Biophysical Journal
Biophysical Perspective

Biophysical Society

A Primer on the Bayesian Approach to High-Density Single-Molecule Trajectories Analysis

Mohamed El Beheiry,[1,2,6] Silvan Türkcan,[3] Maximilian U. Richly,[4] Antoine Triller,[5] Antigone Alexandrou,[4] Maxime Dahan,[1,2,7] and Jean-Baptiste Masson[6,7,*]

[1]Laboratoire Physico-Chimie, Institut Curie, PSL Research University, Paris, France; [2]Department of Radiation Oncology, Sorbonne Universités, Paris, France; [3]Division of Medical Physics, Stanford University School of Medicine, Palo Alto, California; [4]Laboratoire d'Optique et Biosciences, Ecole Polytechnique, Université Paris-Saclay, Palaiseau, France; [5]Biologie Cellulaire de la Synapse, École Normale Supérieure, PSL Research University, Paris, France; [6]Physics of Biological Systems, Institut Pasteur, Paris, France; and [7]Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia

ABSTRACT Tracking single molecules in living cells provides invaluable information on their environment and on the interactions that underlie their motion. New experimental techniques now permit the recording of large amounts of individual trajectories, enabling the implementation of advanced statistical tools for data analysis. In this primer, we present a Bayesian approach toward treating these data, and we discuss how it can be fruitfully employed to infer physical and biochemical parameters from single-molecule trajectories.

In the past years, single-molecule measurements have become a tool of choice for the study of biological systems. In particular, single-molecule tracking techniques play an important role for probing the processes regulating the mobility of biomolecules in living cells. They have been successfully used to investigate the motion of, among others, membrane proteins, molecular motors, mRNAs, and transcription factors (1,2). In practice, single-molecule tracking experiments involve capturing a sequence of images of biomolecules tagged with a fluorescent probe (usually fluorescent proteins, organic dyes, or nanoparticles). Fluorescent probes are then imaged on a sensitive camera and spatially localized (e.g., with a fitting algorithm). The positions are then connected between frames to construct individual molecule trajectories (3). From these trajectories, it is then possible to obtain information on the physical and biological parameters controlling the movement.

The throughput of single-molecule tracking experiments have long been limited to a few tens or hundreds of trajectories. However, the advent of high-density tracking methods, such as sptPALM (4) or uPAINT (5), has changed the scale at which individual motions can be recorded and made it possible to capture hundreds of thousands or even millions of individual trajectories. Importantly, it enables more advanced statistical methods to infer the motion parameters. Moreover, it means that sufficient data are available to determine these parameters in a spatially resolved manner.

In this primer, we discuss the use of Bayesian inference methods to analyze single-molecule trajectories. First, we recall the basic principles of Bayesian inference. Next, we detail its implementation for the analysis and mapping of the stochastic motion of proteins in living cells. Finally, we illustrate the Bayesian approach with experimental results on the dynamics of transmembrane proteins.

Bayesian framework

A general goal in data analysis is the determination of the parameters of a model (e.g., the diffusion coefficient in the case of Brownian particles) given a set of experimental observations (e.g., individual trajectories from single-molecule tracking experiments). Bayesian approaches provide a consistent and reliable framework for extracting information from experimental measurements (6–11). As illustrated in this primer, a great benefit of Bayesian methods is that they easily incorporate hypotheses on the physical and biological properties of the system, as well as on experimental conditions. Classically, Bayesian inference features two steps: the derivation of the posterior probability distribution of the model parameters and sampling from the posterior distribution to estimate the parameters.

Submitted August 19, 2015, and accepted for publication January 15, 2016.

*Correspondence: jbmasson@pasteur.fr

M. El Beheiry and S. Türkcan contributed equally to this work.

Editor: Stanislav Shvartsman

http://dx.doi.org/10.1016/j.bpj.2016.01.018

The starting point is Bayes' law that reads as follows:

$$P(\{U\}|\{T\},M) = \frac{P(\{T\}|\{U\},M)P(\{U\}|M)}{P(\{T\}|M)}, \quad (1)$$

where $\{T\}$ is the set of experimental observations, $\{U\}$ is the set of model parameters (to be evaluated), and M is the model chosen to describe the data. In standard terminology, $P(\{U\}|\{T\},M)$ is the posterior distribution, $P(\{T\}|\{U\},M)$ is the likelihood, $P(\{U\}|M)$ is the prior distribution, and $P(\{T\}|M)$ is the evidence of the model.

