

Systems-level approach to uncovering diffusive states and their transitions from single-particle trajectories

Peter K. Koo^{1,2} and Simon G. J. Mochrie^{1,2,3,*}

¹*Department of Physics, Yale University, New Haven, Connecticut 06520, USA*

²*Integrated Graduate Program in Physical and Engineering Biology, Yale University, New Haven, Connecticut 06520, USA*

³*Department of Applied Physics, Yale University, New Haven, Connecticut 06520, USA*

(Received 13 January 2016; published 28 November 2016)

The stochastic motions of a diffusing particle contain information concerning the particle's interactions with binding partners and with its local environment. However, an accurate determination of the underlying diffusive properties, beyond normal diffusion, has remained challenging when analyzing particle trajectories on an individual basis. Here, we introduce the maximum-likelihood estimator (MLE) for confined diffusion and fractional Brownian motion. We demonstrate that this MLE yields improved estimation over traditional mean-square displacement analyses. We also introduce a model selection scheme (that we call mleBIC) that classifies individual trajectories to a given diffusion mode. We demonstrate the statistical limitations of classification via mleBIC using simulated data. To overcome these limitations, we introduce a version of perturbation expectation-maximization (pEMv2), which simultaneously analyzes a collection of particle trajectories to uncover the system of interactions that give rise to unique normal and/or non-normal diffusive states within the population. We test and evaluate the performance of pEMv2 on various sets of simulated particle trajectories, which transition among several modes of normal and non-normal diffusion, highlighting the key considerations for employing this analysis methodology.

DOI: [10.1103/PhysRevE.94.052412](https://doi.org/10.1103/PhysRevE.94.052412)

I. INTRODUCTION

Single-particle tracking (SPT) provides the ability to probe noninvasively at subdiffraction-limit resolution the spatiotemporal motions of individual fluorescently labeled proteins (FPs) inside living cells. Because the different interactions that a FP undergoes inside a cell give rise to different types of diffusive motion, SPT data encode each FP's interactions with other particles and with its local environment: Biochemical binding interactions can lead to different diffusivities if the FP can bind to different substrates [1], while interactions with the cellular medium can give rise to anomalous diffusion [2–4] or can lead to confined motions [5,6]. Thus, important goals of SPT measurements are (i) to infer these interactions from an analysis of protein trajectories, and (ii) to determine the spatiotemporal kinetics of each interaction.

To uncover the desired information, we must infer from the observed protein trajectories the number of unique diffusive states, each state's diffusion mode, and its diffusion properties, along with which portions of each trajectory correspond to a given diffusive state. Successfully accomplishing this task would allow for the determination of the underlying transition kinetics and the spatiotemporal locations of particular diffusive states and their transitions within the cell.

Previous work [7–16] that seeks to assess dynamic heterogeneity in tracking data has been reviewed by us in Ref. [17]. The traditional approach for analyzing the diffusive properties of individual particle trajectories is by fitting each trajectory's time-averaged mean-square displacement (taMSD) to a corresponding diffusion model [18]. However, the way in which the taMSD is usually calculated results in a statistically complex representation of the underlying diffusion

process, especially for short trajectories (see the Supplemental Material [19]), rendering the taMSD unreliable. Thus, an unweighted least-squares regression against the taMSD yields a statistically inefficient estimation of the diffusion model parameters. Improved estimation can be achieved by analyzing longer trajectories, albeit the same interaction must persist throughout the duration of the trajectory, which is an increasingly unlikely condition in the complex environment inside living cells. Alternatively, ensemble-averaging taMSD curves across particle trajectories, which share the same underlying diffusive properties, are another route for bolstering the statistics and thus better representing the underlying diffusive behavior. However, because the diffusive properties of each trajectory are not known *a priori*, how to sort trajectories into groups that share diffusive properties, and therefore may be averaged together, is not straightforward.

Because of the drawbacks of taMSD analysis, a number of alternatives have emerged for determining diffusion parameters, namely the maximum-likelihood estimator (MLE) [20], optimal least-squares fitting (OLSF) [21], and the covariance-based estimator (CVE) [22]. These approaches have demonstrated improved estimation in comparison with traditional taMSD analysis. Importantly, however, to date these approaches, which do properly account for localization noise sources, have only been shown to be applicable to particle trajectories undergoing normal diffusion.

Recently, systems-level analyses, namely variational Bayes single-particle tracking (vbSPT) [1] and perturbation expectation-maximization (pEM) [17], have demonstrated that the limited statistics of individual particle trajectories can be augmented by simultaneously analyzing a population of particle trajectories to uncover the number of unique diffusive states and their corresponding diffusive properties. However, both of these methods have their own limitations. While vbSPT allows for transitions between different diffusive

*simon.mochrie@yale.edu

states, it fails to properly account for experimental noise sources, compromising vbSPT's ability to reliably extract the correct number of diffusive states and each state's diffusive properties in some situations [17]. On the other hand, while pEM properly accounts for experimental noise sources, it assumes that diffusive properties are constant throughout the duration of each trajectory. Thus, pEM is only suitable to analyze tracks sampled at sufficiently short time scales that transitions between different diffusive states may be neglected. In addition, both methods make a short-time diffusion approximation, thereby effectively assuming that every particle trajectory undergoes normal diffusion. In fact, however, diffusing proteins interact with the complex environment in living cells, which can lead to non-normal diffusive behavior, including, for example, confined diffusion within focal adhesions [6] and membrane corals [5], in which a labeled protein is tethered to a particular fixed location within a cell, and subdiffusive behavior in the bacterial cytoplasm [2–4], which may be the result of the complex viscoelastic properties of this medium [2]. Thus, the short-time diffusion approximation made by pEM and vbSPT does not necessarily hold on experimentally relevant time scales.

In the present paper, we present an overall methodology comprising a number of advances that overcome these limitations: First, we extend Berglund's MLE framework to determine the diffusion parameters for canonical modes of non-normal diffusion, namely confined diffusion and fractional Brownian motion (fBm). Second, we introduce a model selection scheme, that we term mleBIC, which classifies individual trajectories to a given diffusion model. Third, we extend the pEM framework to be able to uncover non-normal diffusion modes and transitions between different diffusive states within particle trajectories. We also give empirical guidelines for the sort of data likely to be necessary to successfully apply our methodology.

Specifically, in Sec. II A, we demonstrate the improved performance of MLE against traditional taMSD analysis on various sets of simulated particle trajectories undergoing non-normal diffusion across a wide parameter spectrum. Since the diffusion mode of each experimental particle track is not known *a priori*, in Sec. II B we introduce a model selection scheme, based on the Bayesian information criterion (BIC), that we call mleBIC, for classifying individual trajectories to a given diffusion model. By applying mleBIC to synthetic trajectories undergoing various modes of non-normal diffusion, both without and with localization noise, we illustrate, by example, the statistical limits of mleBIC's classification. In general, we find that, even though MLE estimation is quite reliable for determining diffusion parameters, classification to determine the correct underlying diffusion model depends strongly on the length of the trajectory, and it is only accurate for sufficiently long trajectories. Moreover, resolving the level of heterogeneity within a population of trajectories that realize different diffusion modes remains challenging. Consequently, the SPT analysis goals defined above—specifically, uncovering the number of diffusive states and their properties and transitions—cannot generally and reliably be achieved from an analysis that treats individual particle trajectories independently.

Therefore, in Sec. II C, we turn to a systems-level analysis: We present a major extension of the pEM framework, that we call pEM version 2 (pEMv2 [23]), that seeks to uncover the system of diffusive behaviors arising from distinct physical interactions by (i) identifying the number of unique diffusive states (normal or non-normal diffusion modes), (ii) determining the diffusive properties of each diffusive state, and (iii) classifying individual trajectories to particular diffusive states to reveal the spatiotemporal dynamics of each diffusive behavior in reference to the cell. In addition to now being applicable to non-normal modes of diffusion, importantly, pEMv2 eases the other important constraint on pEM, namely that the diffusive state remain the same throughout the trajectory. It accomplishes this by splitting long trajectories into equally sized bins of smaller trajectories, thus enabling transitions between different diffusive states to be accounted for. We test the performance of pEMv2 on various sets of synthetic particle trajectories to gain better intuition concerning its capabilities and limitations in reference to the free parameters in the analysis. We show that in many cases, pEMv2 is indeed able to uncover and characterize normal/non-normal diffusion modes and the transitions between them. Thus, pEMv2 represents a powerful analysis tool for accurately characterizing the interactions of diffusing proteins in live cells, and it brings us a major step closer to being able to understand spatiotemporal biochemistry inside living cells via SPT.

II. RESULTS AND DISCUSSION

A. Maximum-likelihood framework

The one-dimensional (1D) stochastic increments of a diffusing particle undergoing a stationary Gaussian process are given according to [25]

$$x(i+1) = x(i) + \Sigma(i,j)^{1/2} W(j), \quad (1)$$

where $x(i)$ is the x coordinate of the particle's position at time step i , $W(j)$ is a standard Brownian motion with the properties $\langle W(j) \rangle = 0$ and $\langle W(i), W(j) \rangle = \delta_{i,j}$, where $\delta_{i,j}$ is the Kronecker delta, and $\Sigma(i,j)$ is the covariance matrix of the particle's x -displacements at time steps i and j . Equation (1) employs the Einstein summation convention in which a sum over j is implied. Throughout this paper, we use parentheses to indicate the elements of a vector or matrix. Additional dimensions, such as a track index, are indicated using a subscript.

It follows from Eq. (1) that the likelihood function, $P(\Delta\mathbf{x}|\Sigma)$, is given by a multivariate Gaussian distribution according to

$$P(\Delta\mathbf{x}|\Sigma) = \frac{1}{(2\pi)^{N/2} |\Sigma|^{1/2}} \exp \left[-\frac{1}{2} \Delta\mathbf{x}^T \Sigma^{-1} \Delta\mathbf{x} \right], \quad (2)$$

where $\Delta\mathbf{x}$ represents the vector of the N particle track displacements, $\{\Delta\mathbf{x}(n)\}_{n=1}^N$, and $\Delta\mathbf{x}^T$ is its transpose. $|\Sigma|$ is the determinant of the covariance matrix, and Σ^{-1} is its inverse. Equation (2) is the likelihood function that we seek to maximize. The dependence of the covariance matrix in Eq. (2) on model parameters for several canonical modes of diffusion is given in Table I.

For normal diffusion, the presence of experimental noise sources, namely static localization noise, which is the

TABLE I. Analytical covariance matrix of particle track displacements separated in time by Δt for canonical diffusion models, namely normal diffusion, confined diffusion, fractional Brownian motion, and an immobile model *with* localization noise corrections, assuming that the camera exposure time equals the frame duration, Δt . D is the diffusion coefficient, L is the confinement size for confined diffusion, and α is the anomalous exponent for fBm.

