP/CTxB/GM1 diffusing in model membranes.

METHODS

We implemented the following methods in MATLAB (The MathWorks, Natick, MA). The source code, documentation, trajectory data, and working examples are freely available (<https://doi.org/10.5281/zenodo.1405647>).

HPW model

We developed a model for a particle that switches between a freely diffusing state and a confinement state localized around a slowly diffusing center. The state is encoded by a hidden variable z , with $z_i = 0$ if the particle is freely diffusing at time t_i and $z_i = 1$ if confined, where $i = 1..N$ denotes the time point (i.e., frame). The state z_{i+1} depends only on z_i with transition probabilities (constant frame rate):

$$\text{free}(z = 0) \xrightleftharpoons[p_{\text{trap}}]{p_{\text{esc}}} \text{confined}(z = 1), \quad (1)$$

where p_{trap} and p_{esc} are the per-frame probabilities of switching into and out of confinement, respectively. The probability of being in state z_{i+1} given state z_i is therefore as follows:

$$\pi(z_{i+1} | z_i) = \text{Bernoulli}(z_{i+1}; z_i(1 - p_{\text{esc}}) + (1 - z_i)p_{\text{trap}}), \quad (2)$$

where Bernoulli ($x; p$) denotes the Bernoulli probability distribution with variable x and parameter p . In the free state, the particle diffuses freely with diffusion coefficient D . In the confined state, the particle is assumed to have a directed component to its diffusive motion, proportional to the distance from the well center C_i , i.e., the force is proportional to $X_i - C_i$ where X_i is the particle position at time t_i . (Note that X_i and C_i are two-dimensional (2D) vectors). During confinement, the center diffuses much slower than the particle itself (diffusion coefficient $D_C \ll D$). When the particle is free, C diffuses with diffusion coefficient D_{est} , where D_{est} is sufficiently high that the center can relocate between different confinement sites. The center is thus still present even when it is not affecting the particle. The stochastic differential equations for this model are as follows:

$$dX_t = -\kappa z_t (X_t - C_t) dt + \sqrt{2D} dW_t \text{ and} \quad (3)$$

$$dC_t = \sqrt{2(D_C z_t + D_{\text{est}}(1 - z_t))} dW_t^{(C)}, \quad (4)$$

where $W_t, W_t^{(C)}$ are independent Weiner processes. During confinement, X_t has Ornstein-Uhlenbeck (OU) dynamics with center C_t . We assume that switching can only occur at the sampling points. We also assume that C_t is slowly varying and therefore ignore its time dependence over the time step Δt . The frame-to-frame dynamics are hence,

$$X_{i+1} - X_i \sim N\left(z_i(C_i - X_i)(1 - e^{-\kappa\Delta t}), D\left((1 - z_i)2\Delta t_i + \frac{z_i}{\kappa}(1 - e^{-2\kappa\Delta t_i})\right)\right), \text{ and} \quad (5)$$

$$C_{i+1} - C_i \sim N(0, 2\Delta t_i(D_C z_i + D_{\text{est}}(1 - z_i))). \quad (6)$$

See Note S1 of the **Supporting Materials and Methods** for full details. If the step size is sufficiently small relative to the confinement strength ($\kappa\Delta t \ll 1$), an Euler-Maruyama approximation is justified, but if the particle explores the well over Δt , this OU solution is required. We refer to this discrete-time stochastic model as the HPW confinement model.

The model has two hidden states to be inferred at all trajectory time points $i = 1, \dots, N$: the state z_i (confined or free) and the position of the HPW center C_i when confined. There are also five parameters to be inferred: two diffusion coefficients (D and D_C), the strength of the HPW (κ), and two transition probabilities (p_{esc} and p_{trap}). D_{est} is treated separately because it only weakly affects the trajectory and does not affect the likelihood or parameter estimates provided it is sufficiently high. Fig. 1 A shows a simulated HPW model trajectory. In the simulation, we include a drift

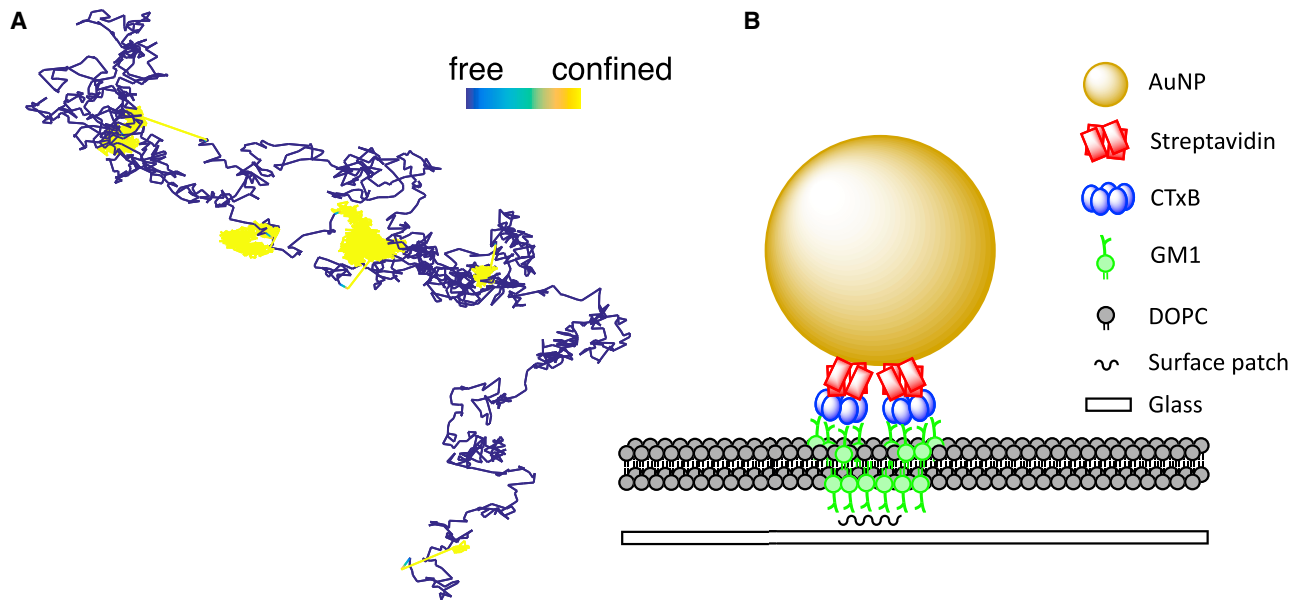


FIGURE 1 Simulated harmonic potential well (HPW) model trajectory. (A) A simulated trajectory colored by state. Model parameters are $D = 0.5 \mu\text{m}^2\text{s}^{-1}$, $D_{est} = 0.5 \mu\text{m}^2\text{s}^{-1}$, $D_C = 0.01 \mu\text{m}^2\text{s}^{-1}$, $\kappa = 3000 \text{ s}^{-1}$, $p_{esc} = 0.001$, $p_{trap} = 0.002$, time step $2 \times 10^4 \text{ s}$, and $N = 5000$ frames. The simulation was performed using Eq. 5 and a modified version of Eq. 6 as discussed in the main text. Trajectory colored blue, denoting free diffusion or yellow, denoting the confined state. Color bar length, $0.1 \mu\text{m}$. (B) A schematic of AuNP/CTxB/GM1 complex in 1,2-dioleoyl-sn-glycero-3-phosphocholine (DOPC) lipid bilayer, based on figure 5 in (4).

term for the center so that C tracks X when not confined. This ensures that C is close to X when the particle switches from free diffusion to confinement, and therefore confinement zones remain within a reasonable field of view. This tracking of X by C is not included in the inference algorithm because diffusion alone is sufficient to allow the Markov chain to find high-probability paths.

MCMC sampler

There are a number of MCMC samplers for linear switching models in the literature; the main distinction is whether variables are integrated out using an inverse Wishart prior (32), or a Markov chain incorporating all variables is used. The latter approach allows greater control of prior information, including use of uninformative priors, whereas the Wishart distribution, motivated by computational convenience, imposes a dependence between variable correlations and scale, which is a concern for inference (33). We developed an MCMC algorithm (Supporting Materials and Methods, Note S1) for the full system of variables to fit the HPW model (Eqs. 5 and 6) to 2D trajectory data, $\mathbf{X} = \{X_i, t_i\}_{i=1}^N$. We chose uninformative priors for all parameters except for the transition probabilities, where we use an informative prior to restrict rapid switching between states (details in Note S1 of the Supporting Materials and Methods). For an SPT trajectory, the algorithm samples the posterior distribution, $\pi(\theta, \mathbf{z}, \mathbf{C} | \mathbf{X})$, giving K samples of the parameters $\theta^{(k)} = \{D^{(k)}, D_C^{(k)}, \kappa^{(k)}, p_{esc}^{(k)}, p_{trap}^{(k)}\}_{k=1}^K$ and hidden states $\{\mathbf{z}^{(k)}, \mathbf{C}^{(k)}\}_{k=1}^K$. Here, for each sample k , $\mathbf{z} = \{z_i\}_{i=1}^{N-1}$ and $\mathbf{C} = \{C_i\}_{i=1}^{N-1}$ are the set of hidden states and center locations (2D vectors) throughout the trajectory.

We determined convergence of the MCMC sampler by calculating the Gelman potential scale reduction factor (PSRF) (34), considering a run converged provided the PSRF was below a threshold in all variables, set to 1.2 on experimental trajectories. The MCMC run length was increased up to a maximum of 4×10^5 steps on trajectories that failed the convergence criteria on shorter runs.