The likelihood embodies the physical model and hypotheses regarding the acquisition of data. In the context of tracking experiments, it encodes the model used to describe the motion of the molecules (including the presence of drift, or the Markovian/non-Markovian nature of the process), the positioning noise induced by the experimental conditions, and various hypotheses regarding characteristic scales of the environment.

Prior probability distributions are critical to Bayesian analysis (6,7,9,10,12). They represent knowledge on the parameters before any measurements, including various physical constraints that may not be present in the likelihood. Furthermore, they can impose that the posterior distribution is invariant under reparametrization (i.e., Jeffreys prior), and they can ensure that the posterior distribution is a well-behaved function. Finally, priors can be used to regularize the inferred parameters as discussed below.

The evidence allows access to the probability of a model. It is mostly used in the context of Bayesian model comparison (6,7,9,10,12) as follows:

$$\gamma_{ij} = \frac{P(\{T\},M_i)}{P(\{T\},M_j)} = \frac{P(\{T\}|M_i)P(M_i)}{P(\{T\}|M_j)P(M_j)}, \quad (2)$$

where $P(M_i)$ is the prior probability of the model M_i and where the model with the highest evidence is chosen. In the context of single-molecule analysis, proper model selection is made semi-empirically with the use of numerical simulations and various statistical estimators (13,14).

The final task of the Bayesian approach is sampling the posterior distribution. The most common estimator of the inferred parameters is the maximum a posteriori (MAP), i.e., the highest-probability parameter value from the posterior. It is usually accessed through direct optimization of the posterior distribution. In low dimensions, posterior sampling can be done by direct integration (9,10). Otherwise, it is commonly performed using Monte Carlo sampling (7,9,10,12,15). Fine tuning of the Monte Carlo parameters is often required to obtain efficient sampling in large dimensions. Numerous discussions on Monte Carlo sampling can be found in (7,9,10,12). It is worth noting that in the following and under most relevant assumptions (discussed in the next section), the posterior distributions, in our case, are well-behaved functions with well-defined maxima (usually not unique but with a main one). Hence, a common approximate method for posterior sampling (allowing faster computation than Monte Carlo sampling) is the Laplace approximation (12) of the posterior distribution at the MAP value of the parameters. It provides a Gaussian approximation of the posterior distribution (11,16,17). In most relevant hypotheses, it leads to overestimation of the error in parameter estimation, which stems from asymmetry of the diffusion posterior distribution that is imperfectly approximated by a Gaussian.

Mapping the parameters of single-molecule motion

As indicated above, a unique benefit of Bayesian approaches is that the likelihood and prior distributions can be used to account for noise and variability in experiments. Hence, before explicitly constructing the different terms required for Bayesian inference, it is important to note a few specific features of single-molecule tracking measurements in living cells. Different sources of noise and variability may be identified in single-molecule experiments. First, the dynamics of individual molecules in cells is primarily driven by thermal noise, leading to Brownian motion. Next, the cellular environment is highly heterogeneous. For example, this is the case of the plasma membrane where lipid microdomains, scaffold proteins, and molecular crowding (among other factors) contribute to variability in diffusive properties. Finally, single-molecule experiments are inherently single-cell experiments, meaning that they are subject to phenotypic variabilities between cells.

Given the vast spatial heterogeneity in the dynamics of molecules, it is natural to map the parameters controlling their motion rather than provide a single value averaged over the entire cell. There are a few implicit assumptions in the mapping procedure. First, it is necessary to mesh the cellular environment, meaning that there are scales at which the physical parameters can be considered approximately constant. In turn, the number of trajectories must be large enough so that sufficient information is gathered to infer the parameters in each zone of the mesh. A second assumption is that the cellular system does not significantly evolve over the durations of the measurements (note that slow evolutions may still be analyzed as in (18)). Experimental constraints associated with these assumptions have largely been alleviated with the advent of high-density single-particle tracking techniques that permit rapid recording of a large number of data points (up to several thousands per μm^2 in a few minutes).