Mode	Covariance matrix (μm^2)
Normal	$\Sigma^{\text{normal}}(i, j) = \begin{cases} 2D\Delta t + 2\sigma^2 - \frac{2}{3}D\Delta t, & j = i \\ -\sigma^2 + \frac{1}{3}D\Delta t, & j = i \pm 1 \\ 0 & \text{otherwise} \end{cases}$
Confined	$\Sigma^{\text{confined}}(i, j) = \tilde{\Sigma}^{\text{confined}}(i, j) + \begin{cases} 2\sigma^2 - \frac{1}{6}(2\tilde{\Sigma}^{\text{confined}}(i, i) - 2\tilde{\Sigma}^{\text{confined}}(i, i + 1)), & j = i \\ -\sigma^2 - \frac{1}{6}(2\tilde{\Sigma}^{\text{confined}}(i, j) - \tilde{\Sigma}^{\text{confined}}(i, j - 1) - \tilde{\Sigma}^{\text{confined}}(i, j + 1)), & j = i \pm 1 \\ -\frac{1}{6}(2\tilde{\Sigma}^{\text{confined}}(i, j) - \tilde{\Sigma}^{\text{confined}}(i, j - 1) - \tilde{\Sigma}^{\text{confined}}(i, j + 1)) & \text{otherwise,} \end{cases}$ <p>where</p> $\tilde{\Sigma}^{\text{confined}}(i, j) = \begin{cases} \frac{L^2}{6} - \frac{16L^2}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} \Phi(1), & j = i \\ \frac{-L^2}{12} + \frac{8L^2}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} \Phi(1)[2 - \Phi(1)], & j = i \pm 1 \\ \frac{8}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} [-2\Phi(j - i + 1) + \Phi(j - i) + \Phi(j - i + 2)] & \text{otherwise,} \end{cases}$ <p>where $\Phi(n) = \exp\left[-\left(\frac{k\pi}{L}\right)^2 Dn\Delta t\right]$</p>
fBm [24]	$\Sigma^{\text{fBm}}(i, j) = \begin{cases} \frac{2D\Delta t^\alpha}{(\alpha+2)(\alpha+1)} [A(1) - 2] + 2\sigma^2, & j = i \\ \frac{D\Delta t^\alpha}{(\alpha+2)(\alpha+1)} [A(2) - 2A(1) + 2] - \sigma^2, & j = i \pm 1 \\ \frac{D\Delta t^\alpha}{(\alpha+2)(\alpha+1)} [A(j - i + 1) - 2A(j - i) + A(j - i - 1)] & \text{otherwise,} \end{cases}$ <p>where $A(n) = (n + 1)^{\alpha+2} + (n - 1)^{\alpha+2} - 2n^{\alpha+2}$</p>
Immobile	$\Sigma^{\text{immobile}}(i, j) = \begin{cases} 2\sigma^2, & j = i \\ -\sigma^2, & j = i \pm 1 \\ 0 & \text{otherwise} \end{cases}$

uncertainty due to a finite number of photons emitted from a fluorophore during a camera's exposure time, and dynamic localization noise, which is the uncertainty caused by the motions of the fluorophore during a camera's exposure time [26], has been shown to contribute nearest-neighbor covariance terms [20]. In the Supplemental Material [19], these calculations are extended to incorporate static localization noise into the covariance terms for non-normal diffusion with the result that

$$\Sigma^{\text{static}}(i, j) = \begin{cases} 2\sigma^2, & j = i, \\ -\sigma^2, & j = i \pm 1, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Assuming that the camera exposure time equals Δt , which is the usual situation in SPT measurements, the dynamic localization noise contribution to the covariance matrix for normal diffusion and confined diffusion may be shown to be given approximately by

$$\Sigma^{\text{dynamic}}(i, j) \approx -\frac{1}{6}(2\tilde{\Sigma}(i, j) - \tilde{\Sigma}(i + 1, j) - \tilde{\Sigma}(i, j + 1)), \quad (4)$$

where $\tilde{\Sigma}$ is the covariance matrix in the absence of noise (Appendix). A derivation of Eq. (4) is given in the Supplemental Material [19]. For fBm, the contribution of dynamic localization noise to the covariance matrix is derived in Ref. [24]. As also shown in the Supplemental Material [19], corrections for static localization noise, Σ^{static} , and dynamic localization

noise, Σ^{dynamic} , contribute additively to the covariance matrix:

$$\Sigma = \tilde{\Sigma} + \Sigma^{\text{static}} + \Sigma^{\text{dynamic}}. \quad (5)$$

Analytical results for the covariance matrix, incorporating localization noise corrections, for three canonical modes of diffusion, including an immobile particle model, are given in Table I. The likelihood function is maximized numerically as described in Sec. IV C.

To validate the performance of our maximum-likelihood framework, we generated various sets of synthetic particle trajectories, corresponding to different modes of diffusion, as described in Sec. IV. For confined diffusion, trajectories were simulated with a number of confinement sizes from 0.25 to 5 μm . For fBm, trajectories were simulated with a number of anomalous exponents from 0.25 to 1.75. For confined diffusion, the trajectories were simulated with a diffusion coefficient of $D_{\text{sim}} = 0.3 \mu\text{m}^2 \text{s}^{-1}$. For fBm, the trajectories were simulated with a "diffusion coefficient" of $D_{\text{sim}} = 0.3 \mu\text{m}^2 \text{s}^{-\alpha}$. Dynamic localization noise was added by first simulating particle positions separated by "micro" time steps of $\delta t = \Delta t/32$, and then by averaging blocks of 32 of these positions together to produce positions separated by time steps of $\Delta t = 32 \text{ ms}$. The net effect is to mimic experimental motion-blurred positions, corresponding to a camera exposure time equal to the frame duration of $\Delta t = 32 \text{ ms}$. Static localization noise was included by adding a normally distributed random number with zero mean and variance, σ_{sim}^2 , to each motion-blurred position, where $\sigma_{\text{sim}} = 0.04 \mu\text{m}$ (Sec. IV). For

each set of diffusion parameters, we generated sets of particle trajectories with track lengths $N = \{30, 60, 120, 240\}$ steps. To maintain the same level of positional information across all sets of synthetic particle trajectories, the total number of particle positions across each simulation set was constant at 12 000 total steps.

We have compared the performance of MLE and taMSD analyses using synthetic particle trajectories both without localization noise (Figs. S1 and S2) and with localization noise (Figs. S3 and S4). The detailed procedures involved in the MLE analysis and the taMSD analysis are given in Sec. IV. For confined diffusion (Figs. S1 and S3), MLE outperforms taMSD. Even though both the MLE and the taMSD diffusivity estimates exhibit a positive bias in their estimations of the diffusion coefficient, both the bias and the error are significantly less for MLE than for taMSD, especially in the presence of localization noise. When analyzing synthetic particle trajectories with localization noise, taMSD-based estimates of the confinement length are erratic. By contrast, even with localization noise, MLE yields reasonable confinement size estimates, provided the reduced confinement size ($L_{\text{reduced}} = \frac{L}{\sqrt{12D\Delta t}}$) is sufficiently small. As could be expected, the range of reduced confinement sizes for which MLE provides reasonable estimates increases with increasing track length, because the increased errors for larger reduced confinement sizes are associated with each particle's limited sampling of its confinement, which is inevitable for short tracks. For the MLE analyses, the static localization noise estimate was slightly negatively biased with a decreasing bias for increasing track length.

For fBm (Figs. S2 and S4), MLE is superior to taMSD. In this case, MLE and taMSD estimates both appear unbiased when particle tracks do not contain localization noise. However, the MLE estimates show noticeably lower errors. In the presence of localization noise, both MLE and taMSD estimates for the diffusivity, anomalous exponent, and static localization noise become biased. However, both the bias and the error are considerably less for MLE than for taMSD. As expected, the bias and the error are reduced for longer trajectories both without and with localization noise.

These collected results unambiguously demonstrate that MLE improves upon taMSD estimates for non-normal diffusion modes, in each case reliably characterizing the underlying diffusion model over a broader range of parameter space. They also emphasize that the presence of static localization noise reduces the quality of both taMSD- and MLE-based estimation, and in some cases it may introduce a bias, underscoring the importance of properly incorporating the effect of localization noise. As expected, bias and errors are reduced for longer (but fewer) individual trajectories, even for a fixed total number of time steps.

B. Performance of Bayesian model selection to classify individual particle trajectories

For experimental particle trajectories, the underlying mode of diffusion is in general unknown *a priori*. Therefore, some criterion must be imposed to select the best model, i.e., to statistically assess which diffusion model best describes any given particle trajectory.

According to Bayesian model selection, classification can be made by inferring the probability of the k th diffusion model, \mathcal{M}_k , from a trajectory, $P(\mathcal{M}_k|\Delta\mathbf{x})$, where $\Delta\mathbf{x}$ represents the vector of displacements from a particle trajectory. According to Bayes' rule, the probability of diffusion model k is as follows: $P(\mathcal{M}_k|\Delta\mathbf{x}) = \frac{P(\Delta\mathbf{x}|\mathcal{M}_k)P(\mathcal{M}_k)}{P(\Delta\mathbf{x})}$, where $P(\mathcal{M}_k)$ is the model prior, $P(\Delta\mathbf{x})$ may be viewed as a normalization constant, given by $\sum_{i \in \mathcal{M}} P(\Delta\mathbf{x}|\mathcal{M}_i)P(\mathcal{M}_i)$, and $P(\Delta\mathbf{x}|\mathcal{M}_k)$ is the model evidence given by $P(\Delta\mathbf{x}|\mathcal{M}_k) = \int P(\Delta\mathbf{x}|\theta, \mathcal{M}_k)P(\theta|\mathcal{M}_k)d\theta$, where $P(\Delta\mathbf{x}|\theta, \mathcal{M}_k)$ is the likelihood distribution and $P(\theta|\mathcal{M}_k)$ is the prior distribution of the parameters, θ , of model k .

Although the model prior, $P(\mathcal{M}_k)$, may be specified to express a preference for a particular model, we elect to take an agnostic approach and assume that all diffusion models are equally probable. In this manner, the model evidence is the only term of interest as the normalization absorbs all other contributions. However, the priors of each model, $P(\theta|\mathcal{M}_k)$, may introduce a bias that becomes more pronounced when the peak of the likelihood distribution is not sharp. A representative likelihood distribution for confined diffusion (Fig. S5) and fBm (Fig. S6), calculated using simulated data for various track lengths, demonstrates that the likelihood distribution for these non-normal diffusion modes is indeed broad near its global maximum.

To minimize the influence from priors, we employ a Laplace approximation to the model evidence and assume a broad multivariate Gaussian prior with a full rank covariance matrix, which leads via standard manipulations to the Bayesian information criterion (\mathcal{B}) given according to [27,28]

$$\mathcal{B} = \ln P(\Delta\mathbf{x}|\mathcal{M}_k) = \ln P(\Delta\mathbf{x}|\hat{\theta}, \mathcal{M}_k) - \frac{N_{\text{params}}}{2} \ln M, \quad (6)$$

where $\hat{\theta}$ are the maximum-likelihood parameters of model k , N_{params} is the number of free parameters, and M is the number of particle track displacements.

In summary, for a given trajectory, the MLE is found for each candidate diffusion model, according Sec. IV, yielding the parameter estimates and ln-likelihood value, from which the BIC can be calculated [Eq. (6)]. The model probability for each diffusion model can be subsequently calculated according to

$$P(\mathcal{M}_k|\Delta\mathbf{x}) = \frac{\exp(\mathcal{B}_k - \bar{\mathcal{B}})}{\sum_{i=1}^K \exp(\mathcal{B}_i - \bar{\mathcal{B}})}, \quad (7)$$

where $\bar{\mathcal{B}}$ is the maximum BIC value across K candidate diffusion models. Thus, classification is determined by the diffusion model that yields the highest model probability. Henceforth, this analysis pipeline is referred to as mleBIC.

