Published in final edited form as:

Nat Methods. 2015 September ; 12(9): 838–840. doi:10.1038/nmeth.3483.

Inferring transient particle transport dynamics in live cells

Nilah Monnier^{1,7,8}, Zachary Barry^{1,8}, Hye Yoon Park^{2,3,4,7}, Kuan-Chung Su^{5,6}, Zachary Katz^{2,7}, Brian P English⁴, Arkajit Dey^{1,7}, Keyao Pan¹, Iain M Cheeseman^{5,6}, Robert H Singer^{2,3,4}, and Mark Bathe¹

¹Department of Biological Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

²Department of Anatomy and Structural Biology, Albert Einstein College of Medicine, Bronx, New York, USA.

³Gruss-Lipper Biophotonics Center, Albert Einstein College of Medicine, Bronx, New York, USA.

⁴Howard Hughes Medical Institute, Janelia Research Campus, Ashburn, Virginia, USA.

⁵Whitehead Institute for Biomedical Research, Cambridge, Massachusetts, USA.

⁶Department of Biology, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

Abstract

Live-cell imaging and particle tracking provide rich information on mechanisms of intracellular transport. However, trajectory analysis procedures to infer complex transport dynamics involving stochastic switching between active transport and diffusive motion are lacking. We applied Bayesian model selection to hidden Markov modeling to infer transient transport states from trajectories of mRNA-protein complexes in live mouse hippocampal neurons and metaphase kinetochores in dividing human cells. The software is available at <http://hmm-bayes.org/>.

High-resolution fluorescence imaging is routinely used to report on the dynamic behavior of single molecules and macromolecular assemblies in diverse biological contexts including membrane receptor dynamics^{1, 2, 3}, long-range mRNA transport^{4, 5} and chromosome segregation⁶. Although a variety of advanced particle-tracking techniques are available^{7, 8, 9}, model-based inference procedures are needed to extract mechanistic information from these trajectories. Classification of local particle dynamics using physical motion models offers insights into subtle features of molecular transport such as the direction and speed of

Correspondence should be addressed to M.B. (mark.bathe@mit.edu).

⁷Present addresses: Department of Genetics, Stanford University School of Medicine, Stanford, California, USA (N.M.); Department of Physics and Astronomy, Seoul National University, Seoul, South Korea (H.Y.P.); Salk Institute for Biological Studies, La Jolla, California, USA (Z.K.); Twitter Inc., San Francisco, California, USA (A.D.).

⁸These authors contributed equally to this work.

Author Contributions

N.M. and M.B. conceived the method; N.M., A.D., K.P. and M.B. developed the theory; N.M. and Z.B. implemented the method; H.Y.P., Z.K., B.P.E. and R.H.S. collected mRNP data sets and advised on their analysis; K.-C.S. and I.M.C. collected kinetochore data sets and advised on their analysis; Z.B. and N.M. analyzed experimental data sets; N.M., Z.B. and M.B. wrote the paper.

Competing financial interests

The authors declare no competing financial interests.

molecular motors that drive transport of intracellular cargo, as well as the identification of specific intracellular locations of cargo confinement or transient binding interactions^{2, 3}. Because intracellular transport often exhibits a high degree of heterogeneity depending on the spatial location of cargo within the cell, combinations of motion models should ideally be considered in the analysis and annotation of single-particle trajectories.

Although kymograph and mean-square displacement (MSD) analyses are commonly used to characterize intracellular motion from particle trajectories^{1, 10}, hidden Markov modeling (HMM) is a powerful alternative owing to its ability to annotate heterogeneous motion locally along a single trajectory^{2, 11, 12, 13}. In contrast to MSD analysis, HMMs account for the possibility of stochastic switching between distinct motion states with single-step temporal resolution without time averaging. This advantage of the HMM approach has been demonstrated for diffusing particles in the analysis of single-receptor dynamics confined by membrane corrals and undergoing transient cytoskeletal-binding interactions^{2, 11, 12} and RNA-binding protein dynamics in bacteria¹³. Incorporation of Bayesian model selection into the inference process additionally enables objective selection of the simplest stochastic motion model that describes a given trajectory¹³. However, existing Bayesian HMMs are limited to modeling purely diffusive motion, whereas intracellular cargo often exhibit combinations of active transport and random diffusive motion. An important example is long-range transport of mRNAs in complex with mRNA-binding proteins (mRNPs), driven by molecular motors along microtubule tracks in neuronal dendrites⁵.

Long-range transport of β -actin mRNP complexes to sites of local protein translation in neurons is implicated in synapse formation and plasticity during development and learning⁵. In live neuronal cultures^{4, 5}, endogenous β -actin mRNP particles undergo heterogeneous periods of anterograde and retrograde transport interspersed with pausing events, with a moderate bias toward anterograde transport⁵. Kymographs of β -actin mRNPs qualitatively confirmed the presence of both stationary and active transport phases, indicating that transport is not fully processive (Fig. 1a, Supplementary Fig. 1 and Supplementary Note 1). Extracting quantitative information from kymographs is a subjective process, however, particularly for short-lived phases of motion (Supplementary Fig. 1). Quantitative analyses of mRNP trajectories using MSD curves averaged within local time windows along each trajectory¹⁴ provided additional evidence for multiple phases of motion (Fig. 1b, Supplementary Fig. 2 and Supplementary Note 1). However, the intrinsically limited temporal resolution of MSD-based techniques that require sliding-window averaging made them unable to resolve short-lived phases of motion, and the application of these techniques yielded variable results depending on user-selected parameters such as window size (Supplementary Fig. 2 and Supplementary Note 1). Although HMM-based procedures can in principle resolve distinct motion states with single-time-step resolution, purely diffusive HMM approaches resulted in erroneous annotations (Fig. 1c) because they neglect the possibility of active transport in the underlying set of motion models considered.