Likelihood

We start our Bayesian analysis by constructing the likelihood. The dynamics of molecules, the characteristic scales of motion, the relaxation time of the velocity ($10^{-18} - 10^{-10}$ s),

and the heterogeneity of the environment suggest that the overdamped Langevin equation (19–21) can be used to model various kinds of motion. It reads as follows:

$$\frac{d\mathbf{r}}{dt} = \frac{\mathbf{F}(\mathbf{r})}{\gamma(\mathbf{r})} + \sqrt{2D(\mathbf{r})}\xi(\mathbf{t}), \tag{3}$$

where $D(\mathbf{r})$ is the diffusion, $\gamma(\mathbf{r})$ is the friction, $\mathbf{F}(\mathbf{r})$ is the force, and $\xi(t)$ is a zero-averaged Gaussian noise. Equation 3 describes the time evolution of the molecule in a heterogeneous environment where the diffusion, friction, and forces vary with space. If the force field is conservative, i.e., the energy is conserved, it can be written as the derivative of a potential $\mathbf{F}(\mathbf{r}) = -\nabla V(\mathbf{r})$. The analysis is not limited to conservative force fields. Before analysis there is no information on the various scales at which these fields vary or even if they are present. Furthermore, the heterogeneous spatial diffusion potentially leads to the well-known (Itô-Stratonovich-Hänggi) dilemma (19–26). The dilemma stems from the dependence of the results of stochastic integrals on the hypothesis made to evaluate them. In the case of the Langevin equation, it leads to the "spurious force" induced by diffusion gradients (as shown in Eq. 4).

The Fokker-Planck equation governing the time evolution of the transition probability of a molecule displacement associated to Eq. 3 reads (20) as follows:

$$\begin{aligned}\frac{\partial P(\mathbf{r},t\,|\,\mathbf{r}_0,t_0)}{\partial t} = &-\nabla\left[\left(\frac{\mathbf{F}(\mathbf{r})}{\gamma(\mathbf{r})} + \lambda\nabla D(\mathbf{r})\right)P(\mathbf{r},t\,|\,\mathbf{r}_0,t_0)\right] \\ &+ \nabla[D(\mathbf{r})\nabla P(\mathbf{r},t\,|\,\mathbf{r}_0,t_0)],\end{aligned} \tag{4}$$

where $\mathbf{r}_0$ is the position of the molecule at time t_0; $\mathbf{r}$ is the position of the molecule at time t; and λ is a constant equal to $\lambda = 0$ if Eq. 3 is interpreted by the Itô definition, $\lambda = 1/2$ if it is interpreted by the Stratonovich definition, and $\lambda = 1$ if it is interpreted by the Hänggi (isothermal) definition.

We can now compute the full likelihood. The motion being Markovian (or at least approximated as such), the probability of a trajectory of length N reads $P(\{\mathbf{r}_i,t_i\}_{i=1..N}) = \prod_{i=1}^{N-1} P(\mathbf{r}_{i+1},t_{i+1}|\mathbf{r}_i,t_i)$. It follows from Eq. 4 that the probability P can be expressed as a path integral (27,28) over all paths $\mathbf{r}(s)$ connecting $\mathbf{r}_0$ to $\mathbf{r}$ as follows:

$$P(\mathbf{r},t|\mathbf{r}_0,t_0) \propto \int \mathcal{D}\mathbf{r}(s)\, e^{-\int ds\, Q(\mathbf{r}(s))}, \tag{5}$$

where

$$\begin{aligned}Q(\mathbf{r}(s)) \equiv &\frac{1}{4D(\mathbf{r}(s))}\left(\frac{d\mathbf{r}(s)}{ds} - \frac{\mathbf{F}(\mathbf{r}(s))}{\gamma(\mathbf{r}(s))} + \lambda\nabla\mathbf{D}(\mathbf{r}(\mathbf{s}))\right)^2 \\ &+ \lambda\nabla\frac{\mathbf{F}(\mathbf{r}(\mathbf{s}))}{\gamma(\mathbf{r}(\mathbf{s}))}\end{aligned}$$

is the classic Gaussian weight governing the probability of displacements over an infinitesimal time interval.

We now use the mapping hypothesis suggesting that trajectories close in space share the same properties. Hence, our analysis is performed in a spatially partitioned space: a mesh. We have recently described three methods of meshing in the context of single-molecule trajectories: square (11,14), quad-tree (hierarchical), and Voronoi tessellation (via unsupervised clustering) (29). The goal of the mesh is to balance multiple constraints: local densities of translocations, characteristic sizes of motion heterogeneities, and local domain size scaling. A common property is that the size of the mesh is (at best) on the order of $\sqrt{D\Delta t}$ where D is the local diffusion coefficient and Δt is the time between frames.