To understand the statistical limits of classification under ideal circumstances, namely particles that have constant diffusion properties throughout the duration of their trajectories, we employed mleBIC across various sets of synthetic particle trajectories *with* static and dynamic localization noise for each canonical diffusion mode (Fig. 1). For short particle trajectories undergoing normal diffusion [Fig. 1(a)], a normal diffusion model was favored with a high probability when $D_{\text{reduced}} = \frac{D\Delta t}{\sigma^2} > 1$. When $D_{\text{reduced}} < 1$, the underlying static localization noise dominates the underlying diffusion, which

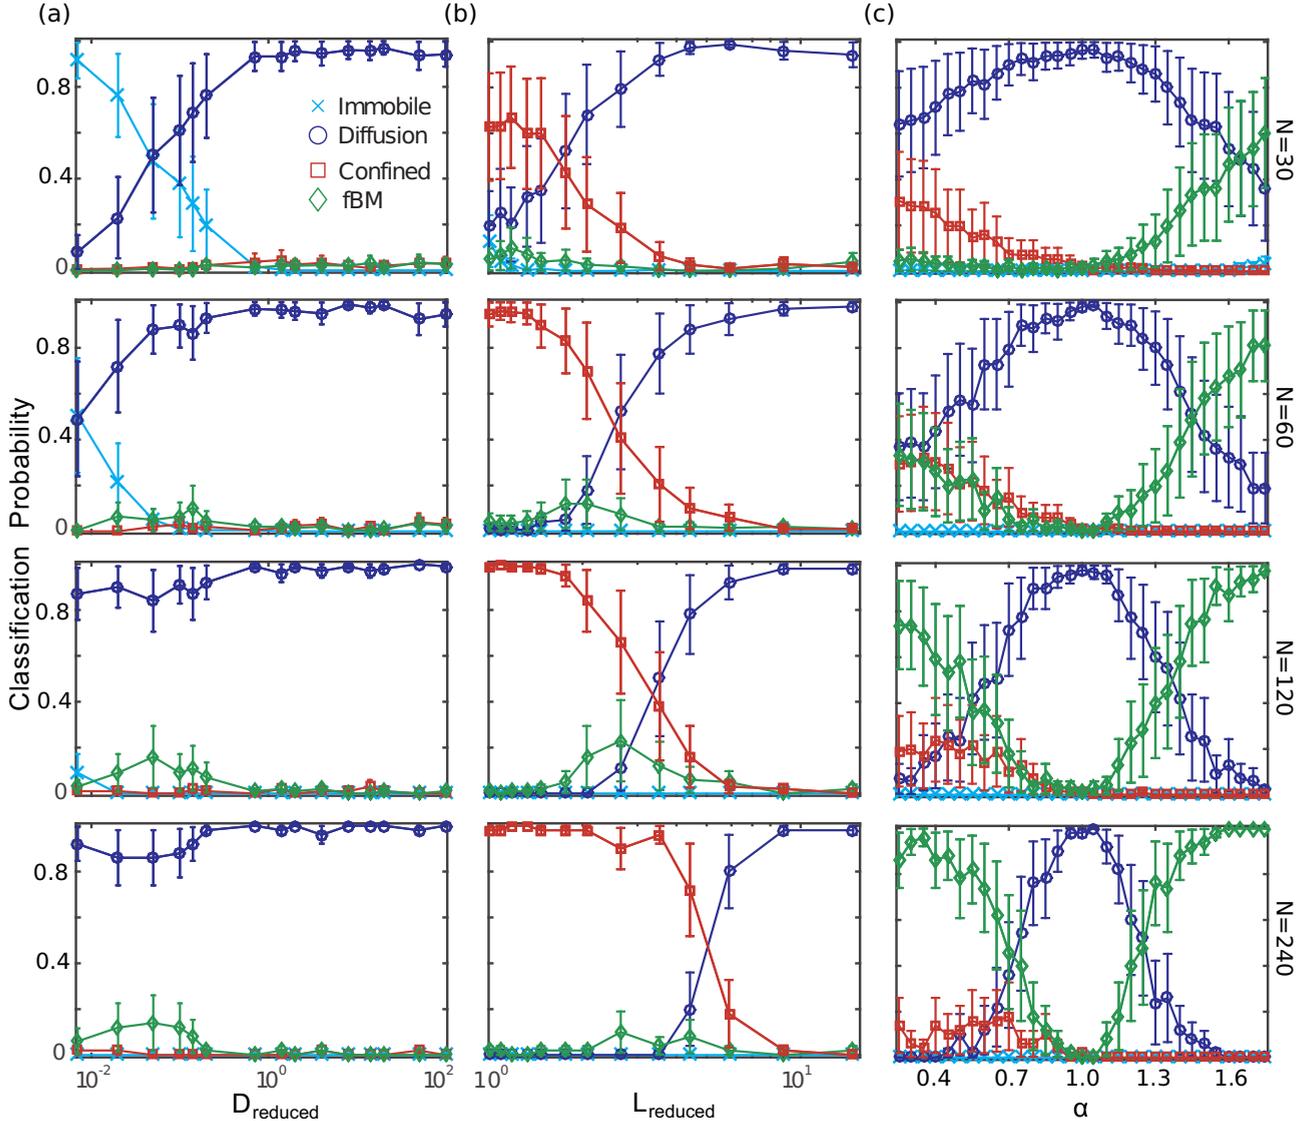


FIG. 1. Classification probability via mleBIC of simulated particle trajectories *with* localization noise for various particle track lengths, undergoing (a) normal diffusion for various underlying diffusivities, (b) confined diffusion for various reduced confinement sizes, and (c) fractional Brownian motion for various anomalous exponents. Each row represents a different particle track length: first row, $N = 30$ steps; second row, $N = 60$ steps; third row, $N = 120$ steps; and last row, $N = 240$ steps. The probability of each model was calculated on the basis of the fraction of tracks classified to that model at each point in parameter space, and it is specified by a unique marker and color: immobile (cyan cross), normal diffusion (blue circles), confined diffusion (red squares), and anomalous diffusion (green diamonds). Error bars represent the observed standard deviation.

leads mleBIC to favor an immobile model. Thus, more statistics are necessary to reject the simpler immobile model.

For particle trajectories undergoing confined diffusion [Fig. 1(b)], when confinement sizes are small, a confined diffusion model is favored. As the confinement size increases, a normal diffusion model becomes favored. At this confined-to-normal crossover, a small preference for anomalous diffusion is found. As expected, longer trajectories provide more opportunities to explore the boundaries of confinement, resulting in a wider region of parameter space for which a confined diffusion model is favored.

For particle trajectories undergoing fBm [Fig. 1(c)], a normal diffusion model is mostly favored when particle trajectories are short ($N \leq 60$). A fBm diffusion model is not

consistently favored until trajectories contain 240 steps, albeit only when the anomalous exponent is below 0.7 or greater than 1.3. As expected, when particle trajectories contain minimal localization noise errors, mleBIC yields improved estimation for fBM (Fig. S7). Thus, the presence of localization noise requires even longer tracks for proper classification, even though MLE can determine reliable estimates for the underlying diffusivity and anomalous exponent (Fig. S4). BIC's built-in parsimony causes it to favor a normal diffusion model when there are not enough data to support a non-normal diffusion model, even when the correct model corresponds to non-normal diffusion. This behavior seems not undesirable.

Similar to taMSD analysis, mleBIC does not take into account transitions between diffusive states. While analyzing

subsets of the data may allow for different diffusive states within a particle trajectory, Figs. 1 and S7 illustrate that accurate classification cannot be made for wide ranges of parameter space, even in the most ideal circumstances. As the trajectories become longer, the statistical power grows, thereby allowing for improved mleBIC classification over a wider parameter space, and misclassification gradually reduces. However, longer particle trajectories that have constant diffusion properties becomes increasingly unlikely, especially when a particle is diffusing in a complex environment such as a living cell. Thus, while mleBIC is certainly an improvement over taMSD analysis, the statistical power of classification by analyzing particle trajectories on an individual basis remains limited.

C. Systems-level analysis of a collection of particle trajectories

To augment the limited statistics provided by individual particle trajectories, pEM simultaneously analyzes a collection of trajectories by employing a systems-level likelihood function to account for a finite number of unique diffusive states, each of which we envision to arise as a result of particular interactions within the cell. Here, we extend the original pEM framework [17] to now include non-normal modes of diffusion, i.e., we lift the short-time-diffusion approximation. A powerful aspect of this version of pEM is that it is essentially a *model-free approach* in that no prior assumptions need to be made concerning which types of diffusion mode are present in the data at hand.

To implement our version of pEM, we first write the systems-level log-likelihood function:

$$\ln \mathcal{L}(\Delta \hat{\mathbf{x}} | \hat{\boldsymbol{\pi}}, \hat{\boldsymbol{\Sigma}}) = \sum_{m=1}^M \ln \left\{ \sum_{k=1}^K \pi_k P(\Delta \mathbf{x}_m | \Sigma_k) \right\}, \quad (8)$$

where M is the total number of tracks, which collectively realize K distinct underlying diffusive states, $\Delta \mathbf{x}_m$ represents the vector of N displacements for particle trajectory m , and $\Delta \mathbf{x}_m = \{\Delta x_m(n)\}_{n=1}^N$. $\Delta \hat{\mathbf{x}} = \{\Delta \mathbf{x}_m\}_{m=1}^M$ is the set of all particle track displacements. $\hat{\boldsymbol{\pi}} = \{\pi_k\}_{k=1}^K$ is the set of variables that represent the fraction of the population of trajectories that realize diffusive state k , which is bounded and normalized: $0 \leq \pi_k \leq 1$ and $\sum_{k=1}^K \pi_k = 1$, and $\hat{\boldsymbol{\Sigma}} = \{\Sigma_k\}_{k=1}^K$ is the set of covariance matrices that define each diffusive state.

Importantly, the theoretical covariance matrix for any diffusion mode that undergoes a stationary Gaussian process, including in the presence of localization noise, has a symmetric Toeplitz form (Table I), so that the element (i, j) of the covariance matrix depends only on $|i - j|$. We can impose the requirement that the covariance matrix for each diffusive state, Σ_k , take on such a symmetric Toeplitz form by averaging the diagonal, one-off-diagonal, two-off-diagonal, etc. elements of the empirical covariance matrix for particle track m to obtain the experimental covariance matrix elements for track m : $\mathbf{C}_m(i, j) = \mathbf{C}_m(|i - j|) = \langle \Delta \mathbf{x}_m(l) \Delta \mathbf{x}_m(l + |i - j|) \rangle$, where the average is taken over all possible values of l for track m .

Furthermore, because the covariance structure of each diffusion mode decreases rapidly to zero for increasing separation between displacements—i.e., with increasing $|i -$

$j|$ —we can reasonably restrict the number of informative covariance matrix elements that we include in the analysis by setting $\mathbf{C}_m(|i - j|) = 0$ for $|i - j| > f$, where f is the number of off-diagonal covariance matrix elements included in the analysis. If $f = 0$, only the diagonal elements of the covariance matrix are permitted to be nonzero, reproducing the theoretical structure of the covariance matrix for simple diffusion in the absence of localization noise. For $f = 1$, a one-off-diagonal element is included, permitting the covariance matrix to properly account for localization noise sources. In principle, different diffusive states, which are characterized by unique diffusion properties, may be distinguished on the basis of different values of the covariance matrix elements. The inclusion of additional off-diagonal terms introduces additional information to help distinguish diffusive states that undergo confined diffusion, fBm, or other modes of non-normal diffusion.