GM1 molecules diffusing in model membranes

We applied the MCMC algorithm to previously published iSCAT SPT data (4), where CTxB-coated AuNPs were introduced to a 1,2-dioleoyl-sn-glycero-3-phosphocholine lipid bilayer containing 0.03% GM1 lipids (Fig. 1 B). A confinement event corresponds to an interaction between an AuNP/CTxB/GM1 complex on the upper leaflet with a lower leaflet GM1 that is immobilized on a glass surface. This was previously referred to as “inter-leaflet coupling and molecular pinning” (4). Both Gaussian and non-Gaussian confinement events were observed; we investigated these events in greater detail using our HPW model.

The data set includes 71 trajectories of 20-nm AuNP/CTxB/GM1 diffusing in a model membrane on a glass substrate and 18 trajectories of 40-nm AuNP/CTxB/GM1 in a model membrane on a mica substrate. There is a dynamic error in the localization accuracy at the 50-kHz sampling rate resulting in apparent superdiffusive behavior, which we removed by subsampling down to 5 kHz (Figs. S1 and S2; Note S2 of the Supporting Materials and Methods). We also removed trajectory artifacts because of multiple AuNPs in the focal area (Supporting Materials and Methods, Note S2). The MCMC algorithm did not converge on five 20-nm AuNP/CTxB/GM1 trajectories (PSRF convergence criteria of 1.2) leaving a set of 66 trajectories for further analysis. MCMC runs on all 18 40-nm AuNP/CTxB/GM1 on mica trajectories converged.

Thresholding hidden states for lifetime analysis

For each trajectory, at each time point i , we computed the probability of confinement $\pi(z_i | \mathbf{X})$ from the MCMC posterior distribution samples. This probability distribution is concentrated near 0 and 1 on experimental trajectories (only 2.7% of the state probabilities were between 0.2 and 0.8; Fig. S3), indicating high confidence in confinement state estimates. To annotate the trajectory by state, we define the binary signal, $z_i^B = 0$ or $z_i^B = 1$, for free diffusion and confinement respectively, using a threshold of 0.5 on $\pi(z_i | \mathbf{X})$. We then identify confinement events as a series of

ones in the (posterior) binary state vector \mathbf{z}^B and free diffusions as a series of zeros, allowing event lifetimes (and per-event spatial statistics) to be computed. When considering event lifetimes, we exclude those containing either the first or last time point of the trajectory, because the full event is not witnessed, hence the state lifetime is unknown.

Confinement event profiling

To analyze confinement events in 20-nm AuNP/CTxB/GM1 trajectories, we utilized spatial statistics (including the mean confinement radius and radial skewness, defined in Table 1) based on the Euclidean distance between the particle and the confinement center. We calculate these statistics for all events of at least 0.01 s (50 frames). Unlike the event lifetime analysis, we allow events that contain either the first or last (or both) time points. Furthermore, we compute statistics including events revisiting a previous confinement zone where applicable (details in Table 1). These restrictions left a set of 271 confinement events when excluding repeat events and 427 when including them. The number of events within a trajectory ranges from 1 (there were six examples in which the particle remained trapped for the entire trajectory) to 11 (without repeats) or 25 (with repeats).

RESULTS

MCMC on simulated data

The HPW model sampler was extensively tested on simulated data. Figs. 2 and 3 show an MCMC run on the simulated trajectory of Fig. 1 A. The parameter posteriors are consistent with the true (i.e., simulation) values (Fig. 2). When confined, the inferred center closely tracks the simulated center (Fig. 3, A and B), and every confinement event is accurately inferred (Fig. 3 C). The inferred model parameters are independent of the algorithm parameter D_{est} in both the noiseless case (Fig. S5) and in the presence of static localization error at the same level as the 20-nm AuNP/CTxB/GM1 trajectories (Fig. S6); p_{trap} , and p_{esc} are typically underestimated because of their informative priors. There also appears to be a small underestimation in the confinement strength κ in the presence of localization error,

which we further investigated by increasing the localization error's SD (Fig. S7). This revealed a clear trend, with an increase in confinement strength underestimation bias with localization error. However, the effect is relatively small at the noise level (SD 2.7 nm) of the experimental data presented in this article. Performance was robust to changes in trajectory length and number of events (Figs. S8 and S9). Estimation of D , D_C , κ , \mathbf{z} , and switching rates were robust to the time series subsampling rate (Figs. S10 and S11); in particular, most events were still detected even with a 10-fold subsampling (Fig. S11). Parameter estimation is also robust to changes in confinement strength κ (Fig. S12). However, we found that confinement strength estimation accuracy decreases dramatically when the confinement center diffusivity D_C approaches 10% of the free diffusivity D (Fig. S13), reflecting the degeneracy of the problem when $D_C \approx D$.

MCMC on 20-nm AuNP/CTxB/GM1 on glass trajectories

An example of the model fit is shown in Fig. 4, with the segmented trajectory shown in Fig. 4 E. Fig. 5 shows the associated parameter posterior estimates. Particle state (confined or free diffusion) is well determined, with state probabilities near zero or one (Fig. 4 C). The parameter estimates for D and κ are tight (low relative SD), whilst the diffusion coefficient of the center is very low, $D_C = 0.010 \pm 0.0009 \mu\text{m}^2\text{s}^{-1}$ (mean \pm SD) compared to $D = 0.52 \pm 0.017 \mu\text{m}^2\text{s}^{-1}$, indicating near complete immobilization of the well. The inferred position of the well center is also practically stationary in both coordinates during periods of confinement consistent with immobilization (Fig. 4, A and B). As an independent measure of changes in mobility, we estimated the effective local diffusion coefficient (Fig. 4 D), which demonstrates a clear shift at around 0.5 s (i.e., the first inferred switch point). By color coding the trajectory according to the probability of being confined per frame (Fig. 4 E), we can extract periods of confinement with non-Gaussian occupation profiles (Fig. 4, F and H). In this trajectory, we observed that one confinement zone is visited twice (Fig. S15) and that the repeat confinements had remarkably similar occupation profiles (Fig. 4, G and H). The probability per frame of switching is reasonably well inferred (Fig. 5 C) despite the small number of events. The probability of escape from a confinement zone is smaller than the probability of trapping, reflecting the short periods of time that the AuNP/CTxB/GM1 complex undergoes free diffusion.

Applying our MCMC algorithm to the 66 trajectories, we obtain parameter estimates across the population (Fig. S16). The mean value of D over all trajectories was $1.15 \pm 0.106 \mu\text{m}^2\text{s}^{-1}$ (mean \pm standard error (SE); population SD $0.86 \mu\text{m}^2\text{s}^{-1}$); MSD analysis (using the @msdalyzer package (35)) gave a smaller estimate,

TABLE 1 Calculation of Confinement Event Statistics

Statistic	Calculation
Confinement radius	$\mathbf{R}_{lm} = \{R_i\}_{i \in T_{lm}}, R_i = \ X_i - \bar{C}_{lm}\ $
Mean confinement radius	$\bar{R}_{lm} = \frac{1}{M_{lm}} \sum_{i \in T_{lm}} R_i$
Radial skewness	$S_{lm} = \frac{\frac{1}{M_{lm}} \sum_{i \in T_{lm}} (R_i - \bar{R}_{lm})^3}{\left[\frac{1}{M_{lm}} \sum_{i \in T_{lm}} (R_i - \bar{R}_{lm})^2 \right]^{3/2}}$
Radial mean-median distance	$ \bar{R}_{lm} - \hat{R}_{lm} $
Radial SD	$\sqrt{\text{var}[\mathbf{R}_{lm}]}$

We denote the time points of the m th trapping event in the l th trajectory T_{lm} . Events have associated particle positions $\mathbf{X}_{lm} = \{X_i\}_{i \in T_{lm}}$ and harmonic well center positions $\mathbf{C}_{lm} = \{C_i\}_{i \in T_{lm}}$. The mean posterior harmonic well center is given by $\bar{C}_{lm} = (1/M_{lm}) \sum_{i \in T_{lm}} C_i$, where M_{lm} is the number of time points in T_{lm} . To remove events that revisit a previous trapping zone, we did not include events if \bar{C}_{lm} was within 30 nm of a previous confinement center ($\bar{C}_{lm}, n < m$) within trajectory 1. $\|\cdot\|$ denotes the Euclidean distance, and \hat{R}_{lm} denotes the median.