Thus, the solution of Eq. 5 can be approximated in mesh subdomains by the Laplace approximation with locally $D(\mathbf{r}) \approx D(\mathbf{r}_c)$, $\gamma(\mathbf{r}) \approx \gamma(\mathbf{r}_c)$, $\mathbf{F}(\mathbf{r}) \approx \mathbf{F}(\mathbf{r}_c)$ and if the force field is conservative $\nabla V(\mathbf{r}) \approx \nabla V(\mathbf{r}_c)$ with $\mathbf{r}_c$ the centroid of the mesh domain. Finally, events happening at smaller scales than the one imposed by the local diffusivity, $\sqrt{D(\mathbf{r})\Delta t}$, cannot be directly analyzed without earlier knowledge (30). Note that these approximations collapse the three definitions of stochastic integrals previously mentioned onto the same value.

Depending on the physical modeling, on the characteristic scales of the fields and on the statistical properties of the data, many different likelihoods may be formulated (11,13,14,16–18,29,31,32). In this primer, we discuss the likelihood to infer heterogeneous diffusion in a potential field, thus with $\mathbf{F}(\mathbf{r}) = -\nabla V(\mathbf{r})$. The full likelihood reads as follows:

$$P\left(\{T_k\}_{(k\in\mathcal{T})}\middle|\{\nabla V_l\}_{(l\in\mathcal{M})},\{D_l\}_{(l\in\mathcal{M})}\right) \propto \prod_{(l\in\mathcal{M})}\left(\prod_{k\in\mathcal{T}}\prod_{\mu:\mathbf{r}_\mu^k\in S_l}\frac{\exp\left(-\frac{\left(\mathbf{r}_{\mu+1}^k-\mathbf{r}_\mu^k-D_l\nabla V_l\Delta t/k_BT\right)^2}{4\left(D_l+\frac{\sigma^2}{\Delta t}\right)\Delta t}\right)}{\left(4\pi\left(D_l+\frac{\sigma^2}{\Delta t}\right)\Delta t\right)^{\frac{d}{2}}}\right), \tag{6}$$

where $\mathcal{M}$ is the mesh, S_l is a subdomain in the mesh, $\mathcal{T}$ is the set or subset of trajectories, k is the superscript index of a trajectory in $\mathcal{T}$, μ is the subscript time index of a point in the trajectory, σ is the positioning noise modeled as a Gaussian process, and d the spatial dimensionality ($d\in\{1,2,3\}$). The local fluctuation-dissipation relation, $D_l = k_BT/\gamma_l$ is enforced directly in this likelihood. Note that, for a large number of points, the potential inferred converges onto the Boltzman equilibrium distribution (11,13,14,16–18,20,21,29,31,32).

Priors

To perform the analysis, we use as a prior the product of two priors: the Jeffreys prior and the field prior. In our case, the Jeffreys prior reads as follows:

$$P(D_l, V) \propto \frac{D_l^2}{\left(D_l \Delta t + \sigma^2\right)^2}. \tag{7}$$

In a practical manner, the Jeffreys prior prevents from inferring negative diffusion in local confinement domains because of positioning noise (31). The field prior is used to penalize strong gradients of the diffusive and potential fields. This prior reads as follows:

$$P(D(\mathbf{r}), V(\mathbf{r})) \propto \exp\left(-\mu \int d\mathbf{r} \| \nabla D(\mathbf{r}) \|^2 - \lambda \int d\mathbf{r} \| \nabla V(\mathbf{r}) \|^2 \right), \tag{8}$$

where μ and λ are weighting coefficients optimized using simulations matching experimental conditions. Note that the parameters (μ, λ) can also be accessed directly from the experimental trajectories using the parametric empirical Bayesian approach (12,33). This approach is an approximate Bayesian procedure in which the parameters of the prior are estimate directly from the data using Bayes law. The field prior encodes physical constraints on the medium properties and effectively reinforces the robustness of the parameter estimation. Among other things, it can impede the consequences of misconnected single-molecule trajectories. It damps the propagation during optimization of local region of high diffusion or potential that are induced by improper linking between two separate biomolecules.