Because pEM discovers the values of these covariance matrix elements for each diffusive state, it is not necessary to specify ahead of time what diffusion modes are present, beyond specifying f . It is in this sense that this version of pEM is model-independent. In the case of K diffusive states, insisting that the covariance matrix must be a symmetric Toeplitz matrix and limiting the number of off-diagonal matrix elements to f means that the number of model parameters is equal to $K(1 + f) + K - 1$. (There are $K - 1$ independent population fractions.)

Maximizing Eq. (8) with respect to $\{\Sigma_k, \pi_k\}_{k=1}^K$ naturally yields the expectation-maximization (EM) algorithm [29]. In the expectation step, the posterior probability, γ_{mk} , that particle trajectory m realizes diffusive state k , given the current estimates for Σ_k and π_k , is calculated according to

$$\gamma_{mk} = \frac{\pi_k P(\Delta \mathbf{x}_m | \Sigma_k)}{\sum_{j=1}^K \pi_j P(\Delta \mathbf{x}_m | \Sigma_j)}. \quad (9)$$

In the maximization step, the posterior probability is used to update the parameter estimates of each diffusive state:

$$\Sigma_k = \frac{1}{M_k} \sum_{m=1}^M \gamma_{mk} \mathbf{C}_m, \quad (10)$$

$$\pi_k = \frac{M_k}{M}, \quad (11)$$

where $\mathbf{C}_m(i, j) = \langle \Delta \mathbf{x}_m(i) \Delta \mathbf{x}_m(j) \rangle$ and $M_k = \sum_{m=1}^M \gamma_{mk}$. The EM algorithm solves these equations iteratively until the change in the ln-likelihood becomes smaller than a set threshold [29].

The extension to higher dimensions than 1 is carried out as follows. We calculate the expectation step by averaging the posterior probability over each dimension using the same parameter estimates. For the maximization step, the maximized parameter estimates are calculated separately for each dimension and then averaged. At each step in the iteration procedure, the complete ln-likelihood is calculated by summing the ln-likelihood from each dimension.

Although the EM algorithm guarantees convergence to a maximum [29], convergence to the global maximum is not guaranteed, depending on the initial parameter values. However, as described in detail in Ref. [17] and summarized

in Sec. IV, suitably perturbing the likelihood surface, namely pEM, is a computationally efficient means to reach the global maximum likelihood.

Since the number of diffusive states is not known *a priori*, we repeat the pEM procedure for different numbers of diffusive states, finding the maximum likelihood in each case. To maintain model parsimony, we again employ the Bayesian information criterion to penalize for the inclusion of additional diffusive states via a systems-level extension of Eq. (6). Specifically, we select the model with the largest value of the systems-level BIC, where now $\ln \mathcal{L}$ is the systems-level likelihood function [Eq. (8)], the number of free parameters is $N_{\text{params}} = K(1 + f) + K - 1$, and M is the total number of particle track displacements across the population of tracks.

The procedure described so far makes the assumption that the diffusive properties remain constant throughout the duration of each trajectory. To extend pEM so that it can be applied to trajectories containing transitions between different diffusive states, we split each trajectory into equal-sized bins, such that each bin contains B sequential steps. The assumption of a constant covariance matrix is still assumed to hold within each such bin, but different bins can realize different diffusive states. In this way, pEMv2 is able to account for transitions between different diffusive states within the overall trajectory. Each bin is treated as a Markovian measurement of the diffusive state, Eq. (2). The temporal resolution corresponds to the bin size.

To summarize, our enhanced version of pEM, which we call pEMv2, examines a population of binned particle trajectories, each containing B steps, to determine the number of unique covariance matrices contained in the population. It accomplishes this goal by classifying each binned trajectory to a particular diffusive state, based on similarities in the covariance structure among trajectories. Using the resultant classification, pEMv2 then updates the parameter estimates for each diffusive state. Iteration of this process allows pEMv2 to learn in an unsupervised manner what unique covariance structures, i.e., what diffusive states, are realized within the population of binned trajectories. Since the number of diffusive states is intrinsically handled by the BIC [Eq. (6)], the user-controllable parameters for pEMv2 are the number of off-diagonal elements to include in the covariance matrix, f , and the bin size, B .

1. Dependence on the number of covariance terms

To investigate the performance of pEMv2, we have generated a number of sets of synthetic particle trajectories containing different numbers of diffusive states and different degrees of similarity between the covariance terms across diffusive states. Table II specifies the four sets of diffusion parameters (cases 1–4), which were used to generate the synthetic data sets. There are no transitions among different diffusive states for cases 1 and 2, i.e., the transition probability matrix (\mathbf{A}) is given by $\mathbf{A} = \delta_{i,j}$. However, for cases 3 and 4, transitions are permitted with the corresponding matrices of transition probabilities given by

$$\mathbf{A}_3 = \begin{pmatrix} 0.995 & 0.001 & 0.004 \\ 0.001 & 0.995 & 0.004 \\ 0.015 & 0.015 & 0.970 \end{pmatrix} \quad (12)$$

TABLE II. Simulation parameters for synthetic particle trajectories generated for cases 1, 2, 3, and 4.

		1	2	3	4
Case 1	Mode	Confined	Normal	fBM	fBM
	$D_k^{\text{sim}} (\mu\text{m}^2 \text{s}^{-1})$	0.05	0.15	0.25	0.4
	$L_k^{\text{sim}} (\mu\text{m})$	0.13			
	α_k^{sim}	1	1	0.9	0.6
	$\sigma_k^{\text{sim}} (\mu\text{m})$	0.04	0.04	0.04	0.04
	π_k^{sim}	0.25	0.25	0.25	0.25
Case 2	Mode	Confined	Normal		
	$D_k^{\text{sim}} (\mu\text{m}^2 \text{s}^{-1})$	0.06	0.06		
	$L_k^{\text{sim}} (\mu\text{m})$	0.1			
	α_k^{sim}	1	1		
	$\sigma_k^{\text{sim}} (\mu\text{m})$	0.04	0.04		
	π_k^{sim}	0.4	0.6		
Case 3	Mode	Confined	Confined	Normal	
	$D_k^{\text{sim}} (\mu\text{m}^2 \text{s}^{-1})$	0.005	0.1	0.3	
	$L_k^{\text{sim}} (\mu\text{m})$	0.05	0.2		
	α_k^{sim}	1	1	1	
	$\sigma_k^{\text{sim}} (\mu\text{m})$	0.04	0.04	0.04	
	π_k^{sim}	0.33	0.33	0.34	
Case 4	Mode	Normal	fBM	Normal	fBM
	$D_k^{\text{sim}} (\mu\text{m}^2 \text{s}^{-1})$	0.001	0.03	0.2	0.45
	$L_k^{\text{sim}} (\mu\text{m})$				
	α_k^{sim}	1	0.7	1	0.9
	$\sigma_k^{\text{sim}} (\mu\text{m})$	0.04	0.04	0.04	0.04
	π_k^{sim}	0.25	0.25	0.25	0.25

for case 3, and

$$\mathbf{A}_4 = \begin{pmatrix} 1 - 3p & p & p & p \\ p & 1 - 3p & p & p \\ p & p & 1 - 3p & p \\ p & p & p & 1 - 3p \end{pmatrix} \quad (13)$$

for case 4, where p is input into the simulation selected from one of $\{0, 0.003, 0.005, 0.01, 0.015, 0.02, 0.03\}$.

The covariance matrix elements of different diffusive states must be sufficiently distinct in order for pEMv2 to resolve them as separate diffusive states. First, therefore, we sought to explore the effect of the number of off-diagonal covariance matrix elements (f) that are included in pEMv2 analysis. Figure 2 shows the measured probability distributions of the average covariance matrix elements, $\langle C(i, j) \rangle$ for $|i - j| = 0, 1, 2, \text{ and } 6$, determined from populations containing 1500 synthetic trajectories, realizing four diffusive states with diffusion parameters corresponding to case 1 (top row), and 1500 synthetic trajectories, realizing two diffusive states with diffusion parameters corresponding to case 2 (bottom row). To recapitulate the variability found experimentally, the trajectory lengths were distributed according to an exponential probability distribution with a characteristic length of 25 steps, with a minimum cutoff of 15 steps and a maximum cutoff of 60 steps. In addition, because there are no transitions in our analyses of cases 1 and 2, we analyzed each complete trajectory as a whole, as in the original version of pEM, without splitting into bins. Case 1 corresponds to four diffusive states—two normal diffusion, one fBM, and one confined diffusion—but

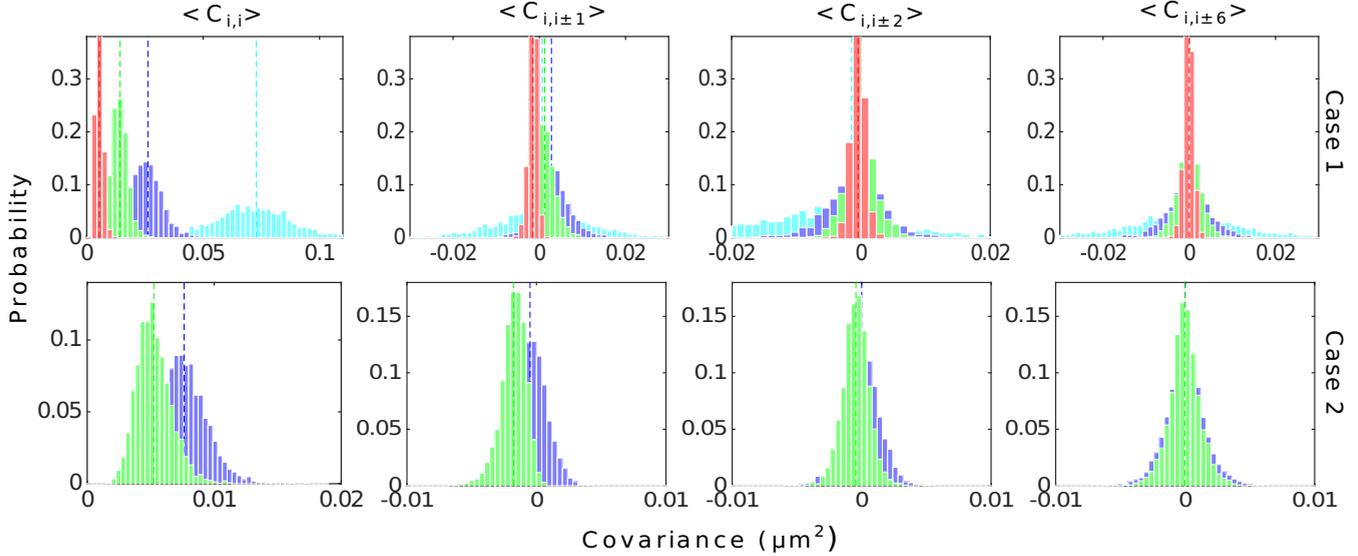


FIG. 2. Empirical probability distributions of the mean covariance matrix elements, $\langle C(i, j) \rangle$, for $j = i, j = i \pm 1, j = i \pm 2,$ and $j = i \pm 6$ for case 1 (top row), where states 1, 2, 3, and 4 are shown in red (second darkest gray), green (second lightest gray), blue (darkest gray), and cyan (lightest gray), respectively, and for case 2 (bottom row), where states 1 and 2 are shown in green (light gray) and blue (dark gray), respectively. Vertical dashed lines indicate the theoretical values with a color/gray value corresponding to each diffusive state.

their diffusion coefficients are well-separated from each other. Case 2 corresponds to two diffusive states with the same diffusion coefficient, one corresponding to normal diffusion and the other to confined diffusion.