To overcome these limitations, we developed a versatile HMM procedure that can be applied both to diffusive switching and to active transport processes interspersed with random pausing events such as in the motor-driven transport of mRNPs. This HMM analysis approach requires the statistical hypothesis that a particle explores a finite set of diffusive

and directed transport motion states whose switching can be modeled as a Markov process. Our procedure, HMM-Bayes, models diffusive and directed motion states along particle trajectories and performs Bayesian model selection to infer the simplest stochastic motion model that is consistent with the observed particle displacements (Supplementary Figs. 3–6). The procedure can be applied to either a single trajectory or a set of pooled trajectories, annotating intermittent periods of diffusive and directed motion locally along each trajectory to reveal when and where switching between distinct types of motion occurs in space and time (Figs. 1d and 2a,b). We validated our approach using simulated trajectories, first confirming that it performed similarly to purely diffusive HMMs when applied to stochastic switching between two distinct diffusion states (Supplementary Figs. 7–9 and Supplementary Notes 2 and 3). In addition, our procedure detected stochastic switching between directed transport and random diffusive motion states (Supplementary Figs. 10–14 and Supplementary Note 3), whereas existing HMM approaches failed to detect transport. As expected¹³, a minimum number of observed particle displacements was required to infer the presence of multiple motion states, and this number depended on the relative values of the motion parameters (Supplementary Figs. 11 and 12, and Supplementary Note 3).

The ability of our procedure to detect directed transport enabled its application to annotate complex neuronal β -actin mRNP transport dynamics (Figs. 1b and 2, Supplementary Note 4, Supplementary Figs. 15–17 and Supplementary Videos 1,2,3). Our procedure yielded trajectories annotated with the local spatial-temporal dynamics of when and where each mRNP exhibited retrograde, anterograde or pausing motion along its trajectory within the cell (Figs. 1d and 2a,b), as well as the lifetime of each individual period of motion and population distributions of velocities, diffusion coefficients and state lifetimes across a collection of heterogeneous mRNP trajectories from multiple cells (Fig. 2c). We found that β -actin mRNPs existed on average in passive pausing states longer than in active transport states, and that anterograde transport had a higher average velocity ($0.76 \pm 0.45 \mu\text{m/s}$, mean $\pm$ s.d.) than retrograde transport ($0.58 \pm 0.35 \mu\text{m/s}$), with large variability in single-molecule mRNP transport rates. Analysis of β -actin mRNP trajectories in cells treated with KCl, which induced inhibition of active transport⁵, confirmed that our procedure annotated these trajectories with only a single diffusive state (Supplementary Fig. 18). The quantifications produced by our analysis procedure, combined with secondary labeling of microtubules, cytoskeletal-associated proteins and ribosomes, may facilitate interpretation of the molecular origins of these heterogeneous β -actin mRNP transport dynamics in future studies.

To confirm the applicability of HMM-Bayes across multiple modes of intracellular motion in biological systems, we applied it to β -actin mRNP trajectories in live mouse fibroblasts¹⁵ and to oscillating kinetochores in dividing HeLa cells¹⁶. As with previous HMM procedures¹³, trajectories can be pooled to improve inference power (Supplementary Figs. 19–22), which enabled the detection of multiple diffusive states of fibroblast β -actin mRNPs (Supplementary Fig. 23 and Supplementary Note 5) and multiple transport states of oscillating kinetochores without the need for manual identification of direction reversals (Fig. 3, Supplementary Figs. 24 and 25, and Supplementary Note 6). Pooling assumes that each trajectory consists of the same set of underlying hidden motion states and parameter values; therefore, transport phases annotated within pooled trajectories are assumed to have the same direction and magnitude of transport velocity across the trajectories. In some cases

trajectories can be projected along specific directions of interest before pooling, as in the analysis of kinetochore dynamics. When transport occurs in random directions across a set of pooled trajectories, a modified version of HMM-Bayes that incorporates χ^2 emission distributions can be applied (Supplementary Note 7 and Supplementary Figs. 26 and 27). Future applications of HMM-Bayes may explore the one-dimensional (1D) projection of 2D and 3D transport along curvilinear objects such as microtubules, or the use of directional statistics to model changes in direction of transport. In addition, hierarchical or k -means clustering of trajectories could be used iteratively with HMM-Bayes to identify similar subsets of trajectories on the basis of common motion types and parameter values.

In conclusion, consideration of directed motion in the HMM annotation process for particle trajectories is important because biological transport is often driven by active motor and cytoskeletal processes. Our procedure annotates stochastic phases of directed transport and random motion. Quantitative information on rates of transport, directional switching, and locations and durations of pausing can be used to explore complex mechanisms of intracellular transport of biological particles. Because the formulation of our procedure is general to single-particle motion analysis, it should also apply to the analysis of cell migration in tissue culture models¹⁷ and developing embryos^{18, 19}. The HMM-Bayes software package is available in open source (<http://hmm-bayes.org/> and Supplementary Software) and benefits from parallel computing when large biological data sets are considered.