Finally, the robustness of the results regarding misevaluation of the positioning noise should be tested. To this end, an empirical noise prior is introduced as follows:

$$P(\sigma) = \frac{1}{\sigma_0} \exp\left(-\frac{\sigma}{\sigma_0}\right), \tag{9}$$

where σ_0 is the experimentally measured positioning noise. It is worth noting here that σ_0 includes all sources of noise, i.e., camera shot noise, single-molecule motion blur (34), localization algorithm, etc. The posterior distribution is then computed by numerically integrating the likelihood over all the possible values of σ as follows:

$$P(\{U\}|\{T\}, M) \propto \int P(\{T\}|\{U\}, \sigma, M) P(\{U\}|M) P(\sigma) d\sigma. \tag{10}$$

Note that in both the likelihood and the prior, the physical model of the environment and the experimental conditions (acquisition time, number of translocations, localization error) are mixed allowing a full probabilistic view of the experiments.

Posterior distribution sampling

Estimators of parameters are linked to the Bayesian procedure (6,7,9,10,12). When numerous localizations are available, the posterior distribution is bound to be narrow. Thus, we use the MAP values as estimators. The MAP is evaluated by optimizing the posterior distribution using simplex or quasi-Newton methods in multiple dimensions (35) (other optimization procedures work also efficiently). When the analysis is performed with independent mesh domains, local posterior distributions of diffusivity and forces are evaluated by direct integration. When the analysis is performed on large-scale domains (or full cell surface) with coupling between mesh domains (either in the likelihood or the prior), approximate methods are often employed. As interaction and diffusivity fields are limited in their spatial extension, the posterior distribution sampling was performed by setting all the parameter values outside the area of interest to their MAP values. This procedure was used to perform direct integration of diffusion and force posterior distribution on local areas and also to evaluate confinement energy. Furthermore, a common approximation used when evaluating the posterior distribution of parameters in local domains is the Laplace approximation. This approximation leads to a Gaussian shape of the posterior distribution. The main effect in the case of single-molecule analysis is an overestimation of the error in diffusion analysis, which stems from the asymmetry of the diffusion distribution imperfectly approximated by a Gaussian.

Quantitative elements

The efficiency of the scheme regarding convergence, reliability of the maps, quality of the posterior distribution, and effect of the prior are discussed in great detail in (11,13,14,16–18,29,31,32). A few common rules can be shortly described. Numerous properties are problem dependent, however, generating a mesh with ≈ 20 points per zone ensures a well-behaved posterior distribution and reliably inferred maps. When handling various levels of single-molecule confinement, the Jeffreys prior ensures that no area will have a negative diffusivity induced by the combined effect of the positioning error and confinement. Finally, field priors have a useful property beyond smoothing. Numerous experiments and single-molecule tracking algorithms lead to trajectory misconnections during tracking that can have significant consequences on analyzed properties. Coupling mesh domains together with the field prior ensures that these events have limited statistical weight.

Globally, when handling any new problem, the inference procedure is optimized on numerical simulations. Trajectories are numerically generated in media mimicking both the biological system and the recording properties.

Parameters are then optimized to infer the fields used to generate the trajectories. It is worth noting that single-molecule experiments enjoy a useful property: most possible configurations can be simply and efficiently simulated.

A biological example: dynamics of membrane proteins

We illustrate the application of our Bayesian approach in the case the dynamics of glycine receptors (GlyR) in living cells. The analysis was performed by means of InferenceMAP (29), an open source C++ software that implements the method discussed above. GlyR diffuses on the postsynaptic membrane of neurons and interacts with the scaffold protein gephyrin via an intracellular β-loop. The GlyR-gephyrin interaction is thought to be a key molecular mechanism of synaptic formation and plasticity in the inhibitory neuronal synapse (31,36,37). Gephyrin scaffolds are stable on the timescale of a few minutes and, as previously shown (31), they can be described as potential wells in which receptors are transiently trapped.