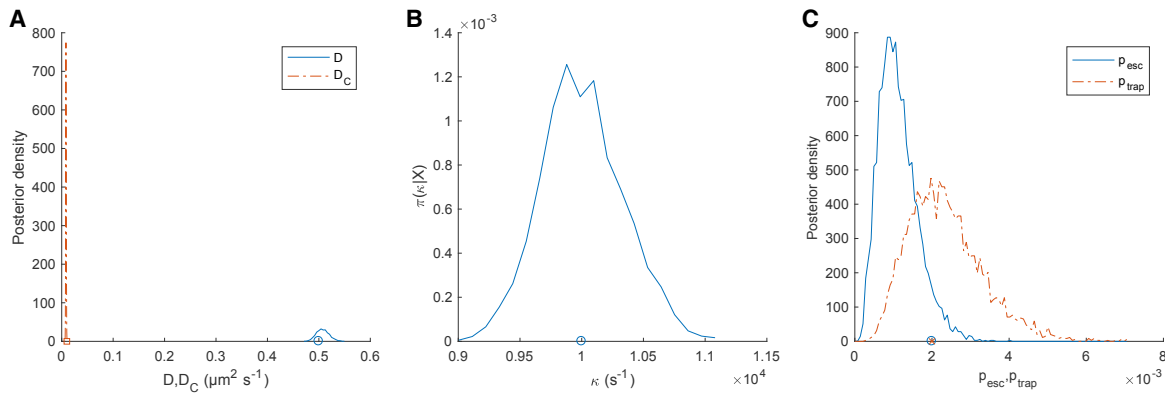


FIGURE 2 Posterior parameter distributions of the HPW model for a simulated trajectory. (A) The posterior distribution for D (blue, solid line) and D_C (red, dashed line), with simulation values indicated (circles) (B) The posterior for κ and simulation value (circle). (C) The posterior for p_{esc} (blue, solid line) and p_{trap} (red, dashed line), with simulation values (cross, circle respectively). The trajectory is as in Fig. 1. MCMC priors are as in Note S1 of the Supporting Materials and Methods. Corresponding MCMC runs are shown in Fig. S4. Data are based on the pooling of five independent chains of 2000 steps with a 1000-step burn-in. MCMC priors are as in Note S1 of the Supporting Materials and Methods. To see this figure in color, go online.

$0.0525 \pm 0.017 \mu\text{m}^2\text{s}^{-1}$ (mean \pm SE). This difference reflects the fact that MSD does not account for confinement, which is the dominant state, whereas our diffusion coefficient estimate does. Our Bayesian analysis provides estimates of parameter confidence per trajectory, which are in fact substantially smaller than the spread between trajec-

tories (Fig. S17); specifically, the ratio of the population variances of D and κ are 257 and 59 times larger than the average trajectory posterior variances, respectively. This indicates the presence of system variability, giving rise to trajectory heterogeneity. To understand its cause, we investigate whether heterogeneity is manifest in the

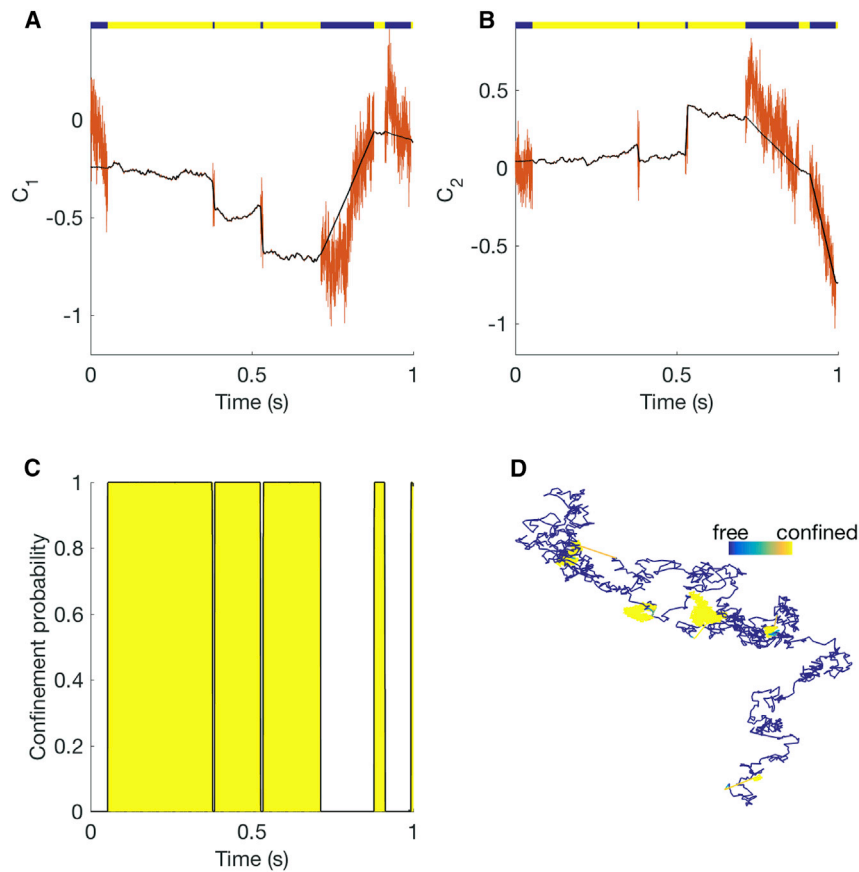


FIGURE 3 Hidden state inference for the HPW model for a simulated trajectory. (A and B) The mean inferred position of the harmonic potential center in x and y directions (black) and simulated (true) center (red). The colored line at the top represents the following particle states: free diffusion (blue) and confinement (yellow). (C) The inferred confinement probability (black line) and simulated (true) confinement state (yellow area). (D) The trajectory is colored by mean inferred confinement state, from $\pi(z_i|\mathbf{X}) = 0$ (blue, free) to $\pi(z_i|\mathbf{X}) = 1$ (yellow, confined). Color bar length, $0.1 \mu\text{m}$. MCMC is as in Fig. 2.

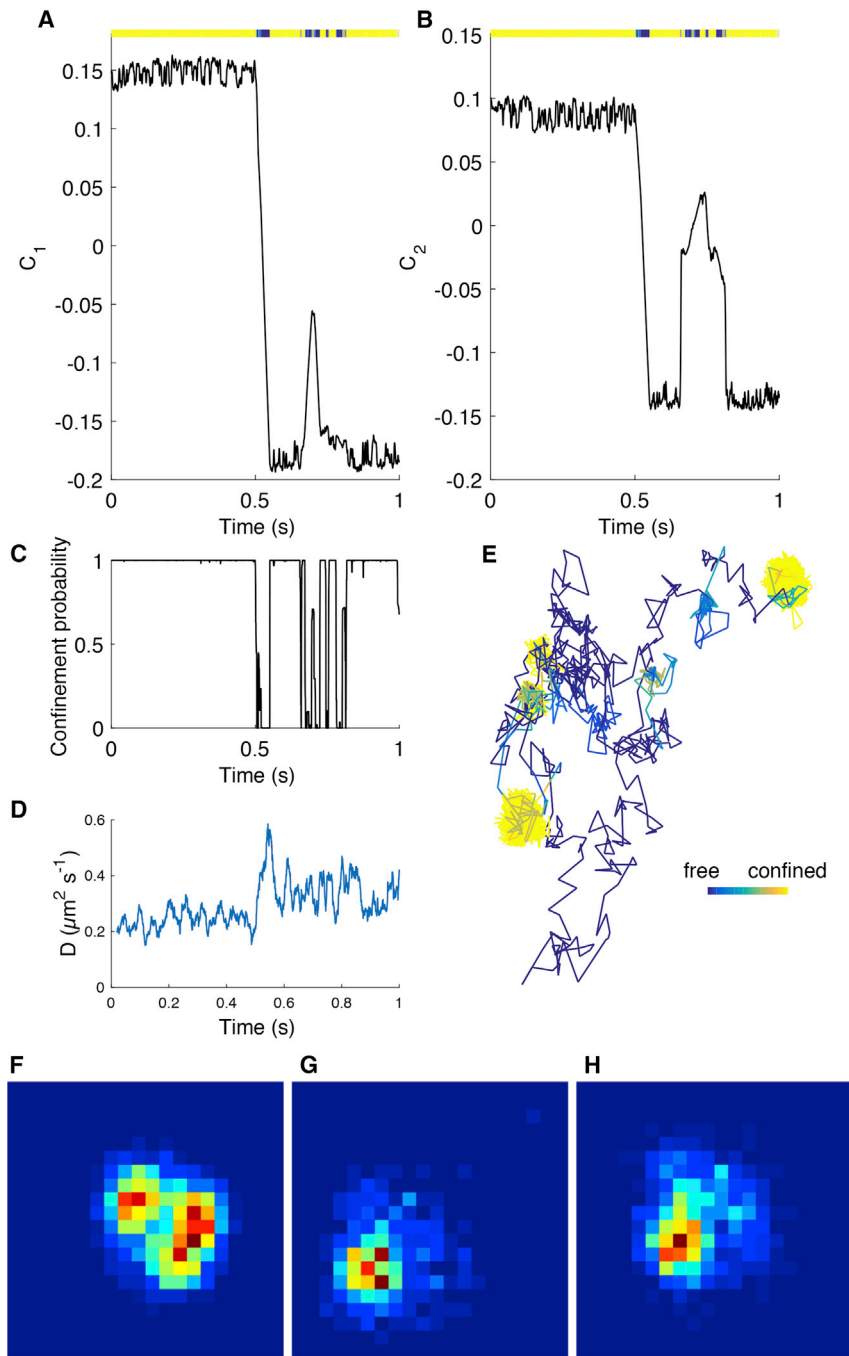


FIGURE 4 Hidden state inference for the HPW model applied to a 20-nm AuNP/CTxB/GM1 trajectory. (A and B) The mean inferred position of the HPW center C (x , y components) and upper colored bar representing $\pi(\mathbf{z}|\mathbf{X})$, (color scale goes from $\pi(z_i|\mathbf{X}) = 0$ (blue, free) to $\pi(z_i|\mathbf{X}) = 1$ (yellow, confined)). (C) The inferred mean confinement state and (D) moving average of local maximal likelihood diffusion coefficient estimate (window size: 100 subsampled frames). (E) The trajectory colored by mean inferred confinement state (color bar length, $0.1 \mu\text{m}$). (F–H) Density-colored 2D spatial histograms of confinement events. The two events in (G) and (H) are spatially collocated. Data is based on the pooling of 10 independent chains. MCMC priors and convergence criteria are as in Note S1 of the [Supporting Materials and Methods](#).

confinement events of individual trajectories, specifically the size, shape, and lifetime of these events.