Methods

Formulation of a particle-trajectory HMM with directed motion

Following previous approaches^{11, 12, 13}, we model particle displacements as emissions and motion models as hidden states of a hidden Markov model (HMM). The emission distributions for the observed displacements are normal distributions whose parameters depend on the hidden motion states, as described in more detail below. The s.d. of the emission distributions depend on the diffusion coefficients of the motion states and on localization error in the measurement of particle positions (Supplementary Note 2). In contrast to previous HMM procedures that assume diffusive motion states with zero mean for the displacement emission distributions^{11, 12, 13}, here we allow for active transport states with nonzero mean. This generalization to nonzero means to account for directed motion introduces multiple additional parameters (Supplementary Table 1) that can easily lead to overfitting of the trajectories and, consequently, erroneous trajectory annotation. We therefore implement Bayesian model selection, also described in more detail below, both on the number of hidden states and on the inclusion of the nonzero velocity parameter of appropriate dimensionality (1D, 2D, 3D) within each motion state, thus considering a considerably larger set of complex motion-switching models (Supplementary Table 1) than previous approaches.

In a single dimension, a particle trajectory consists of a sequence of particle positions x_t separated by a time interval Δt . For a particle undergoing a random walk with drift and in the absence of localization error²⁰, the particle displacements $\Delta x_t = x_{t+1} - x_t$ along this dimension follow a normal distribution with an s.d. that depends on the diffusion coefficient

D according to $\sigma = (2D \Delta t)^{1/2}$ and a mean that depends on the velocity v_x according to $\mu_x = v_x \Delta t$. For a 2D or 3D particle trajectory with particle positions $\mathbf{r}_t = \{x_t, y_t\}$ or $\mathbf{r}_t = \{x_t, y_t, z_t\}$, the displacements become $\mathbf{r}_t = \{x_t, y_t\} = \{x_{t+1} - x_t, y_{t+1} - y_t\}$ or $\mathbf{r}_t = \{x_t, y_t, z_t\} = \{x_{t+1} - x_t, y_{t+1} - y_t, z_{t+1} - z_t\}$ and the velocity of the particle has multiple components, $\mathbf{v} = \{v_x, v_y\}$ or $\mathbf{v} = \{v_x, v_y, v_z\}$. In this work we assume isotropic diffusion, so the diffusion coefficient is always a scalar value. Under these conditions, the displacements are distributed according to a multivariate normal distribution; for example, in three dimensions

$$\Delta \mathbf{r} \sim \frac{1}{(4\pi D \Delta t)^{3/2}} \exp \left(-\frac{(\Delta x - v_x \Delta t)^2 + (\Delta y - v_y \Delta t)^2 + (\Delta z - v_z \Delta t)^2}{4D \Delta t} \right) \quad (1)$$

Error in the experimental measurement of particle positions (localization error) contributes to the s.d. of the observed displacements and must be taken into account when converting between σ and the diffusion coefficient (Supplementary Note 2).

In modeling single-particle transport dynamics in biological systems, we assume that particles may experience local changes in their diffusion coefficient, velocity or both due to interactions with their surroundings. To infer changes in these motion parameters from an observed sequence of particle displacements, we model particle displacements using an HMM, a type of Bayesian network that consists of a Markov chain of hidden variables, $\mathbf{s} = \{s_1, \dots, s_T\}$, and a corresponding sequence of observed variables or emissions, $\mathbf{e} = \{e_1, \dots, e_T\}$, where T is the number of observations. The hidden variables s_t can each take a discrete set of values or states, $\{S_1, \dots, S_K\}$, where K is the number of available states²¹. The probability of transitioning to a particular state at time t depends only on the previous state at time $t - 1$ and is denoted by $\Phi_{ij} = \Pr(s_t = S_j | s_{t-1} = S_i)$, and the probability of starting in state S_i at time $t = 1$ is denoted by $\pi_i = \Pr(s_1 = S_i)$. To model particle trajectories, we let the hidden states be the unobserved motion states of the particle characterized by the parameters D and $\mathbf{v}$; thus, each hidden state S_i represents a particular state of motion with a specific diffusion coefficient and velocity, $S_i = \{D_i, \mathbf{v}_i\}$. Because a particle may also experience periods of pure diffusion without directed transport, hidden motion states with zero velocity and only a single motion parameter, $S_i = \{D_i\}$, must also be considered in the model.

At every time point t , the particle is assumed to exist in one of these possible hidden states, but the specific state is unknown a priori. Instead, we observe a time series of particle displacements $\{\mathbf{r}_1, \dots, \mathbf{r}_T\}$, where T is the total number of time intervals over which the particle is tracked. These particle displacements are modeled as the emissions of the HMM. The probability of observing a particular emission at time t depends only on the state at time t according to a probability distribution $p_{S_i}(e_t) = \Pr(e_t | s_t = S_i)$, which is in this case the normal distribution in equation (1) above, parameterized by the unobserved motion parameters $\{D_i, \mathbf{v}_i\}$ of the hidden state. Thus, any local changes in particle transport dynamics are inferred using the temporal information that is encoded in the time series of emissions.

The full set of parameters, denoted θ , for an HMM with K motion states includes the length- K vector of starting probabilities π_i , the $K \times K$ matrix of transition probabilities ϕ_{ij} between all pairs of states and the motion parameters $\{D_i, \mathbf{v}_i\}$ that characterize each hidden state

$$\theta = \left\{ \{\pi_i\}_{i=1}^K, \{\phi_{ij}\}_{i,j=1}^K, \{D_i, \mathbf{v}_i\}_{i=1}^K \right\} \quad (2)$$

These parameter values are inferred for different values of K on the basis of the observed sequence of particle displacements during the model fitting and selection process described below.