To reconstitute the interaction in a nonneuronal system, we used HeLa cells cotransfected with gephyrin and a transmembrane domain possessing the intracellular β-loop (38) and fused to an extracellular photoactivatable fluorescent probe (dendra2). Using sptPALM in total internal reflection fluorescence illumination, we recorded more than 160,000 individual trajectories in a 300 s acquisition. With the described mapping approach, we generated diffusion, force, and potential energy maps spanning the whole cell (Fig. 1). The map is based on an adaptive mesh, generated with an unsupervised clustering technique described in (29). There are more than 4,000 zones in the mesh, with the minimal spatial resolution being ~120 nm. The full analysis required $\approx$30 min. Analysis of individual clusters only requires a few tens of seconds. The main time-consuming step is the optimization of the potential field, which strongly depends on the number of variables. Approximate methods to optimize large-scale fields, such as randomized optimization of mesh subsets (implemented in InferenceMAP (29)), reduce computation time and efficiently converges to the MAP values of the parameters.

The maps display important features. For the diffusion, values range from 0.06 to 0.32 $\mu m^2 \times s^{-1}$. The spatial variation of diffusivities is likely the result (among other things) of variations in lipid content of the membrane, local changes in membrane composition, crowding, and nonspecific interaction with the environment. Importantly, we verify that the posterior distribution of the low- and high-diffusivity regions are disjoint (Fig. 1 *F*), indicating that the measured differences are statistically significant. The map of the force norms reveals small regions of strong directional bias, which corresponds to potential wells of different depths in the potential energy map (Fig. 1 *E*). Importantly, the sites

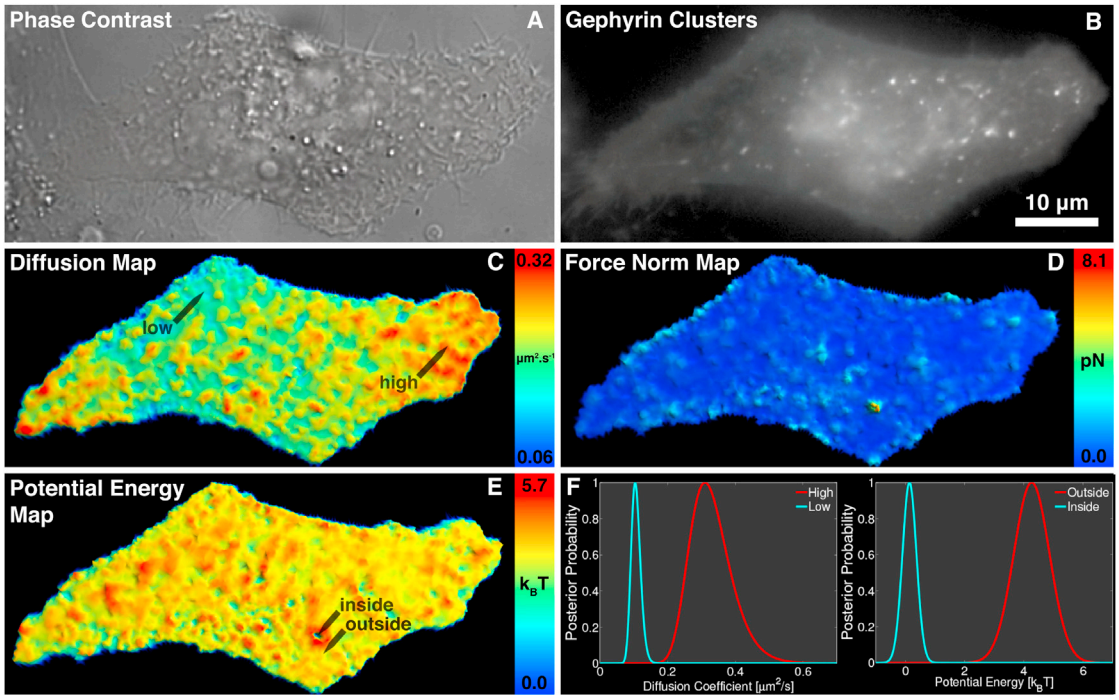

FIGURE 1 Bayesian approach to mapping single-molecule dynamics of a transmembrane protein that interacts with gephyrin clusters via an intracellular β-loop inside a HeLa cell: (*A*) phase contrast image of a HeLa cell; (*B*) fluorescent image of gephyrin clusters, visible as bright spots; (*C*) whole-cell diffusivity; (*D*) force norm; and (*E*) potential energy. (*F*) Example of posterior distributions of the diffusion and potential is shown. Positions are indicated by labeled arrows in (*C*) and (*E*). To see this figure in color, go online.

of these interaction regions are correlated to the location of gephyrin clusters inside the cell, seen in the gephyrin fluorescence image (Fig. 1 *B*). Outside the gephyrin clusters, the force and potential landscape is mostly flat meaning that the motion is predominantly diffusive. We can quantify the depths of the potential wells by sampling the posterior distribution of the potential energy inside and outside of the interaction region (Fig. 1 *F*). The potential energy difference of roughly ~4 kT is evidence of a transient interaction, in which receptors are temporarily confined at the site of a gephyrin cluster.