Lifetime and shape analysis of confinement events

The mean confinement state lifetime (as defined in [Methods](#)) is 0.024 s , but there is a large variation in event lifetimes across trajectories ([Fig. 6, A and B](#)) and significant

heterogeneity across trajectories ($p = 0.02$, Kruskal-Wallis test, 1779 events across 60 trajectories). The lifetimes of free-diffusion events (mean 0.002 s) did not show significant heterogeneity across trajectories, ($p = 0.86$, Kruskal-Wallis test on 60 trajectories, 1770 events). Further, we examined if the population of lifetimes across trajectories conform to an exponential waiting time model, i.e., whether switching between states obey first-order kinetics. A quantile-quantile plot demonstrates that there is a distinct deviation from an

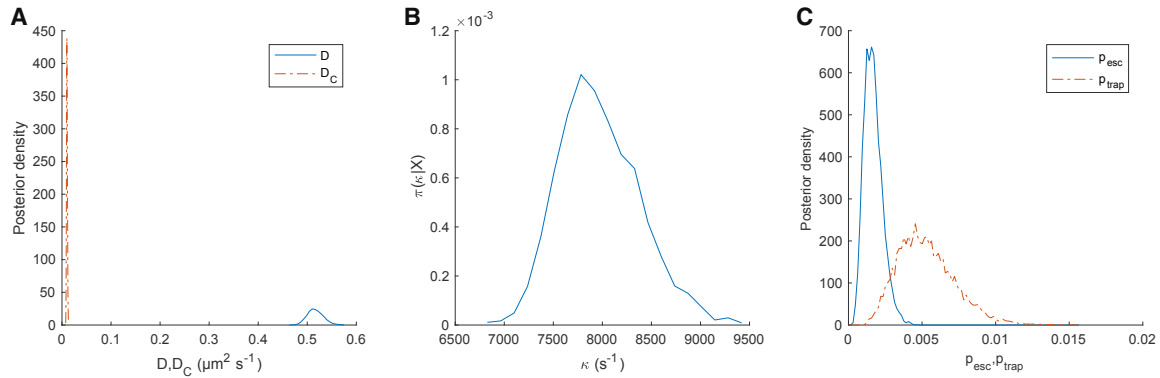


FIGURE 5 Posterior parameters of the HPW model applied to a 20-nm AuNP/CTxB/GM1 trajectory. (A) Posterior distributions for D (blue, solid line) and D_C (red, dashed line). (B) The posterior for κ . (C) The posteriors for p_{esc} (blue, solid line) and p_{trap} (red, dashed line). Distributions consist of samples pooled from 12 independent runs. Corresponding MCMC chains are shown in Fig. S14. Trajectory and MCMC runs are as in Fig. 4. To see this figure in color, go online.

exponential distribution fit (mean event time $\mu = 0.024$ s); specifically, there are a far higher proportion of longer trapping events, indicative of heterogeneity. A mixture of two exponentials is a better fit (Fig. 6 D) suggesting that the

confinement events derive from a heterogeneous population with at least two components with short and long average lifetimes. The minor population of long lifetime events are dispersed over trajectories (Fig. 6 B); in particular,

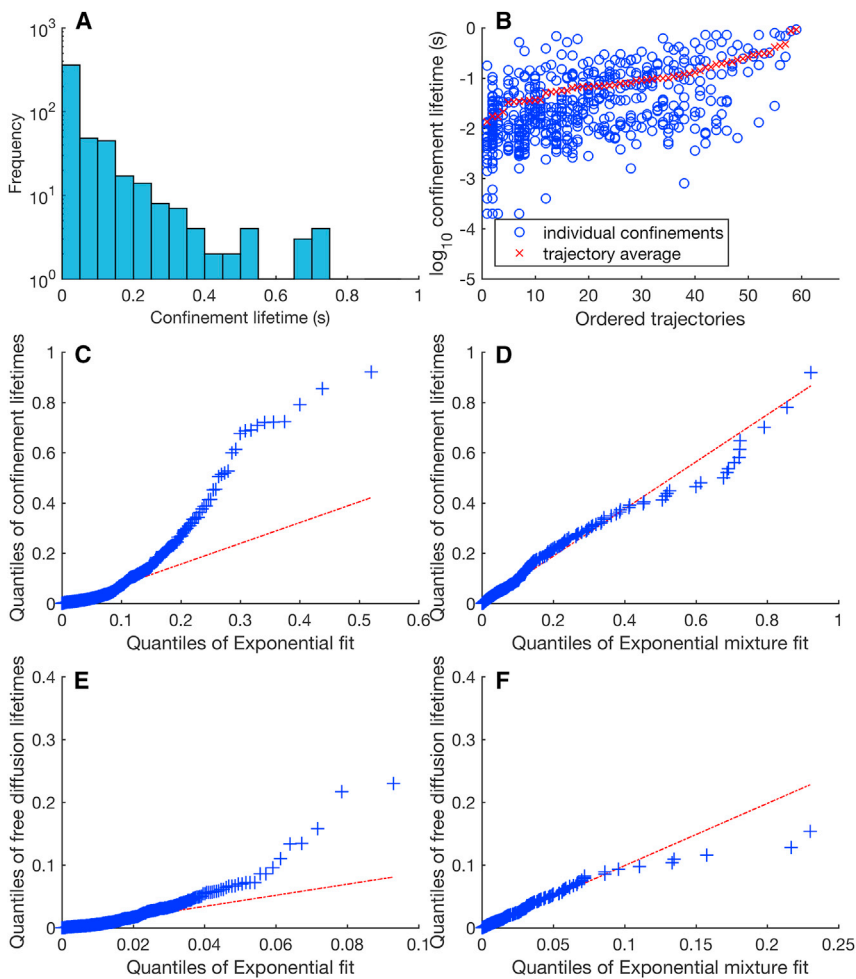


FIGURE 6 Confinement event lifetimes are not exponentially distributed. (A) A histogram of all confinement lifetimes ($n = 1959$ events). (B) A scatterplot of confinement lifetimes against trajectories ordered by mean confinement lifetimes. (C–F) Quantile-quantile plots of state lifetimes against exponential fits. (C) Confinement events against the exponential distribution ($\mu = 0.024$ s, $R^2 = 0$) and (D) confinement events against samples ($n = 10^4$) from a mixture of two exponentials ($\mu_1 = 0.004$ s, $\mu_2 = 0.1$ s; weights 0.80 and 0.20, respectively; $R^2 = 0.98$) are shown. (E) Free diffusion lifetimes ($n = 2011$ events) against the exponential distribution ($\mu = 0.002$ s, $R^2 = 0.98$) and (F) free diffusion lifetimes against samples ($n = 10^4$) from a mixture of two exponentials ($\mu_1 = 0.002$ s, $\mu_2 = 0.01$ s; weights 0.99 and 0.01, respectively; $R^2 = 0.98$) are shown. The red line is an extrapolated linear fit to the first and third quantiles. Plots include all confinement events except those that contained the trajectories' first or last time point. To see this figure in color, go online.

trajectories are not split into two groups with long and short mean confinement times. In contrast, the free diffusion state lifetimes closely follow a mono-exponential distribution (Fig. 6, E and F).

We next analyzed confinement event shape using the spatial statistics defined in Table 1. The mean confinement radius over all trajectories is 18 nm, comparable to the size of the AuNP, although estimator inflation is likely to be present (36). The mean radial skewness is 0.88; for comparison, a 2D Gaussian distribution gives a radial displacement (from the mean) that is Rayleigh distributed with skew $(2\sqrt{\pi}(\pi - 3)/(4 - \pi)^{3/2}) \approx 0.63$. Mean confinement radius and radial skewness show a wide distribution of values across confinement events (Fig. 7, A and B), with significant (one-way analysis of variance: mean confinement radius $p = 1.2 \times 10^{-6}$; radial skewness $p = 9.1 \times 10^{-4}$; 271 events grouped by 66 trajectories) heterogeneity across trajectories (Fig. 7, C and D). Confinement event spatial histograms for all 66 20-nm AuNP/CTxB/GM1 trajectories, ordered by the average within-trajectory mean

confinement radius (Fig. S18) and average radial skewness (Fig. S19) demonstrate the wide variety of confinement shapes.

The observed heterogeneity in confinement time, size, and (previously reported (4)) shape raises two key questions:

- 1) Does the shape of confinement events determine their lifetime?
- 2) Does heterogeneity predominately arise from a mechanism operating at individual confinement sites (local environment dependent) or at the level of trajectories (AuNP/CTxB/GM1 nanoparticle dependent)?