Model-selection framework for the particle-trajectory HMM

When analyzing a particle trajectory, we do not know a priori either the number of motion states K that the particle explores during the time it is observed or the motion parameters associated with these states. Here we use a Bayesian model-selection approach to determine the appropriate number of motion states K without overfitting. We also assume that each of the K motion states may or may not include directed motion with nonzero velocity in addition to random diffusive motion. Because inclusion of a nonzero velocity parameter can also lead to overfitting, we use our Bayesian model-selection framework to choose both the appropriate K and the number of states within K that have nonzero velocities (directed transport states). For example, we test three possible two-state models—one with two diffusive states (model D-D), one with one diffusive state and one transport state (model D-DV), and one with two transport states (model DV-DV)—and penalize the total number of parameters, which increases with the number of transport states. A full set of tested models and their associated numbers of parameters is shown in Supplementary Table 1 for HMMs with up to three hidden states. Note that for a given number of states K , there are $K + 1$ possible models to be tested, depending on how many of the K states are transport states with nonzero velocity. Our model-selection procedure evaluates the relative probabilities of these competing models up to some specified maximum number of states $K_{\max}$. The total number of models tested for a particular value of $K_{\max}$ is thus $(K_{\max}^2 + 3K_{\max})/2$, considerably larger than in the case of purely diffusive HMMs¹³, which test only a single motion-switching model for each number of states K .

Given an observed sequence of particle displacements, the posterior probability of a particular model M_k can be expanded according to Bayes' theorem,

$$\Pr(M_k|\mathbf{e}) = \frac{\Pr(\mathbf{e}|M_k)\Pr(M_k)}{\Pr(\mathbf{e})} \propto \Pr(\mathbf{e}|M_k) \quad (3)$$

where the emissions $\mathbf{e} = \{e_1, \dots, e_T\}$ are the displacements $\{\mathbf{r}_1, \dots, \mathbf{r}_T\}$, as described above, and the final proportionality holds if the prior probabilities of the models, $\Pr(M_k)$, are assumed equal owing to the absence of prior information about the system. Thus, with uniform prior probabilities, only the likelihood $\Pr(\mathbf{e}|M_k)$ must be calculated for each model to determine the relative model probabilities. This likelihood is found by marginalizing over the unknown parameter values θ_k for each model M_k ,

$$\Pr(\mathbf{e}|M_k) = \int \Pr(\mathbf{e}|\boldsymbol{\theta}_k, M_k) \Pr(\boldsymbol{\theta}_k|M_k) d\boldsymbol{\theta}_k \quad (4)$$

where $\Pr(\boldsymbol{\theta}_k|M_k)$ is the prior probability of a particular realization $\boldsymbol{\theta}_k$ of the parameter values for the model M_k , and $\Pr(\mathbf{e}|\boldsymbol{\theta}_k, M_k)$ is the likelihood of the observed emissions given the model and that particular set of parameters. For an HMM, the likelihood $\Pr(\mathbf{e}|\boldsymbol{\theta}_k, M_k)$ must also be marginalized over all possible hidden-state sequences $\mathbf{s}_k$ for the model M_k , as the temporal sequence of hidden states is unknown

$$\Pr(\mathbf{e}|\boldsymbol{\theta}_k, M_k) = \sum_{\mathbf{s}_k} \Pr(\mathbf{e}|\mathbf{s}_k, \boldsymbol{\theta}_k, M_k) \Pr(\mathbf{s}_k|\boldsymbol{\theta}_k, M_k) \quad (5)$$

Because of the HMM structure²¹, the first term depends only on the emission probability distributions according to $\Pr(\mathbf{e}|\mathbf{s}_k, \boldsymbol{\theta}_k, M_k) = \prod_{t=1}^T p_{s_t}(e_t)$, where the emission distributions are parameterized by the D_i and $\mathbf{v}_i$ values for each state S_i as described above, and the second term depends only on the starting and transition probabilities according to $\Pr(\mathbf{s}_k|\boldsymbol{\theta}_k, M_k) = \pi_{s_1} \prod_{t=2}^T \Phi_{s_{t-1}s_t}$. Substituting into equation (4) above, the full equation for the marginal likelihood in terms of the model parameters becomes

$$\Pr(\mathbf{e}|M_k) = \int \left[\sum_{\mathbf{s}_k} \left(\pi_{s_1} \prod_{t=2}^T \Phi_{s_{t-1}s_t} \prod_{t=1}^T P_{s_t}(e_t) \right) \right] \Pr(\boldsymbol{\theta}_k|M_k) d\boldsymbol{\theta}_k \quad (6)$$

Although the integral in equation (6) is intractable in general, the value of the integrand can be evaluated exactly for a given model M_k at any particular value of its associated parameters $\boldsymbol{\theta}_k$, given a uniform parameter prior distribution $\Pr(\boldsymbol{\theta}_k|M_k)$ and using the forward algorithm²¹ to evaluate the summation over hidden-state sequences. Therefore, to compute the likelihood $\Pr(\mathbf{e}|M_k)$ for each model M_k , we use a numerical integration approach in which the integrand is evaluated at stochastically sampled values of $\boldsymbol{\theta}_k$ as described below.