For both cases 1 and 2, the means of the distributions of $\langle C(i, i) \rangle$ and $\langle C(i, i \pm 1) \rangle$ for each diffusive state are well separated. For case 1, however, the means of $\langle C(i, i \pm 2) \rangle$ are all very similar to each other and are close to zero, with the exception of state 4 (cyan). By contrast, for case 2, the means of $\langle C(i, i \pm 2) \rangle$ for state 1 and state 2 remain distinguished from each other. For cases 1 and 2, the means of $\langle C(i, i \pm 6) \rangle$ for each diffusion state are all very similar to each other and are all close to zero, albeit their widths remain distinct.

Figure 3 shows the In-probability of each model size, determined on the basis of BIC score [Eq. (7)], as a function of model size for different numbers of nonzero off-diagonal covariance matrix elements between $f = 1$ and 13. For case 2, the correct number of diffusive states is found ($K = 2$), irrespective of f . For case 1, where $K = 4$, pEMv2 is able to successfully determine the correct numbers of diffusive states, as indicated by the maximum In-probability, except when $f = 13$, for which a three-diffusive-state model is favored for three out of the five data sets analyzed.

A visual representation of how successfully pEMv2 determines the correct diffusive state is given in Fig. 4, which shows 1500 synthetic particle trajectories corresponding to case 1 (top row) and case 2 (bottom row). In the left column, each trajectory is depicted using a color corresponding to the known, simulated diffusive state of the track. In the right column, each trajectory is depicted using a color corresponding to the diffusive state that realizes the maximum posterior probability for that track, determined using $f = 6$ off-diagonal covariance matrix elements for cases 1 and 2. Although there are a few misclassified trajectories, the overwhelming majority of the trajectories

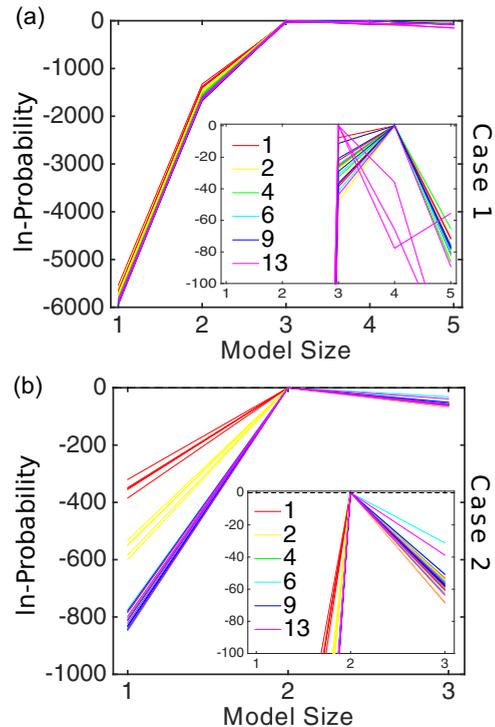


FIG. 3. In-probability (base e) as a function of the number of diffusive states (model size) for (a) case 1 and (b) case 2, determined by pEMv2 analysis using $f = 1, 2, 4, 6, 9,$ and 13 off-diagonal covariance matrix elements. The inset key shows the correspondence between color (gray-scale) value and the value of f . Each data set consists of five sets of simulated particle tracks (shown as a different curve) for each f [shown as a different color (gray-scale) value]. The inset panel shows a zoomed-in representation near the maximum In-probability.

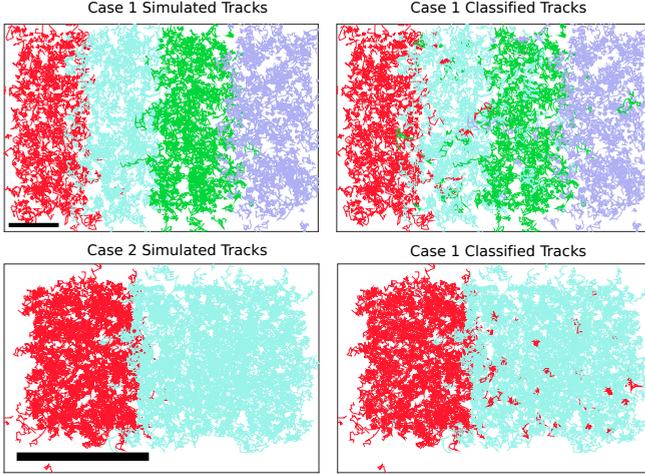


FIG. 4. Representations of 1500 synthetic particle trajectories for case 1 (top row) and case 2 (bottom row). In the left column, each trajectory is depicted using a color (gray-scale) value, corresponding to the known, simulated diffusive state of the track. In the right column, each trajectory is depicted using a color (grey-scale) value, corresponding to the diffusive state, that yields maximum posterior probability, determined on the basis of pEMv2 using $f = 6$ off-diagonal covariance matrix elements. The starting position of each trajectory is on a two-dimensional grid, separated one from another by $1 \mu\text{m}$ (top row) and $0.4 \mu\text{m}$ (bottom row). In both cases, the scale bar represents $5 \mu\text{m}$.

are correctly classified, demonstrating that pEMv2 is capable of reliably uncovering the diffusive states in these cases.

Figure 5 shows the fraction of correctly classified trajectories as a function of the number of off-diagonal covariance matrix elements, confirming that pEMv2 reliably classifies trajectories to the correct diffusive state. The classification accuracy shows only a modest dependence on the number of off-diagonal covariance matrix elements included in the analysis: For case 1, the accuracy of classification is uniformly high for f between 1 and 9, suggesting that the first off-diagonal covariance matrix element ($f = 1$) is decisive in that case. The decrease in classification accuracy for $f = 13$ may be because of the inclusion in this case of a large number of noisy off-diagonal matrix elements, suggesting that it is preferable to not include too many off-diagonal covariance matrix elements. For case 2, the accuracy noticeably improves as f increases from 1 to 4, and it remains high thereafter, suggesting that off-diagonal covariance matrix elements up to $f = 4$ are informative for classification in this case.

The classified covariance matrix elements and the classified taMSD are shown in Fig. 6 for each diffusive state corresponding to case 1. In this instance, using either $f = 1$ or 6 in the analysis leads to the characterization of each diffusive state with high fidelity, with the measured covariance matrix elements and measured taMSDs for each diffusive state, shown as the data points and the solid lines in the figure, almost exactly matching the corresponding true covariance matrix elements and true taMSDs, shown as the dashed lines, which are very nearly coincident with the solid lines. In comparison with mlBIC, which was unable to reliably classify 60-step trajectories undergoing fBm with anomalous exponents of

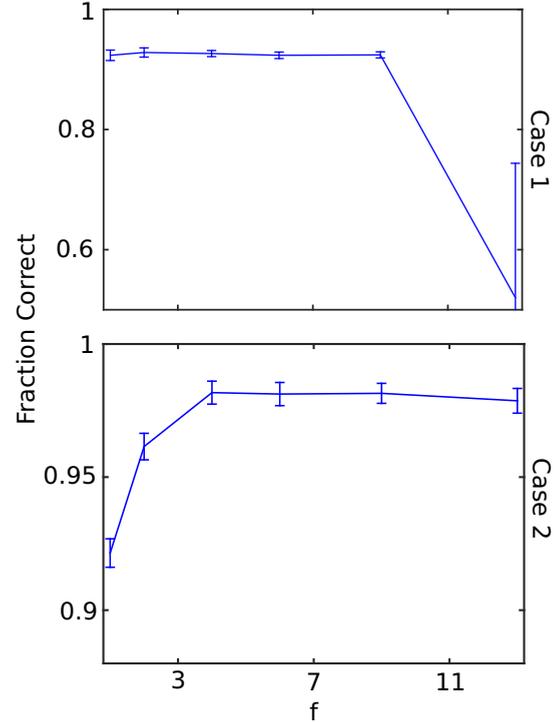


FIG. 5. Fraction of trajectories classified into the correct diffusive state as a function of the number of off-diagonal covariance matrix elements used in the pEMv2 analysis of case 1 (top) and case 2 (bottom). Error bars represent the observed standard deviation across five different sets of simulated particle tracks.

either $\alpha = 0.9$ or even $\alpha = 0.6$, it is striking that pEMv2 is not only able to identify these two diffusive states (states 3 and 4 of case 1) and to accurately categorize individual trajectories into these states (Fig. 4), it is also able to accurately capture the anomalous behavior of their taMSDs (Fig. 6). Thus, the systems-level strategy employed by pEMv2 can find subtle deviations from non-normal diffusive behavior that are statistically challenging to uncover if trajectories are analyzed on an individual basis.

These observations show that the particular value of f used is not critical. In practice, we suggest that a reasonable way to pick f is on the basis of the average covariance matrix elements themselves (see Fig. 6): we suggest picking f to correspond to the off-diagonal term of the ensemble-averaged covariance matrix elements that has essentially converged to zero. This choice should ensure that all informative covariance matrix elements are included in the analysis, but that unnecessary noise is excluded. We ascribe the failure to select the correct model for $f = 13$ to be the result of including unnecessary noise. For the simulations in this paper, $f = 6$ is a reasonable choice.

2. Uncovering transitions by splitting tracks

A population of experimental trajectories that realizes multiple diffusive states is likely to contain at least a subset of trajectories that contain transitions among the diffusive states. The prevalence of transitions depends on their underlying kinetics, i.e., on the transition rates. Our concept for extending

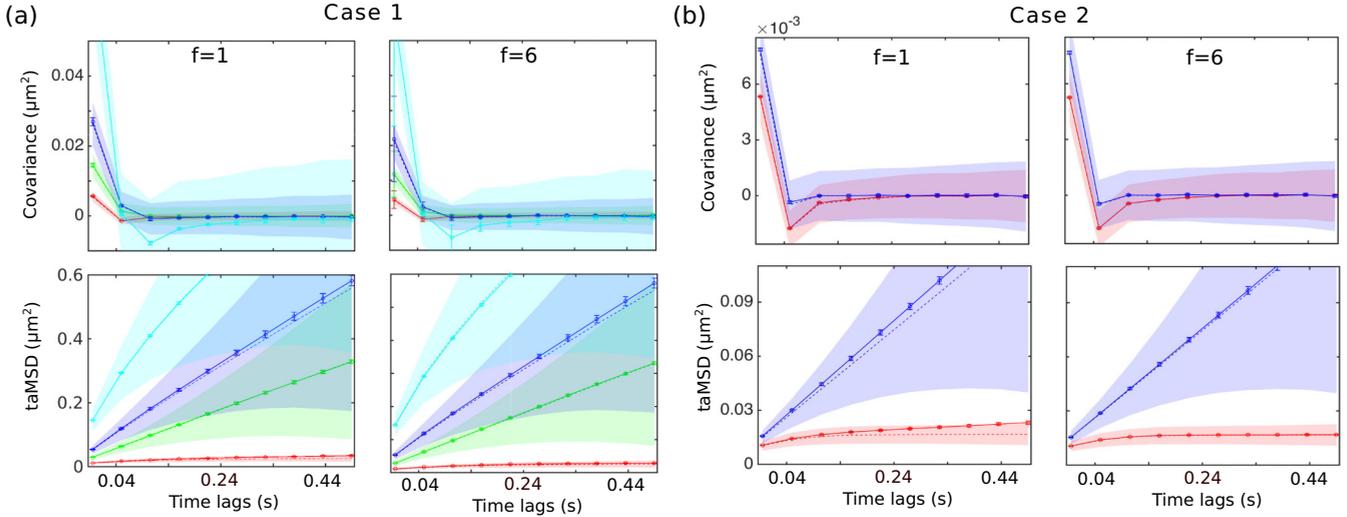


FIG. 6. Average covariance matrix elements and average taMSD from maximum posterior classification as a function of time lag for (a) states 1 (red, bottom curve), 2 (green, second from bottom curve), 3 (blue, second from top curve), and 4 (cyan, top curve) of case 1, and (b) states 1 (red, bottom curve) and 2 (blue, top curve) of case 2. The n th covariance matrix element is calculated according to $\langle [x(i+n+1) - x(i+n)][x(i+1) - x(i)] \rangle$ and the n th taMSD is calculated according to $\langle [x(i+n) - x(i)]^2 \rangle$ for each time lag, $n\Delta t$, where $n = 0, 1, \dots, 9$. Each data point represents the average over five different sets of simulated particle tracks, where the ensemble taMSD and covariance matrix elements were calculated for each state based on classification by pEMv2 using $f = 1$ (left column) and $f = 6$ (right column). The solid lines linking the data points are guides to the eye. The error bars correspond to the standard deviation of the mean across five trials. Shown as the dashed curves are the true matrix elements and the true ensemble-averaged taMSD for each state, determined using the known diffusive states of trajectories, while the shaded bands represent its standard deviation.

our methodology to permit analysis of trajectories with transitions is to split these trajectories into shorter pieces. If the duration of the resultant short trajectories is less than the typical lifetimes of relevant diffusive states, then each short trajectory will with high probability realize a single diffusive state throughout, and the methods described above remain applicable to determine the diffusive states within the population of these short trajectories.