To probe the relationship between confinement event shape and lifetime (the first question), we examined their correlation. We found no correlation between confinement state lifetime and mean confinement radius (Fig. 7 E; Table S1) but a weak (and significant) negative correlation between lifetime and radial skewness (Fig. 7 F; Table S1). This suggests that the mixed-exponential nature of the

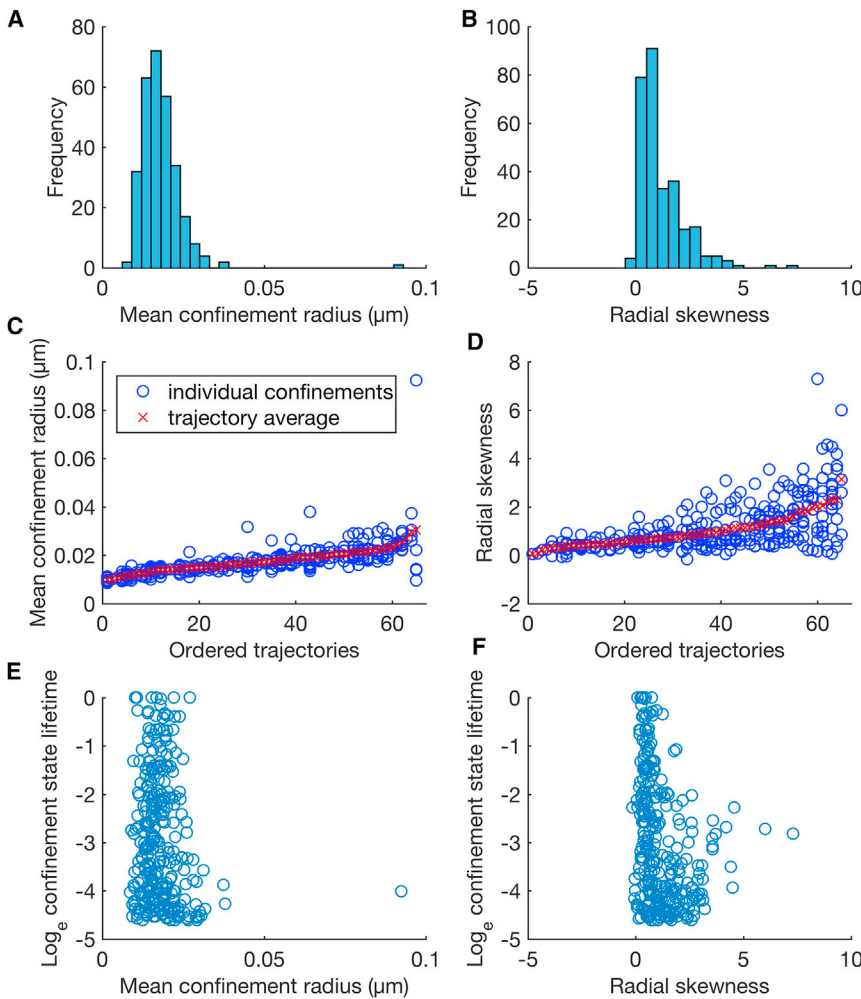


FIGURE 7 Shape statistics for confinement events in 20-nm AuNP/CTxB/GM1 trajectories. (A and B) Histograms over confinement events. (C and D) Spatial statistics for all confinement events, ordered by the average within trajectory statistic. Plots include all confinement events of at least 0.01 s, with events revisiting a previous trapping zone removed (giving 271 events). (E and F) Scatterplots of state lifetime against the given spatial statistic for the same 271 confinement events. To see this figure in color, go online.

binding lifetime is only weakly related to the shape of the binding event, i.e., these arise from different physical mechanisms.

Regarding the second question, the heterogeneity analysis above (see also Fig. 7, C and D) indicates that confinement events are statistically more similar within trajectories than across trajectories. Additionally, the ratio of the mean variance within trajectories to the variance across all events is 0.6 for both mean confinement radius and radial skewness (Table S2). To determine which confinement statistic is most strongly conserved within trajectories, we clustered events by each confinement statistic and quantified the similarity of events within single trajectories (Fig. 8). Confinement size is the most conserved, followed by confinement lifetime when excluding events revisiting a previous confinement zone (Fig. 8 A). Incorporating revisiting events dramatically improves the conservation of confinement size and shape statistics relative to lifetime (Fig. 8 B); this suggests that, although shape is conserved, confinement time is variable between events at the same location. This shape conservation at the same site is evident from the confinement event spatial histograms for long events (Fig. 9). Of note, the mean-median distance statistic failed to show significant heterogeneity across trajectories ($p = 0.06$, one-way analysis of variance) reflecting its lack of conservation when excluding events revisiting a previous confinement zone (Fig. 8 A) but was as conserved as lifetime when all events were included (Fig. 8 B). Thus, in answer to the second question, heterogeneity arises at both the trajectory- and confinement-site level, with different effects on confinement lifetime and shape. This suggest that nanoparticle confinement events are described by two degrees of freedom.

Analysis of 40-nm AuNP/CTxB/GM1 trajectories on mica

As a control, we analyzed 18 trajectories of 40-nm AuNP/CTxB/GM1 diffusing in supported lipid bilayers (SLBs) on a mica substrate. The previous analysis demonstrated that no confinement for this treatment was present (4). We applied our HPW model MCMC algorithm to this data and detected no confinement events (Fig. S20); the posterior confinement probability was <0.01 for all time, in all trajectories. The mean D was $1.2048 \pm 0.09 \mu\text{m}^2\text{s}^{-1}$ (mean \pm SE), comparable to 20-nm AuNP/CTxB/GM1 on glass trajectories ($1.15 \pm 0.20 \mu\text{m}^2\text{s}^{-1}$). The mean MSD-derived (with @msdalyzer (35)) D was $0.87 \pm 0.12 \mu\text{m}^2\text{s}^{-1}$. These values are in closer agreement than for the 20-nm AuNP/CTxB/GM1 on glass data set, which is expected because of the lack of confinement. However, we again observed trajectory heterogeneity in the diffusion coefficients with a ratio of population variance/mean trajectory variance of 123, indicative of individual AuNP-dependent diffusion coefficients.

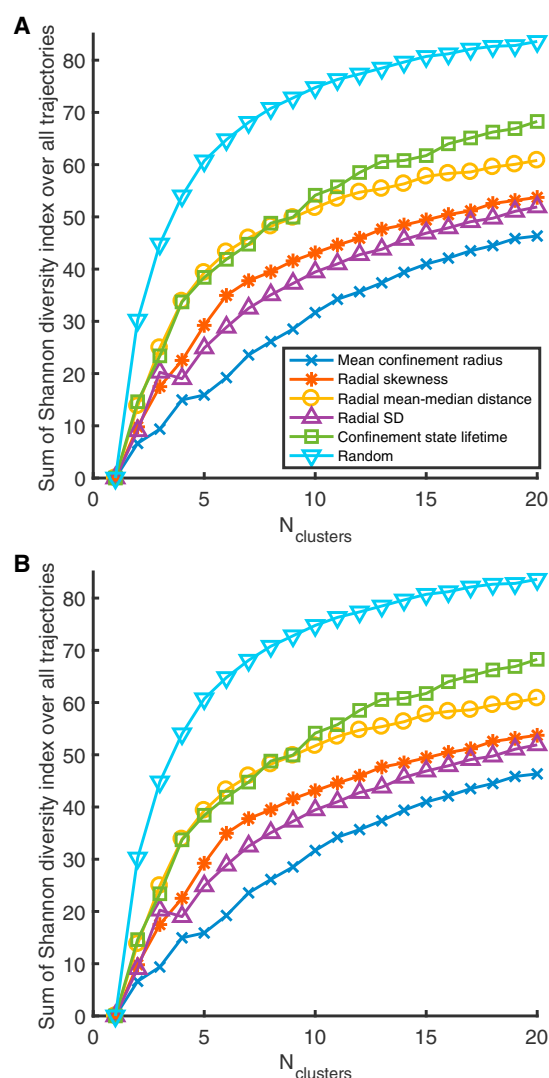


FIGURE 8 Clustering of confinement event statistics across individual confinement events were clustered (k-means++ algorithm (42) with squared Euclidean distance metric) based on event statistics. For each trajectory, l , the Shannon diversity index, $H_l = -\sum_{j=1}^{N_{\text{clusters}}} p_j \log p_j$, was calculated (p_j is the proportion of the events in trajectory l that appeared in cluster j). The sum of the Shannon diversity index over all trajectories is then a measure of the dissimilarity of events within trajectories (the lower the Shannon index the higher the similarity). The event statistics (shown as the legend) are defined in Table 1. For each choice of N_{clusters} , 50 separate clusterings were performed (because the k-means++ algorithm stochastically assigns initial values for cluster centroids), and the sum of the Shannon diversity index was averaged over these clusterings. (A) The clustering of events was obtained as described in “Confinement Event Profiling” in the main text, except with events containing the first or last time points excluded (214 events total). (B) The clustering of events is the same as (A), except with events that revisited a previous confinement zone included. To see this figure in color, go online.