Numerical integration of the likelihood

We use Markov chain Monte Carlo (MCMC), specifically the Metropolis MCMC algorithm^{22, 23, 24}, to sample the posterior distribution in equation (6). The value of the integrand in equation (6), $f(\boldsymbol{\theta}_k)$, is evaluated at a random starting point $\boldsymbol{\theta}_k^{(0)}$ in parameter space, then new parameters $\boldsymbol{\theta}_k^{(1)}$ are proposed from a multivariate normal distribution in parameter space centered at $\boldsymbol{\theta}_k^{(0)}$, and the new $\boldsymbol{\theta}_k^{(1)}$ is accepted with probability $\min(1, f(\boldsymbol{\theta}_k^{(1)})/f(\boldsymbol{\theta}_k^{(0)}))$. Additional details on these steps are given below. Repeating this process results in sampling a stationary distribution of parameters that is proportional to $f(\boldsymbol{\theta}_k)$.

For our particle-trajectory HMM with directed transport, the full set of parameters to be sampled is given in equation (2) above. As discussed in the previous section, however, we also consider models in which some of the hidden motion states are purely diffusive states with zero velocity. For any given tested model M_k , let K be the total number of states and K_V be the number of states with nonzero velocity, where $0 < K_V < K$. Then the full set of parameters that must be sampled is $\boldsymbol{\theta}_k = \{\pi_1, \dots, \pi_K, \Phi_{11}, \dots, \Phi_{KK}, D_1, \dots, D_K, \mathbf{v}_1, \dots, \mathbf{v}_{K_V}\}$. At each step in an MCMC run, one or more parameters in this set can be selected for updating.

We found that updating a randomly selected block of related parameters at each step²⁴ exhibited the fastest and most robust convergence compared with other move-sets, such as updating all parameters simultaneously or updating only a single randomly selected parameter at each step. Motivation for the block updating approach is that parameters with correlated effects on $f(\theta_k)$ should be updated at the same time to increase the probability of escaping from local maxima in the likelihood landscape²⁴. Here we split the parameters into three blocks: the probability parameters $\{\pi_1, \dots, \pi_K, \Phi_{11}, \dots, \Phi_{KK}\}$, the diffusion coefficients $\{D_1, \dots, D_K\}$ and the velocities $\{\mathbf{v}_1, \dots, \mathbf{v}_{K_V}\}$. At each step, one of these blocks is selected randomly with equal probability, updates are proposed for all of the parameters within that block, and the update is accepted or rejected according to the Metropolis criterion above.

To initialize each MCMC run, we choose a random point in parameter space that is within a range defined by the set of observed particle displacements. The maximum-likelihood velocity parameters in each dimension will not be greater than the largest observed single-displacement velocity or less than the smallest observed single-displacement velocity for that dimension. Therefore, initial guesses for the velocity parameters $v_{x,i}$, $v_{y,i}$, and $v_{z,i}$ for each state S_i in $\{S_1, \dots, S_{K_V}\}$ are drawn from uniform distributions on this range; for example, $v_{x,i}^{(0)} \sim U(\min_t(x_t)/t, \max_t(x_t)/t)$. Because the velocities are related to the means of the emission probability distributions by $\mu_{x,i} = v_{x,i} t$, in practice we fit the means $\{\mu_1, \dots, \mu_{K_V}\}$ and convert to velocities as a final step in the analysis. Similarly, the diffusion coefficients $\{D_1, \dots, D_K\}$ are related to the s.d. of the emission distributions by $\sigma_i = (2D_i t)^{1/2}$, so in practice we fit the s.d. $\{\sigma_1, \dots, \sigma_K\}$. These σ parameters must be greater than 0 and will generally not be greater than the largest observed width of the displacement distribution across the dimensions, so we draw initial guesses from a uniform distribution, $\sigma_i^{(0)} \sim U(0, \max_t(\xi_t) - \min_t(\xi_t))$, where ξ parameterizes the spatial dimensions of the particle trajectory, for example, one of $\{x, y, z\}$ for a 3D trajectory. Finally, initial guesses for the probability parameters $\{\pi_1, \dots, \pi_K, \Phi_{11}, \dots, \Phi_{KK}\}$ are all set to the same value $1/K$. Multiple restarts of MCMC are run with different initial guesses for the mean and s.d. parameters, with at least 100 restarts for models with more than a single state.

During each MCMC run from a particular initial parameter guess, parameter values are updated using a normal proposal distribution centered on the previous value of each parameter and with an s.d. δ that is specific to each parameter type, denoted as δ_μ , δ_σ , δ_π and δ_Φ . These proposal distribution widths are set adaptively during the series of MCMC initialization runs to maintain a target acceptance rate, computed separately for each block of parameters, of between 0.3 and 0.5, which is generally recommended for efficiently exploring the likelihood landscape²⁵. We initialize δ_μ and δ_σ to 1/50 of the initial guess ranges given above for the μ and σ parameters and then update δ_μ and δ_σ at the end of each initialization run on the basis of the acceptance rates calculated for that run. Because the σ parameters are constrained to be positive, parameter updates are automatically rejected if a σ value below 0 is proposed. The probability parameters are constrained in the range $[0,1]$ and must satisfy the conditions $\sum_{i=1}^K \pi_i = 1$ and $\sum_{j=1}^K \Phi_{ij} = 1$ for all i ; therefore, these probabilities are renormalized after every update of the probability block, and the values of δ_π and δ_Φ are updated on the basis of the minimum distance of the π_i or Φ_{ij} probabilities,

respectively, from either of the boundaries 0 or 1 to maintain a relatively consistent acceptance rate along the MCMC chain.