To investigate the feasibility of this concept, we simulated particle tracks with three diffusive states, corresponding to case 3 in Table II, that transition among each other with transition rate matrix A_3 [Eq. (12)]. The protocol used for generating transitions is described in Sec. IV. The corresponding lifetimes of states 1, 2, and 3 are 200 (~ 6.4 s), 200 (~ 6.4 s), and 33 (~ 1 s) steps, respectively. We then divided the simulated trajectories into sets of short trajectories containing 5, 10, 15, 20, 25, 30, 60, 90, or 120 steps, respectively, while keeping the total number of steps and hence the total positional information constant at 12 000 total steps across all trajectories. We then applied the pEMv2 methods described above to each population of different-length short trajectories, implicitly assuming that each short trajectory remains in the same diffusive state throughout. The number of off-diagonal matrix elements used in the analysis was fixed at $f = 6$.

Figure 7(a) shows the BIC-based ln-probability of various model sizes for simulated tracks with lengths 15, 30, 60, 90, and 120 steps. The ln-probability selects the correct number of diffusive states ($K = 3$) only when $N = 15$ steps. For trajectories containing 30 or more steps, the BIC-based probability incorrectly favors a four-diffusive-state model, presumably in an effort to describe trajectories containing transitions. Given that the three-state model is correct, Figure 7(b) shows the

fraction of the total number of steps that are assigned to the correct diffusive state for each set of short trajectories, plotted as a function of the track length of each set. Evidently, the fraction of steps correctly assigned decreases as the trajectories became longer. This trend is surely due to the fact that longer trajectories provide more opportunities to transition, as indicated by the increasing number of transitions per track with increasing trajectory length, shown in Fig. 7(c).

To permit pEMv2 to deal with tracks containing transitions, we implemented a procedure that splits long trajectories into shorter trajectories. For the 120-step data set, Fig. 7(d) shows that the ln-probability of various model sizes for tracks, split into 5, 10, 15, 20, or 30 steps, yields the correct model ($K = 3$) for bin sizes less than 30 steps, in agreement with Fig. 7(a). Given that the three-state model is correct, Fig. 7(e) shows the fraction of the total number of steps that are assigned to the correct diffusive state as a function of the bin size. Evidently, this procedure yields a significant improvement in the fraction of steps correctly assigned compared to analysis of the 120-step data set, shown in Fig. 7(b), presumably as a result of decreasing the number of transitions per track from ~ 1.1 per track for the 120-step data set to ~ 0.1 transitions per track for the bin size of 10 steps (Fig. 7). Moreover, pEMv2 is now able to provide a significant improvement in the quality of the estimates for the covariance elements and taMSD for each diffusive state (Fig. 8), as well as reasonable estimates of the transition matrix (Fig. S8).

Is there an optimal bin size? Indeed, Fig. 7(e) shows that the fraction of steps correctly assigned exhibits a maximum at a bin size of 10 steps and decreases for smaller and larger bin sizes. It turns out that using smaller bin sizes may render the results of pEMv2 more susceptible to misclassification

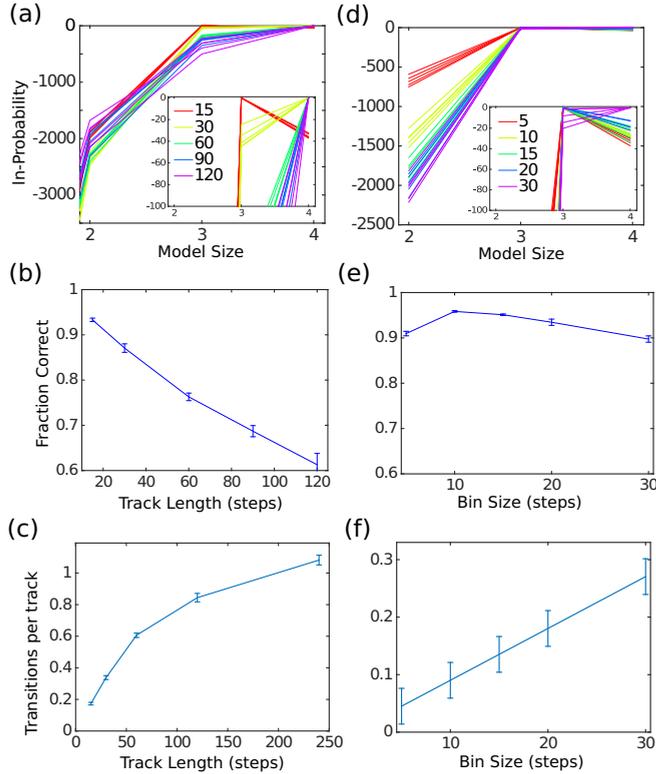


FIG. 7. Performance of pEMv2 on particle tracks that transition between diffusive states given by case 3. In-probability (base e) determined by pEMv2 analysis for five sets of simulated particle tracks (each shown as a different curve) with diffusive states given according to case 3 with (a) track lengths of 15, 30, 60, 90, and 120 steps, and (d) track lengths of 5, 10, 15, 20, and 30 steps created by splitting the 120-step data set in (a). Each simulation set is shown in a different color (grey-scale) value. Average fraction that the classified diffusive states matches the simulated diffusive state for (b) various track lengths and (e) various bin sizes. For the purposes of this comparison, when a bin contains a transition, the “true” diffusive state is chosen to be the state with the highest number of displacements. Average transition rate per track for (c) various track lengths and (f) various bin sizes. (b), (c), (e), and (f) Error bars represent the observed standard deviation across the five data sets.

(Fig. S9). Since information of confinement manifests as anticorrelations between neighboring displacements each time a particle “bounces” off of the confinement barrier—if the bin size is too small, then this information is only contained in the few bins that capture such a “bouncing” event, while other bins would follow an apparent normal diffusion. On the other hand, although including more steps in the bin size allows for more anticorrelation “bouncing” events, a large bin size also has the undesirable effect of increasing the number of transitions per track, which can also lead to poorer performance. Thus, the bin size should be chosen to be as large as possible, subject to the constraint that the mean number of transitions per trajectory should not be too large. In this example of case 3, satisfactory results are obtained by using a level of binning that yields an average of 0.2 transitions per trajectory.

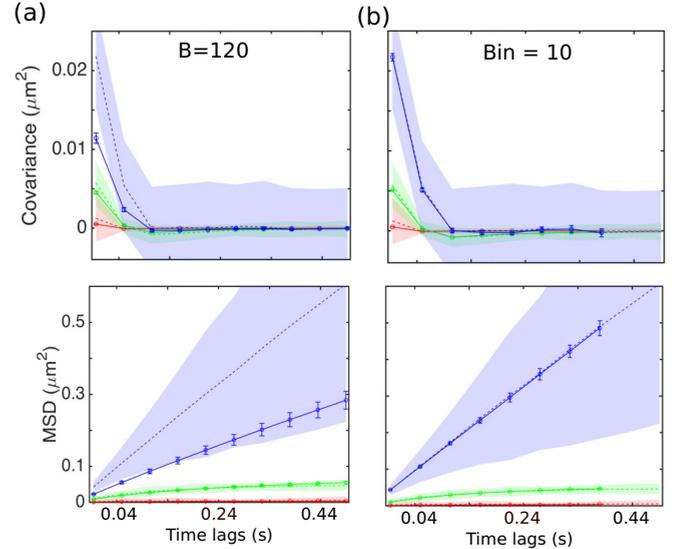


FIG. 8. Average covariance matrix elements and average taMSD from maximum posterior classification as a function of time lag for states 1 (red, bottom curve), 2 (green, middle curve), and 3 (blue, top curve) with a bin size of (a) 120 steps and (b) 10 steps. The n th covariance matrix element is calculated according to $\langle [x(i+n+1) - x(i+n)][x(i+1) - x(i)] \rangle$ and the n th taMSD is calculated according to $\langle [x(i+n) - x(i)]^2 \rangle$ for each time lag, $n\Delta t$, where $n = 0, 1, \dots, 9$. Each data point represents the average over five different sets of simulated particle tracks, analyzed using pEMv2 using $f = 6$. The solid lines linking the data points are guides to the eye. The error bars correspond to the standard deviation of the mean across five trials. Shown as the dashed curves are the true matrix elements and the true ensemble-averaged taMSD for each state, determined using the known diffusive states of trajectories, while the shaded bands represent their standard deviations.

3. Determining the optimal bin size

To further elucidate how pEMv2’s performance depends on the level of transitions and how to determine the optimal bin size in an unsupervised manner, we generated a number of data sets containing 3000 synthetic particle tracks with diffusive states given according to case 4 (Table II), and with varying mean numbers of transitions per track ($R = \{0, 0.36, 0.6, 1.2, 1.8, 2.4, 3.6\}$ transitions per track). All of the track lengths were constant with $N = 120$ steps. For each data set, we applied pEMv2 with bin sizes ranging from 5 to 30 steps. The mean number of transitions per trajectory for each bin size is shown as a function of bin size in Fig. 9(a).

By applying pEMv2 to each of these data sets, the BIC’s In-probability found the correct model size ($K = 4$) when the transition rates were low ($R < 0.6$ transitions per track), irrespective of the bin size used (Fig. S10). Even when the transition rates increase ($R = 2.4$ and 3.6 transitions per track), the BIC continues to favor the correct four-diffusive-state model for smaller bin sizes. However, the BIC favors an incorrect five-diffusive-state model when analyzing data that use bins containing 30 steps.