CONCLUSION

To conclude, Bayesian inference constitutes a natural framework for single biomolecule trajectory analysis. It combines physical modeling of single-molecule motion with hypotheses on data recording into a full probabilistic framework. In future developments, it will be interesting to adapt the Bayesian approach described above to include the internal dynamics of the biomolecule (switching between different diffusive states), the multiscale nature of the environments, non-Markovian processes and, from a computational standpoint, adaptation of the methods to large data sets.

AUTHOR CONTRIBUTIONS

M.E.B., S.T., and M.R. performed the experiments; A.T., M.D., A.A., and J.B.M. designed the experiments; M.E.B., S.T., and J.B.M. programmed the various scheme and inference software; J.B.M. designed the inference procedures; and M.E.B., A.A., M.D., and J.B.M. wrote the article.

ACKNOWLEDGMENTS

We thank Christian Specht and Charlotte Salvatico from the Ecole Normale Superieure, Mathieu Coppey from the Institut Curie, Diego Krapf from Colorado State University, and Brian English and the Transcription Imaging Consortium from Janelia Research Campus for helpful discussions.

This work was supported by funding from the state program "Investissements d'avenir," managed by Agence Nationale de la Recherche (Grant ANR-10-BINF-05 "Pherotaxis" and Grant ANR-10-INSB-04 France BioImaging), the Institut Curie International PhD Program, Paris-Science-Lettres (program ANR-10-IDEX-0001-02 PSL), from ANR Grant SYNAPTUNE, from the Région Ile de France Nanosciences Competence Center, the Délégation Générale de l'Armement, the RTRA Triangle de la Physique, the Bayer Science and Education Foundation, the ERC advanced research grant "PlasltInhib," and the program "Investissements d'Avenir" (ANR-10-LABX-54 MEMO LIFE and ANR-11-IDEX-0001-02 PSL Research University).