Following the MCMC initialization runs, a longer MCMC run is performed, starting from the parameter values that yielded the highest likelihood during the initialization runs, and is subsequently used to report the maximum-likelihood parameter values for each tested model as well as to define the sampling distribution for numerical integration of the desired integral in equation (6). For a sampling distribution $q(\theta_k)$, the estimator for the value of the integral is equal to the mean value of the ratio of the integrand $f(\theta_k)$ to $q(\theta_k)$ over the sampled values of θ_k (ref. 22),

$$\hat{I}_k = \left\langle \frac{f(\theta_k)}{q(\theta_k)} \right\rangle_q \quad (7)$$

where the subscript q indicates that the mean is calculated over values of θ_k sampled from the distribution $q(\theta_k)$. The sampling distribution that gives an estimator with minimum variance²² is the distribution proportional to $f(\theta_k)$, which is approximated by the MCMC samples. We use the MCMC samples to define a sampling distribution $q(\theta_k)$ in which the parameters $\{\sigma_1, \dots, \sigma_K\}$ and $\{\mu_1, \dots, \mu_{K_V}\}$ are each sampled from a normal distribution with the same mean and s.d. as the sampled values of that parameter during the second half of the long MCMC run, so that these parameters are sampled from the important regions of parameter space identified by MCMC. The probability parameters $\{\pi_1, \dots, \pi_K, \phi_{11}, \dots, \phi_{KK}\}$ are sampled from K -dimensional simplexes to ensure that each sampled point meets the conditions $\sum_{i=1}^K \pi_i = 1$ and $\sum_{j=1}^K \phi_{ij} = 1$.

This integration approach results in an estimate from equation (7) of the integral in equation (6). The probability density $q(\theta_k)$ of the sampling distribution is computed analytically for each sampled point θ_k using the normal distributions and uniform simplexes above. The value of $f(\theta_k)$ is computed for each sampled point θ_k using the forward algorithm as discussed above, and assuming that the prior probability $\Pr(\theta_k|M_k)$ in equation (6) is constant over a bounded region in parameter space. For the probability parameters, which are inherently bounded on K -dimensional simplexes, the prior is uniform over each simplex. For the emission mean and s.d. parameters, the uniform bounded region is centered on the mean of the sampled values of that parameter during the second half of the long MCMC run above, with a width equal to 200 times the s.d. of the MCMC samples, as in previous work^{10, 26}. Finally, after computing $\Pr(\mathbf{e}|M_k)$ for each model, the final model probabilities in equation (3) are calculated by normalizing these values of $\Pr(\mathbf{e}|M_k)$ across the set of tested models.

Pooling multiple trajectories

The preceding sections describe the application of our procedure to an individual particle trajectory. However, in many biological applications, large numbers of short trajectories are obtained from independent particles that may be assumed to undergo the same dynamical processes with the same motion parameters for the purposes of multiple-hypothesis testing. When this assumption is valid, pooling the trajectories to perform a joint analysis increases the inference power of the Bayesian HMM approach and therefore its ability to resolve

complex models with multiple states and parameters¹³. Including multiple independent trajectories in the likelihood calculation is straightforward, as their individual likelihoods can be multiplied¹³. For W trajectories, each with a set of observed displacements or emissions $\mathbf{e}_w$, Bayes' rule becomes

$$\Pr(M_k | \{\mathbf{e}_w\}_{w=1}^W) = \frac{\Pr(\{\mathbf{e}_w\}_{w=1}^W | M_k) \Pr(M_k)}{\Pr(\{\mathbf{e}_w\}_{w=1}^W)} \propto \Pr(\{\mathbf{e}_w\}_{w=1}^W | M_k) \quad (8)$$

Analogous to equation (4) above, the marginal likelihood is then

$$\Pr(\{\mathbf{e}_w\}_{w=1}^W | M_k) = \int \Pr(\{\mathbf{e}_w\}_{w=1}^W | \boldsymbol{\theta}_k, M_k) \Pr(\boldsymbol{\theta}_k | M_k) d\boldsymbol{\theta}_k \quad (9)$$

which, assuming independent trajectories, becomes

$$\Pr(\{\mathbf{e}_w\}_{w=1}^W | M_k) = \int \left[\prod_{w=1}^W \Pr(\mathbf{e}_w | \boldsymbol{\theta}_k, M_k) \right] \Pr(\boldsymbol{\theta}_k | M_k) d\boldsymbol{\theta}_k \quad (10)$$

The calculation of each individual $\Pr(\mathbf{e}_w | \boldsymbol{\theta}_k, M_k)$ in equation (10) still follows equation (5) above. The MCMC sampling approach above now explores the parameters $\boldsymbol{\theta}_k$ for each model M_k based on the full set of pooled trajectories rather than for each trajectory individually. Therefore, this pooling approach requires that the trajectories explore the same set of possible motion states with the same motion parameters. The hidden-state sequences $\mathbf{s}_w$, on the other hand, are marginalized out by the forward algorithm, and the most likely state sequence can be obtained for each trajectory independently using the most likely parameters of the most likely model. Thus, the pooled trajectories are not required to have the same hidden-state sequences and can transition between the different motion states at different times.

Simulations of single-particle trajectories

Diffusive single-particle trajectories were simulated by drawing random step lengths in each dimension from a normal distribution with zero mean and s.d. equal to $(2D \Delta t)^{1/2}$, where D is the diffusion coefficient and Δt is the time interval for each step. Directed motion was modeled by adding a fixed displacement of length $\mathbf{v} \Delta t$ to the diffusive component of motion at each time step, where $\mathbf{v}$ is the velocity vector.