DISCUSSION

We developed a Bayesian algorithm to infer an HPW confinement HMM and used it to partition SPT trajectories into periods of free diffusion and confinement. When

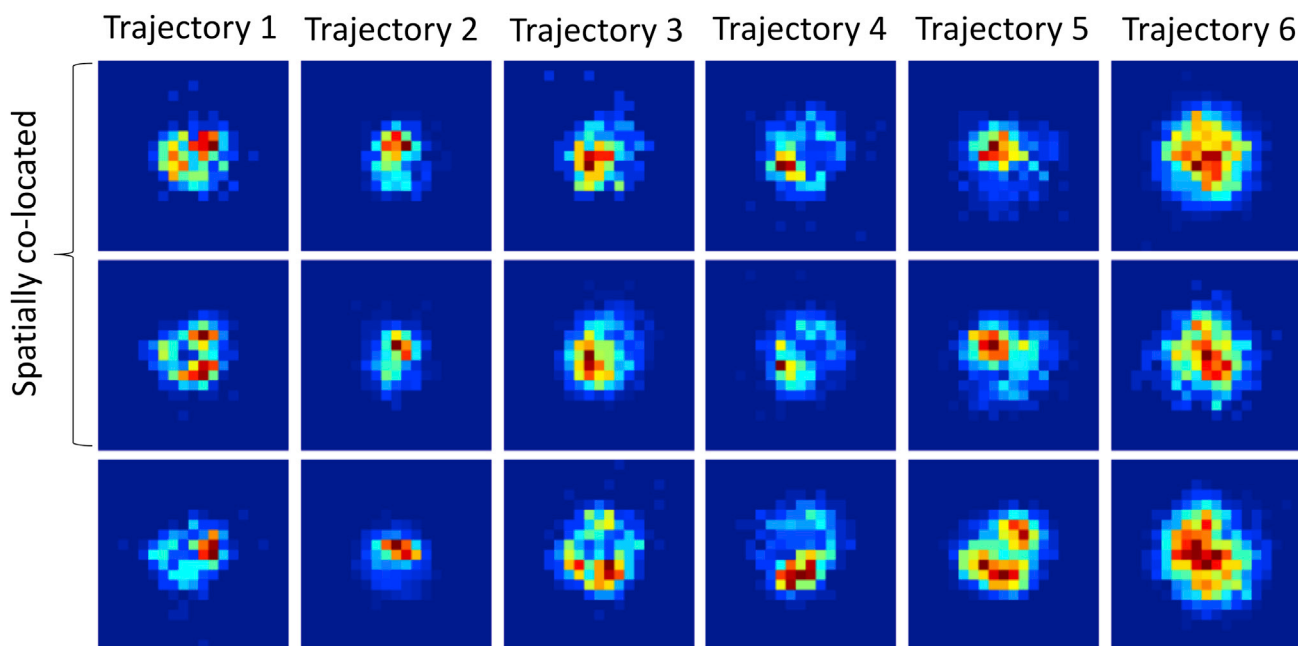


FIGURE 9 Spatial conservation of confinement events revisiting the same site. Each column shows particle position histograms for two spatially collocated confinement events and one event at a different location in the same trajectory. The spatially collocated confinement events are distinct, i.e., the particle moved away from the trapping zone between the displayed events. Each plot has a side length of $0.1 \mu\text{m}$.

applied to experimental AuNP/CTxB/GM1 trajectories, we detected clear periods of confinement and free diffusion (Fig. 4). It was previously proposed that confinement event shape heterogeneity (Gaussian versus non-Gaussian confinement) in this data set was due to transient multivalent binding of the tag (4). Our analysis of confinement events attained using the HPW model revealed the following heterogeneity trends:

- (1) Confinement size and shape are conserved within trajectories (Fig. 7, C and D), and repeat events at the same site show similarities (Figs. 5, G and H and 9).
- (2) Confinement event lifetimes are heterogeneous across trajectories and comprise a mixture of at least two exponentials with short (4 ms) and long (100 ms) mean lifetimes (Fig. 6, C and D).
- (3) Spatial heterogeneity and lifetime heterogeneity are effectively uncorrelated, suggesting they arise from different mechanisms.

Based on these observations, we propose a refinement to the transient multivalent tag-binding hypothesis. Namely, the characteristics of individual confinement events are determined by the following two factors: the size and geometry of the GM1 platform on the lower leaflet (determining residence times) and the number and distribution of CTxB complexes bound to the surface of the AuNP (determining size and shape of confinement event) (Fig. 1 B).

These dependencies are consistent with CTxBs remaining attached to the surface via GM1s throughout the entire trajectory (Fig. 10); this is supported by the high affinity of the

CTxB/GM1 bond with a dissociation rate in SLBs of $(2.8 \pm 0.1) \times 10^{-4} \text{ s}^{-1}$, giving a mean binding lifetime of $3.6 \times 10^4 \text{ s}$ (37). We propose that differences in the geometry of bound CTxB on the surface of the nanoparticle causes trajectory-conserved variation in the observed confinement as follows: tightly packed (or single) CTxBs have more freedom to “wobble” (Fig. 10 B), and broadly spaced, multiple (bound) CTxBs have less freedom (Fig. 10 C), giving a large, respectively small confinement radius for binding events. Additionally, non-Gaussian confinement events occur when there is a second (or potentially multiple) CTxB/GM1 attachment that is not immobilized, which restricts movement to a rotation or nonuniform “wobbling” around the immobilized binding site (Fig. 10 D). These hypotheses are consistent with the fact that there are around 25 CTxBs per 20-nm AuNP (4); it is expected that there will be variability in both their number and spatial distribution. We observe a strong correlation of the diffusion coefficient with the mean confinement radius ($r = 0.61$; Fig. 11 A), but not with the confinement event lifetime ($r = 0.27$; Fig. 11 B). This is consistent with the hypothesis that with more attachments, the AuNP experiences higher drag, whereas the mean confinement radius decreases because of stronger geometric constraints. Our analysis thus suggests that variation in the number and spatial configuration of bound CTxB contributes to the nature of the AuNP interaction with the upper leaflet of the bilayer, thereby giving each individual AuNP/CTxB/GM1 complex a confinement signature (Figs. 8, S18, and S19) and diffusion coefficient, with the latter also being evident in the confinement-free 40-nm AuNP data.

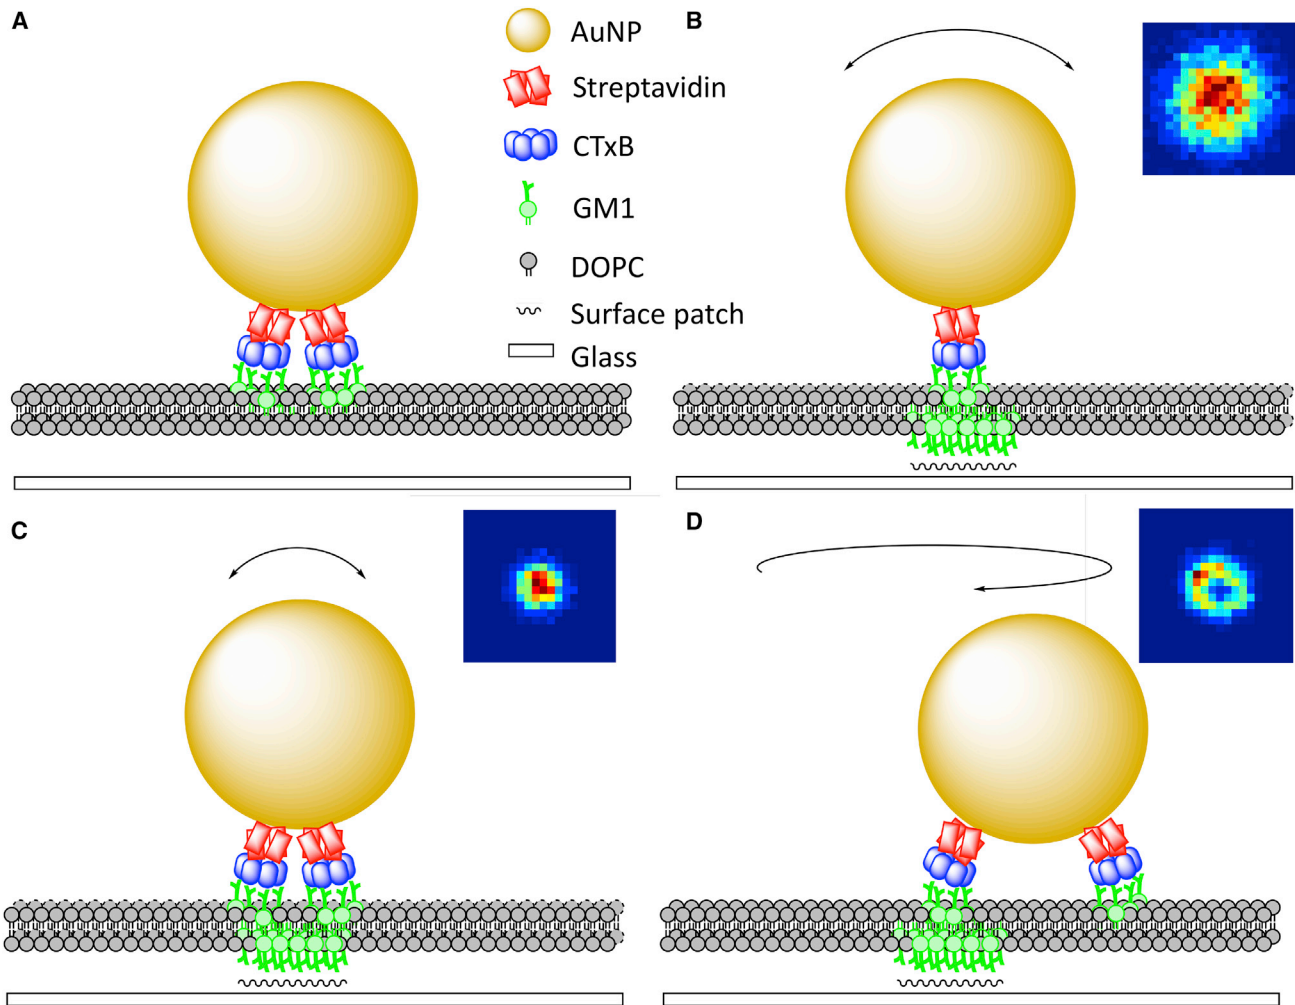


FIGURE 10 Schematic of AuNP/CTxB/GM1 structures leading to Gaussian and non-Gaussian confinement profiles. (A) Free diffusion, (B) wide Gaussian-like confinement, (C) narrow Gaussian-like confinement, and (D) non-Gaussian confinement. Insets in (B)–(D) are example histograms of particle positions pooled over confinement events within selected trajectories (e.g., Figs. S18 and S19). Insets have a side length of $0.1 \mu\text{m}$. The schematic is based on a figure in (4). To see this figure in color, go online.