REFERENCES

1. Kusumi, A., T. A. Tsunoyama, …, T. K. Fujiwara. 2014. Tracking single molecules at work in living cells. *Nat. Chem. Biol.* 10:524–532.
2. Pinaud, F., S. Clarke, …, M. Dahan. 2010. Probing cellular events, one quantum dot at a time. *Nat. Methods.* 7:275–285.
3. Chenouard, N., I. Smal, …, E. Meijering. 2014. Objective comparison of particle tracking methods. *Nat. Methods.* 11:281–289.
4. Manley, S., J. M. Gillette, …, J. Lippincott-Schwartz. 2008. High-density mapping of single-molecule trajectories with photoactivated localization microscopy. *Nat. Methods.* 5:155–157.
5. Giannone, G., E. Hosy, …, L. Cognet. 2010. Dynamic superresolution imaging of endogenous proteins on living cells at ultra-high density. *Biophys. J.* 99:1303–1310.
6. Harney, H. L. 2010. Bayesian Inference: Parameter Estimation and Decisions. Springer, New York.
7. Barber, D., A. Taylan Cemgil, and S. Chiappa. 2011. Bayesian Time Series Model. Cambridge University Press, New York.
8. Hummer, G. 2005. Position-dependent diffusion coefficients and free energies from Bayesian analysis of equilibrium and replica molecular dynamics simulations. *New J. Phys.* 7:34.
9. Von Toussaint, U. 2011. Bayesian inference in physics. *Rev. Mod. Phys.* 83:943–999.
10. Hines, K. E. 2015. A primer on Bayesian inference for biophysical systems. *Biophys. J.* 108:2103–2113.
11. Masson, J.-B., D. Casanova, …, A. Alexandrou. 2009. Inferring maps of forces inside cell membrane microdomains. *Phys. Rev. Lett.* 102:048103.
12. MacKay, D. J. C. 2003. Information Theory, Inference, and Learning Algorithms. Cambridge University Press, Cambridge, UK.
13. Türkcan, S., and J.-B. Masson. 2013. Bayesian decision tree for the classification of the mode of motion in single-molecule trajectories. *PLoS One.* 8:e82799.
14. Türkcan, S., A. Alexandrou, and J.-B. Masson. 2012. A Bayesian inference scheme to extract diffusivity and potential fields from confined single-molecule trajectories. *Biophys. J.* 102:2288–2298.
15. Metropolis, N., A. Rosenbluth, …, E. Teller. 1953. Equation of state calculations by fast computing machines. *J. Chem. Phys.* 21:1087–1092.
16. Türkcan, S., J.-B. Masson, …, A. Alexandrou. 2012. Observing the confinement potential of bacterial pore-forming toxin receptors inside rafts with nonblinking Eu(3+)-doped oxide nanoparticles. *Biophys. J.* 102:2299–2308.
17. Voisinne, G., A. Alexandrou, and J.-B. Masson. 2010. Quantifying biomolecule diffusivity using an optimal Bayesian method. *Biophys. J.* 98:596–605.
18. Türkcan, S., M. U. Richly, …, J.-B. Masson. 2013. Probing membrane protein interactions with their lipid raft environment using single-molecule tracking and Bayesian inference analysis. *PLoS One.* 8:e53073.
19. Lau, A. W. C., and T. C. Lubensky. 2007. State-dependent diffusion: thermodynamic consistency and its path integral formulation. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 76:011123.
20. Risken, H. 1997. The Fokker-Planck Equation: Methods of Solutions and Applications. Springer, New York.
21. Gardiner, C. W. 2003. Handbook of Stochastic Methods: For Physics, Chemistry and Life Sciences. Springer, New York.
22. Hänggi, P. 1978. Stochastic processes I: asymptotic behaviour and symmetries. *Helvetica Physica Acta.* 51:183–201.
23. Hänggi, P. 1978. Stochastic processes II: response theory and fluctuation theorem. *Helvetica Physica Acta.* 51:202–219.
24. Farago, O., and N. Grønbech-Jensen. 2014. Langevin dynamics in inhomogeneous media: re-examining the Itô-Stratonovich dilemma. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 89:013301.
25. Farago, O., and N. Grønbech-Jensen. 2014. Fluctuation-dissipation relation for systems with spatially varying friction. *J. Stat. Phys.* 156:1093–1110.
26. Fox, R. F. 1987. Stochastic calculus in physics. *J. Stat. Phys.* 46:1145–1157.
27. Kleinert, H. 2006. Path Integrals in Quantum Mechanics, Statistics, Polymer Physics and Financial Markets. World Scientific, Singapore.
28. Schulman, L. S. 2005. Techniques and Applications of Path Integration. Dover Publications, New York.

29. El Beheiry, M., M. Dahan, and J.-B. Masson. 2015. InferenceMAP: mapping of single-molecule dynamics with Bayesian inference. *Nat. Methods.* 12:594–595.

30. Zwanzig, R. 1988. Diffusion in a rough potential. *Proc. Natl. Acad. Sci. USA.* 85:2029–2030.

31. Masson, J.-B., P. Dionne, …, M. Dahan. 2014. Mapping the energy and diffusion landscapes of membrane proteins at the cell surface using high-density single-molecule imaging and Bayesian inference: application to the multiscale dynamics of glycine receptors in the neuronal membrane. *Biophys. J.* 106:74–83.

32. Richly, M. U., S. Türkcan, …, A. Alexandrou. 2013. Calibrating optical tweezers with Bayesian inference. *Opt. Express.* 21:31578–31590.

33. Bishop, C. 2007. Pattern Recognition and Machine Learning. Springer, New York.

34. 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.

35. Press, W. H., S. A. Teukolsky, …, B. P. Flannery. 2002. Numerical Recipes in C++: The Art of Scientific Computing. Cambridge University Press, Cambridge, UK.

36. Tyagarajan, S. K., and J.-M. Fritschy. 2014. Gephyrin: a master regulator of neuronal function? *Nat. Rev. Neurosci.* 15:141–156.

37. Choquet, D., and A. Triller. 2003. The role of receptor diffusion in the organization of the postsynaptic membrane. *Nat. Rev. Neurosci.* 4:251–265.

38. Specht, C. G., I. Izeddin, …, A. Triller. 2013. Quantitative nanoscopy of inhibitory synapses: counting gephyrin molecules and receptor binding sites. *Neuron.* 79:308–321.