Imaging mRNA in live neurons

Recent advances in live-cell mRNA fluorescence labeling techniques have enabled the visualization of β -actin mRNA transport dynamics in live cells using the MS2 bacteriophage capsid protein system^{4, 5}. GFP-tagged MS2 capsid proteins (MCP-GFP) associate with tandem RNA stem-loops knocked into the 3' untranslated region (UTR) of the β -actin gene in the β -actin-MS2-binding site knock-in (MBS) mouse⁴. All experiments using animals were carried out under the approval of the Albert Einstein College of Medicine Institutional Animal Care and Use Committee (IACUC). We used 6- to 8-week-old male and female

mice to set up timed pregnancies of double-homozygous MCP $\times$ MBS mice. Hippocampal neurons were cultured from MCP $\times$ MBS mouse pups at postnatal day 0–2 as described previously⁵. Briefly, we dissected out hippocampi, dissociated them with trypsin and plated 85,000 cells onto poly(D-lysine)-coated dishes (MatTek). The cultures were maintained in Neurobasal-A medium supplemented with B-27, GlutaMAX and Primocin (InvivoGen) at 37 °C and 5% CO₂ for 14–22 d before imaging. For live neuron imaging, we removed the medium from cell cultures and replaced it with HEPES-buffered solution (HBS) containing 20 mM HEPES-HCl, pH 7.4, 119 mM NaCl, 5 mM KCl, 2 mM CaCl₂, 2 mM MgCl₂ and 30 mM glucose before the experiment. Time-lapse images were taken on an Olympus IX-71 inverted microscope with a UApo 150 $\times$, 1.45-numerical aperture (NA) oil-immersion objective (Olympus), an MS-2000 XYZ automated stage (ASI) and an iXon electron-multiplying charge-coupled device (EMCCD) camera (Andor). The cells were kept at 37 °C with 60% humidity in an environmental chamber (Precision Plastics). The GFP was excited by the 488-nm line from an argon ion laser (Melles Griot). Emission was filtered with a 525/30 band-pass filter (Semrock). Wide-field images were acquired at 10 f.p.s. using MetaMorph software (Molecular Devices). Individual localizations of mRNPs were obtained via single-particle tracking using the u-track-2.0 package⁷.

Imaging mRNA in live fibroblasts

Mouse embryonic fibroblast cell lines derived from the MBS mouse stably express tdMCP-GFP to label all endogenous β -actin mRNA as described previously¹⁵. Microtubules were labeled by transient expression of mCherry- α -tubulin. mRNPs were visualized with TIRF excitation on the same microscope setup described above at 35 ms per frame in streaming acquisition mode. mRNP trajectories were analyzed with the DiaTrack 3.03 package²⁷, and coordinates were exported for HMM analysis.

Imaging metaphase kinetochore dynamics in tissue culture cells

HeLa cells (Cheeseman lab) were maintained in DMEM supplemented with 100 U/ml streptomycin, 100 U/ml penicillin, 2 mM glutamine and 10% (vol/vol) FCS. Cell lines were routinely validated and checked for mycoplasma contamination using Mycoalert (Lonza). Cells were cultured at 37 °C with 5% CO₂. Cells expressing GFP-LAP fusions to CSAP²⁸ and CENP-A²⁹ were generated using retroviral infection of HeLa cells with pBABE-blasticidin-based vectors. Images were acquired on a Nikon eclipse microscope equipped with a CCD camera (Clara, Andor). For time-lapse imaging, cells were imaged in CO₂-independent medium (Invitrogen), supplemented as above for the DMEM culture medium, at 37 °C. Images were acquired every 6 s using three z sections at 0.7- μ m intervals using a 40 $\times$, 1.3-NA Plan Fluor objective (Nikon). GFP fluorescence was observed using appropriate filters. Kinetochore positions were manually tracked using the MTrackJ program³⁰.

Code availability

The HMM-Bayes software package and associated documentation are available as Supplementary Software and online at <http://hmm-bayes.org/>.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Research reported in this publication was supported by grants from the US National Institutes of Health (NIH) National Institute of Mental Health (U01 MH106011) and the US National Science Foundation Physics of Living Systems (PHY 1305537) to M.B., an NIH grant from the National Institute of Neurological Diseases and Stroke (NS083085-19) to R.H.S., a Scholar award from the Leukemia & Lymphoma Society and an NIH grant from the National Institute of General Medical Sciences (GM088313) to I.M.C., and a Schroedinger fellowship from the Austrian Science Fund to K.-C.S. We also thank M. Linden for helpful discussions.