The characteristics of the lower-leaflet GM1 platform contribute a confinement site dependence, with conservation of shape and size upon revisiting the same site (Figs. 8 and 9). On the other hand, confinement lifetime at the same site is variable. The GM1 in the lower leaflet is immobilized by hydroxyl pinning sites on the glass surface, with these sites having an estimated size of $<10 \text{ nm}$ (4). However, aggregation of GM1 with domain sizes of $15\text{--}60 \text{ nm}$ in SLBs has been observed in atomic force microscopy experiments (38). Large sites consist of more aggregated GM1 in the lower leaflet. Our mean confinement radius is 18 nm , which would comprise both AuNP/GM1/CTxB nanoparticle degrees of freedom around the binding site and displacements of the GM1 platforms between the leaflets. This suggests that either pinning sites are small, or no relative movement is possible. Larger pinning sites may trap multiple CTxB molecules on the AuNP, leading to more Gaussian behavior as rotational degrees of freedom

are lost and possibly longer (on average) trapping times. This could be the cause of the negative correlation between non-Gaussian confinement shape and event lifetime (the lower the radial skewness statistic, the longer the typical confinement time (Fig. 7 F)). However, the double exponential mixture distribution of confinement lifetimes cannot be explained by these mechanisms. Mean lifetime does not partition by trajectories (with long event lifetimes being distributed throughout the trajectories (Fig. 6 B)) suggesting a random process is responsible. The simplest explanation is that binding of the AuNP/CTxB/GM1 at the pinning sites is heterogeneous, e.g. there could be a multistep binding sequence with the second (long lifetime) step proceeding in only a fraction of the binding events. We note that the six trajectories that remain confined throughout are a third population, because even on the long-lifetime distribution, observing binding events of 1 s or longer is negligible (probability 4.5×10^{-5}).

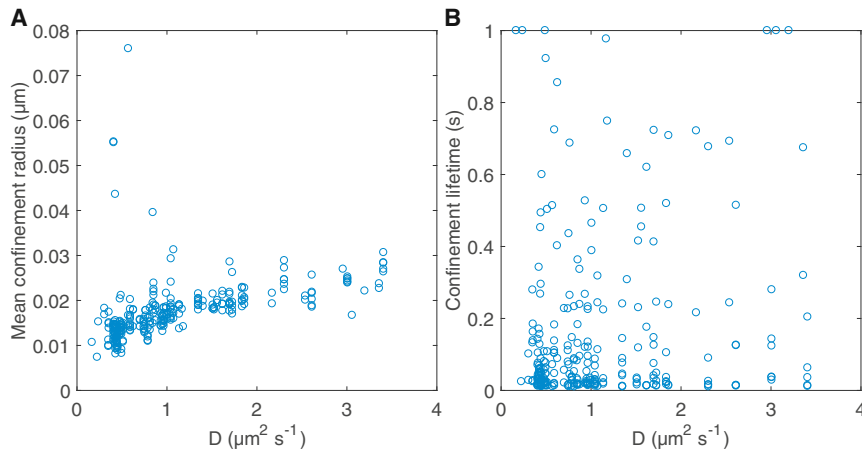


FIGURE 11 Correlation between D and confinement statistics. (A) The mean confinement radius and (B) confinement lifetime. Plots include all confinement events of at least 0.01 s. To see this figure in color, go online.

Outlook and future work

Analysis of SPT trajectories with HMMs has advantages over other methods for detecting confinement in single trajectories. In particular, they do not rely on tuning algorithm parameters through a comparison with Brownian motion. Additional parameters (such as the confinement strength κ , center C , and switching times as inferred here) can also be extracted, which allows for interpretation and comparison of confinement event characteristics across and within trajectories. However, appropriate HMMs are necessary for successful analysis. Specifically, models must approximate well the behavior of different dynamic states in the data. For instance, confinement is often associated with a decrease in the effective diffusion coefficient, suggesting that models that switch diffusivities (25–30) should also be able to detect confinement in these iSCAT particle trajectories. However, we found that a two-state diffusion coefficient switching HMM (30) could not segment these trajectories (data not shown). This implies that the effective diffusion coefficient of the AuNP/CTxB/GM1 complex does not change sufficiently during confinement events. The effective diffusion coefficient under confined Brownian motion is dependent on the temporal and spatial resolution of the data, suggesting that the high temporal sampling rate and positional accuracy of iSCAT data do not reduce the effective diffusion coefficient during confinement. In fact, the sampling timescale leads to frame-to-frame displacements that are on the order of the confinement radius, which required us to use OU dynamics in our MCMC algorithm. Failing to account for the effect of the frame-to-frame displacement/confinement size ratio on the integration scheme accuracy leads to a bias in the parameter estimates (data not shown).

This suggests that specifically modeling confinement, rather than modelling as a change in the diffusion coefficient, is necessary for data such as this. We believe the only other published confinement HMM for SPT analysis

is by Bernstein et al. (31). Our models are similar, with both switching between free diffusion and confinement in a HPW. However, there are key differences. Firstly, we incorporate diffusion of the harmonic well to relax the constraint of a circular potential. The importance of this will depend on the spatial-temporal resolution of the data and whether the confinement zone is static or has time-dependent shape variation or drift. For instance, transient confinement in lipid rafts, which both diffuse and have an irregular shape, has been hypothesized. Secondly, the inference frameworks are different; we use a Bayesian approach to determine full posterior distributions for parameters and hidden states per trajectory, thus quantifying the level of uncertainty in these estimates. Bernstein et al. (31) use a maximal likelihood approach inferring a point estimate, using a particle filter within an expectation-maximization algorithm. Using a Bayesian analysis was essential in our study to show that significant intertrajectory heterogeneity was present. Such information was key to identifying the AuNP signatures and separating environment effects from movement characteristics specific to individual AuNPs. In absence of this information, we would only be able to pool the trajectory estimates to define the population statistics. Furthermore, Bernstein et al. (31) infer multiple, static confinement zones separately, as opposed to our single, moving confinement-zone solution. The merits of these distinct approaches to the problem of multiple confinement zones in different experimental applications is a rich area for future work. Finally, we comment that we have implemented a “hop diffusion” model for multiple adjacent confinement zones using reversible-jump MCMC (39), which may be more appropriate for some systems.

As with all analysis, model accuracy needs to be balanced against computational complexity. In practice, this balance is also affected by the data; data at higher spatial and temporal resolution allows more detailed models to be inferred because subtle model differences

can then be distinguished. We demonstrated our algorithm on high-resolution iSCAT data, with static localization error 2.7 nm , mean particle diffusivity $1.15 \pm 0.106 \mu\text{m}^2\text{s}^{-1}$, and exposure time equal to 10^{-5} s (4). By subsampling the data at rate 10 (Figs. S1 and S2; Note S2 of the Supporting Materials and Methods), we mitigate any effects due to dynamic localization error. Under these conditions, confinement models and switching-diffusivity models can be distinguished. For other data types with lower localization accuracy and shorter trajectories, such as fluorescent probes, explicit consideration of measurement noise may be necessary. This is suggested by our simulations, in which we observed a bias in confinement strength κ as the measurement error increased beyond that of iSCAT data (Fig. S7). Our MCMC algorithm could be extended using an HMM with an unknown (hidden) particle position to model measurement noise as in (30); static and dynamic errors could also be separated (21,40).

The observation that individual AuNP/CTxB/GM1 complexes have a specific spatial signature means that distinguishing the effects of the tag from other factors, such as the cell membrane environment, is difficult. It follows that homogenous tags should improve characterization of the membrane environment. Because the variability in the tag signature presumably arises from the random placing of CTxB molecules on the AuNP surface, using particles with a structured surface is predicted to reduce or potentially eliminate this problem. Virions are ideal, given their highly geometric 3D structure. Interferometric label-free tracking of virions has been demonstrated at 3 s temporal resolution (41); thus, achieving the high spatial and temporal resolution of recent iSCAT (such as in the data set explored here) with viral particle tags is a distinct possibility. The resolution of the tag's signature and the length of events will then determine the resolution of that trajectory and thus the length scale to which SPT can discriminate different types of particle movement. Whether this can be taken below the size of the tag, reminiscent of super resolution, remains to be ascertained.