Assuming the correct model size ($K = 4$), Fig. 9(b) shows the average maximum-likelihood values as a function of the bin size for each data set. When transition rates are low, the

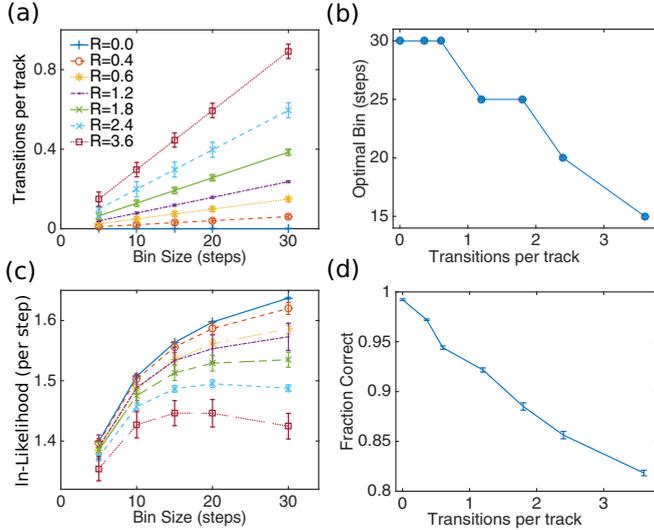


FIG. 9. Dependence of the optimal bin size on transition rate. (a) Average number of transitions per track vs bin size for various transition rates, $R = \{0, 0.36, 0.6, 1.2, 1.8, 2.4, 3.6\}$ transitions per track [each shown in a different color (gray-scale) value]. (b) Average ln-likelihood value per step vs bin size for various transition rates. (c) Optimal bin size vs the mean number of transitions per track, determined by the maximum ln-likelihood per per step. (d) Average fraction that the classified diffusive state matches the simulated diffusive state as a function of the mean number of transitions per track. Each error bar represents the observed standard deviation across five different sets of simulated particle tracks.

average ln-likelihood per step increases monotonically with bin size, suggesting that in these cases the optimal bin size is larger than the maximum binning used. For larger numbers of transitions per track, however, a maximum ln-likelihood per step is observed within the range of bin sizes examined. Figure 9(c) shows that the optimal bin size, determined as the maximum ln-likelihood per step from Fig. 9(b), decreases as the number of transitions per track increases. Not surprisingly, the more transitions that are present, the smaller the bin size should be. Figures S12 and S13 show the pEMv2 classification of representative trajectories of the $R = 3.6$ data set for various bin sizes. When the bin size is five steps, spurious transitions are frequently found, which we ascribe to the relatively larger statistical fluctuations that necessarily accompany smaller bin sizes. As the bin size becomes larger, statistical fluctuations are reduced. However, if the bin size becomes too large ($B = 30$), the corresponding higher rate of transitions per track limits pEMv2's ability to classify diffusive states accurately. Evidently, the optimal bin size balances the accuracy of the covariance matrix elements, which become better-determined with larger bin sizes, against the number of transitions per track, which mix the covariance matrix elements of different diffusive states, leading to poorer pEMv2 performance.

Although pEMv2, using the optimal bin size, is able to uncover the correct numbers of diffusive states and characterize each diffusive state reliably (Fig. S11), the overall accuracy of pEMv2's classification decreases as the number of transitions per track increases, as is indicated by the fraction of correctly classified steps, plotted in Fig. 9(d). Even though the optimal

bin size lowers the *effective* number of transitions per track, the decreased performance may be due to the higher *absolute* number of transitions for the data sets with higher R [Fig. 9(a)]. Notwithstanding, the ensemble behavior of each diffusive state can still be captured accurately when the optimal bin size determined by the maximum likelihood per displacement is used (Fig. S11).

III. CONCLUSIONS

In this paper, we introduced the likelihood functions for two canonical modes of non-normal diffusion, namely confined diffusion and fractional Brownian motion. We showed that the maximum-likelihood estimates provide a significant improvement in comparison with traditional MSD analysis. We introduced a model selection scheme, namely mleBIC, to determine the underlying diffusion model that best represents the motions of a diffusing particle. We demonstrated that while mleBIC is quite successful at classifying tracks without localization noise; classification of tracks with localization noise was limited, especially for short trajectories. Although in this paper we restricted consideration to particles undergoing normal diffusion, confined diffusion, and fBm and immobile particles, extensions to other diffusion models can be added easily by incorporating these models into mleBIC once the likelihood functions are known.

To take advantage of a systems-level approach, we introduced an updated version of pEM analysis, namely pEMv2, that determines the number of unique covariance structures within a population of particle trajectories, thereby bolstering the statistics of individual trajectories. A key output from the pEMv2 algorithm is the posterior probability, γ_{mk} , that particle trajectory m realizes diffusive state k . For the selected model, one simple and useful way to categorize a particular trajectory to a particular diffusive state is to assign the trajectory to the diffusive state that realizes the largest posterior probability, as in Fig. 4.

When analyzing simulated trajectories that transition between different normal/non-normal diffusive states, pEMv2 was able to determine the covariance structure of each diffusive state quite reliably. We also demonstrated the rationales for the selection of the free parameters in pEMv2, which includes the number of covariance features and the bin size. The number of off-diagonal covariance matrix elements to include can be set to the value for which the observed ensemble-averaged covariance matrix elements have just decayed to zero, thereby only including informative covariance terms in the analysis. We have shown that an optimal bin size may be determined by rerunning pEMv2 for various bin sizes, and selecting the bin size that yields the highest likelihood for a given model size. In practice, because the model size is unknown *a priori* for experimental data, we envision running pEMv2 for different model sizes and different bin sizes to find these conditions. Importantly, pEMv2 is rooted in physical principles of stochastic processes. Applying nonphysical clustering methods to the same data, such as k -means clustering, led to poor characterization of the underlying diffusive states [17].

Since pEMv2 does not make any intrinsic assumptions of the underlying diffusion model other than that it follows a Gaussian process, pEMv2 is essentially a diffusion-model-free

approach. Characterization of each covariance structure to determine the diffusion mode and properties can then be performed *post hoc*. Specifically, traditional analyses can then be applied for each diffusive state, such as calculation of the ensemble-average taMSD and the ensemble-average velocity autocorrelation function. Such a procedure provides a more reliable representation of the diffusive behavior compared to individual trajectories, which suffer from limited statistics.

One drawback to pEMv2 is that information of the diffusive state across every bin is treated independently. Thus, when the bin size becomes small, spurious states may occur more frequently. A key benefit of a hidden Markov model (HMM) approach is that spurious transitions can be intrinsically penalized by maximizing a likelihood function that includes a transition matrix. We envisage that, in the future, pEMv2 can be extended to a HMM of multivariate Gaussians to mitigate spurious transitions. In turn, this approach will improve the temporal resolution by making it possible to reduce the bin sizes. However, if transition rates are inhomogeneous across the cell, any HMM approach that assumes a single transition matrix for the entire cell would not be able to properly capture that inhomogeneity. Notwithstanding, current HMMs applied to SPT data, namely vbSPT and HMM-SPT [9], apply a HMM of univariate Gaussians, which is equivalent to using a bin size of two steps ($f = 0$) and thus only using information of the first covariance term. Thus, these HMM analyses overlook localization noise, which introduces correlations between nearest-neighbor displacements, rendering each displacement non-Markovian. Moreover, neither method can properly account for non-normal diffusion models such as confined diffusion and fBM.

Unfortunately, there is no strict rule concerning how many tracks are needed for pEMv2 to return accurate results. Rather, the amount of data needed depends on the complexity of the diffusive states involved, as discussed previously [17]. In practice, we recommend that pEMv2 users complement their SPT analysis of experimental data with an analogous analysis of simulated tracks that recapitulate the diffusive complexity determined by pEMv2. In this way, the user can determine the reliability of pEMv2 for the data at hand, and thereby gain confidence in the results provided by pEMv2.

With the ability to handle normal/non-normal diffusive states that contain transitions between different diffusive states, we envision that pEMv2 can help to uncover more accurate information regarding the diffusive states that occur inside live cells with single-molecule resolution. This analysis sets the benchmark for all future single-particle tracking analyses to begin to understand the spatiotemporal biochemistry of diffusing particles inside live cells with single-molecule resolution.

IV. METHODS

A. Simulation procedure

Synthetic particle trajectories undergoing normal diffusion are generated using the recursion given by Eq. (1), with $\Sigma(i, j) = 2D\Delta t\delta_{i, j}$ and $x(0) = 0$. To generate synthetic particle trajectories undergoing normal diffusion confined in a finite square geometry with size $-L$ to L , we simulate displacements

that follow normal diffusion. At each time step, if the new position falls outside of the finite domain, then the simulated position is set such that the difference between the proposed position and the boundary is reflected, i.e., the Neumann boundary condition, but the total distance traveled remains the same as if the wall were not present. Here, the starting position of each trajectory is at the center of the confinement boundary.

Synthetic particle trajectories undergoing fractional Brownian motion are generated using the recursion given by Eq. (1) with $\nu = 0$ and Σ given by Appendix. Here, the square root of the covariance matrix is determined with the Cholesky decomposition, i.e., $\Sigma = LL^T$, where L is the Cholesky lower triangular matrix. We then generate a vector of normally distributed random numbers $\mathbf{W} = \{W(d)\}_{d=1}^D$, where D is the number of displacements of the particle trajectory, and we apply a matrix multiplication according to $\Delta\mathbf{x} = L\mathbf{W}$. The positions are then reconstructed by calculating the cumulative sum of the displacements $x(i) = x(0) + \sum_{j=1}^i \Delta\mathbf{x}(j)$, with $x(1) = 0$. For each particle trajectory, the process is carried out separately for two spatial dimensions, which are then combined to form the true two-dimensional (2D) positions of the synthetic particle trajectory.

To incorporate transitions between diffusive states within each trajectory, we first generated a random Markov chain with a known transition matrix, A , to specify the state sequence of each particle track displacements. For each state, the displacements are simulated according to the properties of the diffusive state. The positions of the particle trajectories are then reconstructed by calculating the cumulative sum of the displacements $x(i) = x(0) + \sum_{j=1}^i \Delta\mathbf{x}(j)$, with $x(1) = 0$. Each time the Markov state goes to a confined diffusion state, the confinement boundaries are reset with the initial position at the center. When the Markov state switches to another diffusive state, information of the confinement boundaries is forgotten.

Dynamic localization noise is incorporated into the positions by simulating 32 microstep displacement ($\delta t = \Delta t/32$) time steps and averaging 32 successive positions. The net effect is an exposure time equal to the frame duration of 32 ms. Static localization noise is included by adding a normally distributed random number with zero mean and variance, σ_{sim}^2 , to each motion-blurred position. To generate a collection of particle trajectories, the population fractions are used to determine the number of particle trajectories that are initialized to each diffusive state. Here, population fractions serve as the percentage with which the initial state of each trajectory begins.

B. MSD analysis

For a stationary sequence of T 2D particle positions, $\mathbf{x} = \{x(t), y(t)\}$ for $t = 1$ through T , each separated one from the next by a time, Δt , the taMSD is given according to [5,30]

$$\overline{\delta(\Delta_n, T)} = \frac{1}{T - \Delta_n} \sum_{t=1}^{T-\Delta_n} [x(t + \Delta_n) - x(t)]^2 + [y(t + \Delta_n) - y(t)]^2,$$

TABLE III. MSD models for canonical modes of diffusion.

Mode	MSD model
Normal	$4D\Delta t + 4\sigma_0^2 - \frac{2}{3}D\Delta t$
Confined	$\frac{L^2}{3} - \frac{32L^2}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} \exp\left[-\left(\frac{k\pi}{L}\right)^2 Dn\Delta t\right] + 4\sigma^2$
fBm	$4D\Delta t^\alpha + 4\sigma^2$

where $\overline{\delta(\Delta_n, T)}$ is the taMSD for the n th time lag, $\Delta_n = n\Delta t$, and the bar on top of $\delta(\Delta_n, T)$ is used to distinguish the time average.