References

1. Saxton MJ, Jacobson K. Annu. Rev. Biophys. Biomol. Struct. 1997; 26:373–399. [PubMed: 9241424]
2. Chung I, et al. Nature. 2010; 464:783–787. [PubMed: 20208517]
3. Jaqaman K, et al. Cell. 2011; 146:593–606. [PubMed: 21854984]
4. Lionnet T, et al. Nat. Methods. 2011; 8:165–170. [PubMed: 21240280]
5. Park HY, et al. Science. 2014; 343:422–424. [PubMed: 24458643]
6. Vladimirov E, et al. Dev. Cell. 2013; 27:60–71. [PubMed: 24135231]
7. Jaqaman K, et al. Nat. Methods. 2008; 5:695–702. [PubMed: 18641657]
8. Chenouard N, et al. Nat. Methods. 2014; 11:281–289. [PubMed: 24441936]
9. Saxton MJ. Nat. Methods. 2014; 11:247–248. [PubMed: 24577274]
10. Monnier N, et al. Biophys. J. 2012; 103:616–626. [PubMed: 22947879]
11. Das R, Cairo CW, Coombs D. PLoS Comput. Biol. 2009; 5:e1000556. [PubMed: 19893741]
12. Cairo CW, et al. J. Biol. Chem. 2010; 285:11392–11401. [PubMed: 20164196]
13. Persson F, Lindén M, Unoson C, Elf J. Nat. Methods. 2013; 10:265–269. [PubMed: 23396281]
14. Arcizet D, Meier B, Sackmann E, Rädler JO, Heinrich D. Phys. Rev. Lett. 2008; 101:248103. [PubMed: 19113674]
15. Katz ZB, et al. Genes Dev. 2012; 26:1885–1890. [PubMed: 22948660]
16. Cheeseman IM, Desai A. Nat. Rev. Mol. Cell Biol. 2008; 9:33–46. [PubMed: 18097444]
17. Deforet M, et al. Nat. Methods. 2012; 9:1081–1083. [PubMed: 23064519]
18. Tomer R, Khairy K, Amat F, Keller PJ. Nat. Methods. 2012; 9:755–763. [PubMed: 22660741]
19. Dray N, et al. Curr. Biol. 2013; 23:1335–1341. [PubMed: 23810535]
20. Saxton MJ. Biophys. J. 1993; 64:1766–1780. [PubMed: 8369407]
21. Ewens, WJ.; Grant, GR. Statistical Methods in Bioinformatics. Springer; 2005.
22. Robert, CP.; Casella, G. Monte Carlo Statistical Methods. Springer; 2004.
23. Grimmett, G.; Stirzaker, D. Probability and Random Processes. 3rd. Oxford University Press; 2001.
24. Gilks, WR.; Richardson, S.; Spiegelhalter, DJ. Markov Chain Monte Carlo in Practice. Chapman & Hall/CRC; 1996.
25. Gelman A, Gilks WR, Roberts GO. Ann. Appl. Probab. 1997; 7:110–120.
26. He J, Guo SM, Bathe M. Anal. Chem. 2012; 84:3871–3879. [PubMed: 22423978]
27. Vallotton P, Olivier S. Microsc. Microanal. 2013; 19:451–460. [PubMed: 23448973]
28. Backer CB, Gutzman JH, Pearson CG, Cheeseman IM. Mol. Biol. Cell. 2012; 23:2122–2130. [PubMed: 22493317]
29. Gascoigne KE, et al. Cell. 2011; 145:410–422. [PubMed: 21529714]
30. Meijering E, Dzyubachyk O, Smal I. Methods Enzymol. 2012; 504:183–200. [PubMed: 22264535]

Figure 1. Particle-trajectory analysis methods applied to neuronal mRNPs

(a) Left, cultured mouse hippocampal neurons in which endogenous β -actin mRNA molecules are labeled with GFP-tagged MS2 (ref. 5), overlaid in color with trajectories of tracked mRNP complexes. Center, enlargement of the boxed region at left. A transported mRNP complex and its associated trajectory (yellow) are shown. Right, kymograph of the selected mRNP complex along the neuronal dendrite. (b) Sliding-window MSD analysis of the mRNP trajectory in a, showing the slope (α) of a log-log plot of MSD versus time lag within sliding windows of ten-step width. The temporal sequence of motion states obtained using a threshold α (dotted line) to classify directed transport is illustrated by the vertical bar on the right. Inferred transport states and diffusive states are pink and blue, respectively, whereas the nonannotated steps of the trajectory are gray. (c) Analysis of the mRNP trajectory in a with a diffusive-only HMM approach¹³. The inferred motion model with two diffusive states (D-D) is shown at the upper left. Blue and green annotations correspond to the states with higher and lower diffusion coefficients, respectively. (d) Analysis of the mRNP trajectory in a with HMM-Bayes. The inferred motion model with one diffusive state and one transport state (D-DV) and the inferred state annotation are shown, with blue and pink annotations corresponding to the diffusive and transport states, respectively.

Figure 2. HMM-Bayes analysis of neuronal mRNPs

(a,b) β -actin mRNP trajectories inferred by HMM-Bayes to undergo directed transport, illustrated as in Figure 1d. The three distinct motion states in b, one diffusive and two transport states (D-DV-DV), are depicted in blue, pink and orange, respectively. (c) Distributions of diffusion coefficients, velocity magnitudes and the mean lifetimes of diffusive and transport motion states as inferred by HMM-Bayes across 22 β -actin mRNP trajectories (from 13 cells) that undergo switching between diffusive motion and directed transport. The diffusion coefficient distribution is over all states, including both diffusive states and transport states, observed across all trajectories. The velocity magnitude distribution is over all transport states observed across all trajectories.

Figure 3. HMM-Bayes analysis of oscillating metaphase kinetochores

(a) Wide-field image of a HeLa cell expressing GFP-CENP-A to label kinetochores and GFP-CSAP to label the spindle poles (SP). Pairs of kinetochore trajectories during metaphase are overlaid in color, from yellow to dark green over time. (b) Enlarged pair of kinetochore trajectories corresponding to the boxed trajectories in a. (c) Cartoon of kinetochore (KT) motion showing projection of positions and displacements onto the direction perpendicular to the metaphase plate (MP). (d,e) Analysis of pooled kinetochore trajectories from the cell in a with HMM-Bayes. (d) The inferred state annotation from the motion model, two transport states (DV-DV) with opposite directions of transport (pink and orange), for one of the kinetochore trajectories from the pair in b is shown. (e) The inferred state annotations for the four pairs of trajectories in a (numbered 1–4 from top to bottom in a) are shown as colored bars as in Figures 1 and 2.