CONCLUSIONS

We use an HMM-based analysis to partition SPT trajectories into periods of free diffusion and confinement. Our algorithm infers the switching times between these two states, the diffusion coefficient D , and the characteristics of the following confinement events: the HPW strength κ , the position of the HPW center C , and the center diffusion coefficient D_C . We demonstrate the utility of the method on simulated and experimental data; on simulated data, confinement zones were accurately detected, and HPW centers accurately tracked while experimental trajectories were partitioned with high confidence. The model could potentially detect various biological phenomena such as lipid microdomains (or "rafts"), receptor clustering, and hop diffusion.

SUPPORTING MATERIAL

Supporting Materials and Methods, twenty figures, and two tables are available at [http://www.biophysj.org/biophysj/supplemental/S0006-3495\(18\)31056-7](http://www.biophysj.org/biophysj/supplemental/S0006-3495(18)31056-7).

AUTHOR CONTRIBUTIONS

Analyzed the data, P.J.S. and N.J.B.; Wrote the article, P.J.S. and N.J.B.; Developed the software, P.J.S.

ACKNOWLEDGMENTS

We thank Philipp Kukura for providing the data set for analysis and for helpful discussions and suggestions. We also thank Christian Eggeling and Gabrielle de Wit for helpful discussions and suggestions.

P.J.S. received a grant from the University of Warwick to study at Warwick Systems Biology Doctoral Training Centre funded by the Engineering and Physical Sciences Research Council and Biotechnology and Biological Sciences Research Council.

REFERENCES

- Manzo, C., and M. F. Garcia-Parajo. A review of progress in single particle tracking: from methods to biophysical insights. *Rep. Prog. Phys.* 78:124601.
- Shen, H., L. J. Tauzin, ..., C. F. Landes. Single particle tracking: from theory to biophysical applications. *Chem. Rev.* 117:7331–7376.
- Ortega-Arroyo, J., and P. Kukura. Interferometric scattering microscopy (iSCAT): new frontiers in ultrafast and ultrasensitive optical microscopy. *Phys. Chem. Chem. Phys.* 15:625.
- Spillane, K. M., J. Ortega-Arroyo, ..., P. Kukura. 2014. High-speed single-particle tracking of GM1 in model membranes reveals anomalous diffusion due to interleaflet coupling and molecular pinning. *Nano Lett.* 14:5390–5397.
- Hsieh, C. L., S. Spindler, ..., V. Sandoghdar. 2014. Tracking single particles on supported lipid membranes: multimobility diffusion and nanoscopic confinement. *J. Phys. Chem. B.* 118:1545–1554.
- Lin, Y. H., W. L. Chang, and C. L. Hsieh. 2014. Shot-noise limited localization of single 20 nm gold particles with nanometer spatial precision within microseconds. *Opt. Express.* 22:9159–9170.
- Clausen, M. P., and B. C. Lagerholm. 2011. The probe rules in single particle tracking. *Curr. Protein Pept. Sci.* 12:699–713.
- Qian, H., M. P. Sheetz, and E. L. Elson. 1991. Single particle tracking. Analysis of diffusion and flow in two-dimensional systems. *Biophys. J.* 60:910–921.
- Michalet, X. 2010. Mean square displacement analysis of single-particle trajectories with localization error: Brownian motion in an isotropic medium. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 82:041914.
- Michalet, X., and A. J. Berglund. 2012. Optimal diffusion coefficient estimation in single-particle tracking. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 85:061916.
- Monnier, N., S. M. Guo, ..., M. Bathe. 2012. Bayesian approach to MSD-based analysis of particle motion in live cells. *Biophys. J.* 103:616–626.
- Clausen, M. P., and B. C. Lagerholm. 2013. Visualization of plasma membrane compartmentalization by high-speed quantum dot tracking. *Nano Lett.* 13:2332–2337.
- Lagerholm, B. C., D. M. Andrade, ..., C. Eggeling. 2017. Convergence of lateral dynamic measurements in the plasma membrane of live cells

- from single particle tracking and STED-FCS. *J. Phys. D Appl. Phys.* 50:063001–063020.
14. Saxton, M. J. 1993. Lateral diffusion in an archipelago. Single-particle diffusion. *Biophys. J.* 64:1766–1780.
 15. Simson, R., E. D. Sheets, and K. Jacobson. 1995. Detection of temporary lateral confinement of membrane proteins using single-particle tracking analysis. *Biophys. J.* 69:989–993.
 16. Meilhac, N., L. Le Guyader, ..., N. Destainville. 2006. Detection of confinement and jumps in single-molecule membrane trajectories. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 73:011915.
 17. Rajani, V., G. Carrero, ..., C. W. Cairo. 2011. Analysis of molecular diffusion by first-passage time variance identifies the size of confinement zones. *Biophys. J.* 100:1463–1472.
 18. Coppola, S., G. Caracciolo, and T. Schmidt. 2014. Exact occupation probabilities for intermittent transport and application to image correlation spectroscopy. *New J. Phys.* 16:113057.
 19. Renner, M., L. Wang, ..., A. Triller. A simple and powerful analysis of lateral subdiffusion using single particle tracking. *Biophys. J.* 113:2452–2463.
 20. Koo, P. K., M. Weitzman, ..., S. G. Mochrie. Extracting diffusive states of Rho GTPase in live cells: towards in vivo biochemistry. *PLoS Comput. Biol.* 11:e1004297.
 21. Calderon, C. P., and K. Bloom. Inferring latent states and refining force estimates via hierarchical dirichlet process modeling in single particle tracking experiments. *PLoS One.* 10:e0137633.
 22. Koo, P. K., and S. G. Mochrie. Systems-level approach to uncovering diffusive states and their transitions from single-particle trajectories. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 94:052412.
 23. Wagner, T., A. Kroll, ..., M. Wiemann. Classification and segmentation of nanoparticle diffusion trajectories in cellular micro environments. *PLoS One.* 12:e0170165.
 24. Yin, S., N. Song, and H. Yang. Detection of velocity and diffusion coefficient change points in single-particle trajectories. *Biophys. J.* 115:217–229.
 25. Das, R., C. W. Cairo, and D. Coombs. 2009. A hidden Markov model for single particle tracks quantifies dynamic interactions between LFA-1 and the actin cytoskeleton. *PLoS Comput. Biol.* 5:e1000556.
 26. Chung, I., R. Akita, ..., I. Mellman. 2010. Spatial control of EGF receptor activation by reversible dimerization on living cells. *Nature.* 464:783–787.
 27. Ott, M., Y. Shai, and G. Haran. 2013. Single-particle tracking reveals switching of the HIV fusion peptide between two diffusive modes in membranes. *J. Phys. Chem. B.* 117:13308–13321.
 28. Persson, F., M. Lindén, ..., J. Elf. 2013. Extracting intracellular diffusive states and transition rates from single-molecule tracking data. *Nat. Methods.* 10:265–269.
 29. Monnier, N., Z. Barry, ..., M. Bathe. 2015. Inferring transient particle transport dynamics in live cells. *Nat. Methods.* 12:838–840.
 30. Slator, P. J., C. W. Cairo, and N. J. Burroughs. 2015. Detection of diffusion heterogeneity in single particle tracking trajectories using a hidden Markov model with measurement noise propagation. *PLoS One.* 10:e0140759.
 31. Bernstein, J., and J. Fricks. 2016. Analysis of single particle diffusion with transient binding using particle filtering. *J. Theor. Biol.* 401:109–121.
 32. Fox, E., E. B. Sudderth, ..., A. S. Willsky. 2011. Bayesian nonparametric inference of switching dynamic linear models. *IEEE Trans. Signal Process.* 59:1569–1585.
 33. Alvarez, I., J. Niemi, and M. Simpson. 2014. Bayesian inference for a covariance matrix. *Ann. Conf. on Applied Statistics in Agriculture.* 26:71–82.
 34. Gelman, A., J. B. Carlin, ..., D. B. Rubin. 2013. Bayesian Data Analysis, Third Edition. CRC Press, Boca Raton, FL.
 35. Tarantino, N., J. Y. Tinevez, ..., E. Laplantine. 2014. TNF and IL-1 exhibit distinct ubiquitin requirements for inducing NEMO-IKK supra-molecular structures. *J. Cell Biol.* 204:231–245.
 36. Churchman, L. S., H. Flyvbjerg, and J. A. Spudich. 2006. A non-Gaussian distribution quantifies distances measured with fluorescence localization techniques. *Biophys. J.* 90:668–671.
 37. Margheri, G., R. D'Agostino, ..., M. Del Rosso. 2014. The β -subunit of cholera toxin has a high affinity for ganglioside GM1 embedded into solid supported lipid membranes with a lipid raft-like composition. *Lipids.* 49:203–206.
 38. Shi, J., T. Yang, ..., P. S. Cremer. 2007. GM1 clustering inhibits cholera toxin binding in supported phospholipid membranes. *J. Am. Chem. Soc.* 129:5954–5961.
 39. Slator, P. J. 2015. Characterising cell membrane heterogeneity through analysis of particle trajectories. Ph.D. thesis (University of Warwick).
 40. Berglund, A. J. 2010. Statistics of camera-based single-particle tracking. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 82:011917.
 41. Ewers, H., V. Jacobsen, ..., V. Sandoghdar. 2007. Label-free optical detection and tracking of single virions bound to their receptors in supported membrane bilayers. *Nano Lett.* 7:2263–2266.
 42. Arthur, D., and S. Vassilvitskii. 2007. K-Means++: the advantages of careful seeding. Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete algorithms. 8:1027–1025.