We generate the taMSD for the first 14 time lags, and we employ an unweighted nonlinear least-squares fit with diffusion models given in Table III, where σ_0 represents the static localization noise and σ represents the combined static and dynamic localization noise terms.

C. MLE analysis

For a given diffusion model, the maximum likelihood, or equivalently the minimum negative ln-likelihood, is found by employing a constraint, gradient-based, numerical optimization algorithm in MATLAB (Mathworks), namely `fmincon`. At each optimization step, however, the ln-likelihood function requires the calculation of the ln-determinant and the inverse of the covariance matrix. When particle tracks are long or the elements of the covariance matrix are very small, the ln-determinant of the covariance matrix can run into numerical underflow issues. To make this optimization procedure more robust, we employ an eigenvalue decomposition of the covariance matrix, $\Sigma = P\Lambda P^T$, where P is a matrix of the eigenvectors with their corresponding eigenvalues given along the diagonal of Λ . The ln-determinant is given by the product of the eigenvalues or equivalently the sum of the ln-eigenvalues, i.e., $\ln \det(\Sigma) = \ln \prod_i^D \lambda(i) = \sum_i^D \ln \lambda(i)$. The inverse is given by $\Sigma^{-1} = P\Lambda^{-1}P^T$.

In summary, for a given single-particle trajectory, maximum-likelihood estimation yields the parameter estimates, ln-likelihood value, and Hessian for each candidate diffusion model. From this information, the BIC [Eq. (6)] can be calculated for each diffusion model. Once the BIC for each diffusion model has been calculated, the model probability for each diffusion model can be calculated according to Eq. (7).

TABLE IV. Analytical covariance matrix of particle track displacements separated in time by Δt for canonical diffusion modes, namely normal diffusion, confined diffusion, and fractional Brownian motion, with D the diffusion coefficient, L the confinement size, and α the anomalous exponent. The hat, $\hat{\Sigma}$, represents the covariance matrix without localization error corrections.

Mode	Covariance matrix ($\mu\text{m}^2 \text{s}^{-1}$)
Normal	$\hat{\Sigma}^{\text{normal}}(i, j) = 2D\Delta t\delta_{i, j}$
Confined	$\hat{\Sigma}(i, j)^{\text{confined}} = \begin{cases} \frac{L^2}{6} - \frac{16L^2}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} \Phi(1), & j = i \\ \frac{-L^2}{12} + \frac{8L^2}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} \Phi(1)[2 - \Phi(1)], & j = i \pm 1 \\ \frac{8}{\pi^4} \sum_{k=1, \text{odd}}^{\infty} \frac{1}{k^4} [-2\Phi(j - i + 1) + \Phi(j - i) + \Phi(j - i + 2)] & \text{otherwise,} \end{cases}$
	where $\Phi_n = \exp\left[-\left(\frac{k\pi}{L}\right)^2 Dn\Delta t\right]$
fBM	$\hat{\Sigma}^{\text{fBM}}(i, j) = D\Delta t^\alpha (j - i + 1 ^\alpha + j - i - 1 ^\alpha - 2 j - i ^\alpha)$

Classification is determined by the diffusion model, which yields the highest model probability.

D. pEMv2 analysis

pEMv2 analysis was performed with the MATLAB script provided in the Supplemental Material [19]. Briefly, our pEMv2 procedure employs the EM algorithm on the original set of particle trajectories with random initial parameter values. pEMv2 then reemploys the EM on a Monte Carlo bootstrap set of the original particle trajectories, which serves to perturb the likelihood surface with the aim that a local maximum may no longer be a maximum in the perturbed likelihood surface. Upon completion of a perturbation trial, we verify whether a higher likelihood has truly been found by calculating the likelihood using the pEMv2-converged parameters with the original dataset. If the pEMv2-converged parameters indeed yield a higher likelihood, then the EM parameters are updated by reemploying the EM algorithm initialized with the new pEMv2 parameter estimates on the original dataset. Otherwise, the pEMv2 estimates remain unchanged. This process is repeated until a predetermined number of perturbations have been executed and yield no advance.

To generate each set of random initial values, K random numbers between 0 and 1 are drawn from a uniform distribution. The initial population fractions, $\{\pi_k^0\}_{k=1}^K$, are given by normalizing these random numbers so that the sum is equal to 1. The first covariance values are set using the initial randomly chosen population fractions and the empirical cumulative covariance distribution function. By dividing the cumulative distribution function into K regions proportional to the initial population fractions, the initial covariance value of diffusive state k is then picked as the diffusivity corresponding to the midpoint of region k of the cumulative probability distribution, namely to $\sum_{j=1}^{k-1} \pi_j^0 + \frac{\pi_k^0}{2}$. Particle tracks are then classified to each diffusive state by their distance to the first covariance values. The remaining covariance values for each diffusive state are selected by averaging the classified covariance terms. Thus, we achieve an initialization that serves as a nonparametric method to randomly sample from the observed distribution of diffusion coefficients. We found that this method produces better random initializations than a k -means clustering over the whole covariance matrix. In practice, we found k -means clustering converges to similar values over a wide range of parameter space. In addition,

k -means tends to weed out diffusive states with low population fractions, when two diffusive states are close in proximity.

To improve pEMv2's performance, we applied five random initialization trials of the EM and used the parameters of the trial that yielded the highest likelihood value. We then applied 100 perturbation trials. This was done for each diffusive state starting from $K = 1$ and incrementing K until pEMv2 finds a lower BIC value. Upon completion of pEMv2, the returned parameters include the population fractions and covariance matrices of each diffusive state, along with the posterior probabilities of each particle trajectory.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation Grant No. PHY1305509, and by the Raymond and Beverly Sackler Institute for Physical and Engineering Biology.

APPENDIX: COVARIANCE MATRIX FOR PARTICLE TRACK DISPLACEMENTS WITHOUT LOCALIZATION NOISE

Table IV presents the analytical covariance matrices for particle track displacements without localization noise.

-
- [1] F. Persson, M. Lindén, C. Unoson, and J. Elf, *Nat. Methods* **10**, 265 (2013).
- [2] S. C. Weber, M. A. Thompson, W. Moerner, A. J. Spakowitz, and J. A. Theriot, *Biophys. J.* **102**, 2443 (2012).
- [3] M. T. Valentine, P. D. Kaplan, D. Thota, J. C. Crocker, T. Gisler, R. K. Prud'homme, M. Beck, and D. A. Weitz, *Phys. Rev. E* **64**, 061506 (2001).
- [4] B. R. Parry, I. V. Surovtsev, M. T. Cabeen, C. S. O'Hern, E. R. Dufresne, and C. Jacobs-Wagner, *Cell* **156**, 183 (2014).
- [5] A. Kusumi, Y. Sako, and M. Yamamoto, *Biophys. J.* **65**, 2021 (1993).
- [6] O. Rossier, V. Octeau, J.-B. Sibarita, C. Leduc, B. Tessier, D. Nair, V. Gatterdam, O. Destaing, C. Albigès-Rizo, R. Tampé *et al.*, *Nat. Cell Biol.* **14**, 1057 (2012).
- [7] M. J. Saxton, *Biophys. J.* **72**, 1744 (1997).
- [8] T. Savin and P. S. Doyle, *Phys. Rev. E* **76**, 021501 (2007).
- [9] 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 *et al.*, *Nat. Methods* **12**, 838 (2015).
- [10] M. Ott, Y. Shai, and G. Haran, *J. Phys. Chem. B* **117**, 13308 (2013).
- [11] R. Das, C. W. Cairo, and D. Coombs, *PLoS Computat. Biol.* **5**, e1000556 (2009).
- [12] H. Yang, *J. Chem. Phys.* **129**, 074701 (2008).
- [13] M. H. G. Duits, Y. Li, S. A. Vanapalli, and F. Mugele, *Phys. Rev. E* **79**, 051910 (2009).
- [14] P. J. Slator, C. W. Cairo, and N. J. Burroughs, *PloS One* **10**, e0140759 (2015).
- [15] J. Apgar, Y. Tseng, E. Fedorov, M. B. Herwig, S. C. Almo, and D. Wirtz, *Biophys. J.* **79**, 1095 (2000).
- [16] Y. Tseng, T. P. Kole, and D. Wirtz, *Biophys. J.* **83**, 3162 (2002).
- [17] P. K. Koo, M. Weitzman, C. R. Sabanaygam, K. L. van Golen, and S. G. J. Mochrie, *PLoS Comput. Biol.* **11**, e1004297 (2015).
- [18] M. J. Saxton and K. Jacobson, *Annu. Rev. Biophys. Biomol. Struct.* **26**, 373 (1997).
- [19] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.94.052412> for supplemental figures and text. The Supplemental Material includes Refs. [30–40].
- [20] A. J. Berglund, *Phys. Rev. E* **82**, 011917 (2010).
- [21] X. Michalet, *Phys. Rev. E* **82**, 041914 (2010).
- [22] C. L. Vestergaard, P. C. Blainey, and H. Flyvbjerg, *Phys. Rev. E* **89**, 022726 (2014).
- [23] pEMv2 is written in MATLAB (Mathworks) and is freely available at <https://GitHub.com/MochrieLab/pEMv2>.
- [24] M. P. Backlund, R. Joyner, and W. E. Moerner, *Phys. Rev. E* **91**, 062716 (2015).
- [25] D. T. Gillespie, *Am. J. Phys.* **64**, 225 (1996).
- [26] T. Savin and P. S. Doyle, *Biophys. J.* **88**, 623 (2005).
- [27] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning* (Springer, New York, 2006), Vol. 1.
- [28] K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (MIT Press, Cambridge, MA, 2012).
- [29] A. P. Dempster, N. M. Laird, and D. B. Rubin, *J. R. Statist. Soc. Ser. B*, 1 (1977).
- [30] H. Qian, M. P. Sheetz, and E. L. Elson, *Biophys. J.* **60**, 910 (1991).
- [31] A. Einstein, *Investigations on the Theory of the Brownian Movement* (Courier Dover, Mineola, NY, 1956).
- [32] J.-H. Jeon and R. Metzler, *J. Phys. A* **43**, 252001 (2010).
- [33] J. D. Hamilton, *Time Series Analysis* (Princeton University Press, Princeton, NJ, 1994), Vol. 2.
- [34] A. S. Goldberger, *A Course in Econometrics* (Harvard University Press, Cambridge, MA, 1991).
- [35] N. Monnier, S.-M. Guo, M. Mori, J. He, P. Lénárt, and M. Bathe, *Biophys. J.* **103**, 616 (2012).
- [36] G. Voisinne, A. Alexandrou, and J.-B. Masson, *Biophys. J.* **98**, 596 (2010).
- [37] H. Qian, in *Processes with Long-Range Correlations* (Springer-Verlag, Berlin, Heidelberg, 2003), pp. 22–33.
- [38] B. B. Mandelbrot and J. W. Van Ness, *SIAM Rev.* **10**, 422 (1968).
- [39] K. I. Mortensen, L. S. Churchman, J. A. Spudich, and H. Flyvbjerg, *Nat. Methods* **7**, 377 (2010).
- [40] D. Montiel, H. Cang, and H. Yang, *J. Phys. Chem. B* **110**, 19763 (2006).