

RESEARCH ARTICLE

The impact of temporal sampling resolution on parameter inference for biological transport models

Jonathan U. Harrison*, Ruth E. Baker

Wolfson Centre for Mathematical Biology, Mathematical Institute, University of Oxford, Oxford, United Kingdom

* harrison@maths.ox.ac.uk



Abstract

Imaging data has become an essential tool to explore key biological questions at various scales, for example the motile behaviour of bacteria or the transport of mRNA, and it has the potential to transform our understanding of important transport mechanisms. Often these imaging studies require us to compare biological species or mutants, and to do this we need to quantitatively characterise their behaviour. Mathematical models offer a quantitative description of a system that enables us to perform this comparison, but to relate mechanistic mathematical models to imaging data, we need to estimate their parameters. In this work we study how collecting data at different temporal resolutions impacts our ability to infer parameters of biological transport models by performing exact inference for simple velocity jump process models in a Bayesian framework. The question of how best to choose the frequency with which data is collected is prominent in a host of studies because the majority of imaging technologies place constraints on the frequency with which images can be taken, and the discrete nature of observations can introduce errors into parameter estimates. In this work, we mitigate such errors by formulating the velocity jump process model within a hidden states framework. This allows us to obtain estimates of the reorientation rate and noise amplitude for noisy observations of a simple velocity jump process. We demonstrate the sensitivity of these estimates to temporal variations in the sampling resolution and extent of measurement noise. We use our methodology to provide experimental guidelines for researchers aiming to characterise motile behaviour that can be described by a velocity jump process. In particular, we consider how experimental constraints resulting in a trade-off between temporal sampling resolution and observation noise may affect parameter estimates. Finally, we demonstrate the robustness of our methodology to model misspecification, and then apply our inference framework to a dataset that was generated with the aim of understanding the localization of RNA-protein complexes.

OPEN ACCESS

Citation: Harrison JU, Baker RE (2018) The impact of temporal sampling resolution on parameter inference for biological transport models. *PLoS Comput Biol* 14(6): e1006235. <https://doi.org/10.1371/journal.pcbi.1006235>

Editor: Qing Nie, University of California Irvine, UNITED STATES

Received: December 18, 2017

Accepted: May 28, 2018

Published: June 25, 2018

Copyright: © 2018 Harrison, Baker. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Code implementing methods mentioned in the article are available at https://github.com/shug3502/pmmc_inference_for_vjps as described in the Supporting Information.

Funding: This work was supported by funding from the Engineering and Physical Sciences Research Council (EPSRC) (grant no. EP/G03706X/1) and the Biotechnology and Biological Sciences Research Council (BBSRC) (grant no. BB/R000816/1). REB is a Royal Society Wolfson Research Merit Award holder, and would like to

thank the Leverhulme Trust for a Leverhulme Research Fellowship. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

We consider how the temporal resolution of imaging studies affects our ability to carry out accurate parameter estimation for a stochastic biological transport model. This model provides a mechanistic description of motile behaviour and is often used to interrogate transport processes, such as the motion of bacteria. Parameter inference is necessary to characterise different types of transport and to make predictions about biological behaviour under different conditions. Typically, observations of the transport process, at the level of individual trajectories, are made at discrete times. This can lead to errors in parameter estimation because we do not have complete trajectory information. We present a framework for Bayesian inference for these models of biological transport processes. Using this framework, we study the effects of collecting data more or less frequently, and with varying measurement noise, on what we can learn about the biological system via parameter estimation.

Introduction

Biological transport processes occur on a wide range of spatial and temporal scales, and a common mechanism for transport involves two phases: fast active transport, and a quasi-stationary reorientation phase. This pattern of movements has been observed at a range of scales from the intracellular transport of cellular components such as mRNA particles moving on a microtubule network [1], to the run-and-tumble motion of bacteria such as *Escherichia coli* [2–4], and the flights of birds between nesting sites [5]. To capture appropriately these two phases of motion, a class of models known as velocity jump process (VJP) models [2, 5–8] (also known as correlated random walks [9–13] or Levy Walks [14–17]) have been developed.

Estimating the parameters of these models can give us mechanistic information relating to the underlying biological process, such as the rate of reorientation. Being able to obtain accurate estimates, with appropriate uncertainty, for these parameters allows us to compare different biological species or mutants, and gain an understanding of the underlying mechanistic behaviour. Importantly, parameterising models and quantifying the uncertainty in parameter estimates, as can be achieved via Bayesian inference, enables us to use models to make quantitative predictions of behaviour in new conditions with quantifiable uncertainty. By performing experiments to test model predictions, we can evaluate the areas in which a given model fails to describe experimental data, and so iteratively refine our understanding of a given system or phenomenon.

In this work, we consider the effects that experimental design can have on the information we can obtain from a data set, in terms of using that data to estimate parameters of a mechanistic model. In particular, for time series data describing a biological transport process, we vary the time between successive measurements. We demonstrate a framework for estimating the parameters of a VJP model for data of this form in the presence of noise, and examine how the posterior estimates of the model parameters change for more coarsely sampled and noisier datasets. Our framework formulates the VJP model as a process with hidden states, as in a hidden Markov model (HMM), which allows us to use particle Markov Chain Monte Carlo (pMCMC) methods to perform exact Bayesian inference. We use our framework to suggest sensible experimental design choices in the context of microscopy studies, where there may be a trade-off between how frequently it is possible to image, and the noise resulting from more or less frequent observations of the process. We present a comparison between the pMCMC framework described in this work and approximate Bayesian computation (ABC) for

parameter inference in this context. We demonstrate robustness to model misspecification in a situation where we attempt inference on synthetic data generated with a different model to the assumed model. Finally, we apply our method to an imaging dataset of tracks of RNA-protein complexes allowing us to estimate motility parameters for a mechanistic model.

Velocity jump process models

VJP models have been developed to describe biological transport processes where there is persistent or biased random motion [2, 5, 18]. Correlated, persistent motion is observed experimentally in a variety of contexts [19, 20]. VJP models are most appropriate for motion consisting of multiple phases, such as a fast directed phase and a stationary or reorientation phase, and these models describe how an object moves in one direction before reorienting and moving in a new direction. This type of “run and reorientate” motion is exhibited, for example, by *Escherichia coli* [3, 4], fibroblasts [20], and RNA-protein complexes [1].

We present here a mathematical description of a VJP model. Suppose we have a running time distribution, with probability density function (pdf) f_r , and a waiting time distribution, with pdf f_w . Random variables drawn from these distributions dictate the lengths of time spent in the fast active transport, or running, phase of the VJP and the reorientation phase, respectively. After the reorientation phase, a velocity for the new run is chosen according to a transition kernel, f_v . This transition kernel could incorporate a distribution of speeds as well as directions, or could rely on a fixed speed and specify only the directionality of the run. In the fixed speed case,

$$f_v(\mathbf{v}) = \frac{\delta_0(|\mathbf{v}| - c)}{|\mathbf{v}|} f_\phi(\phi),$$

where f_ϕ is a reorientation kernel describing the angle change, and c is the constant running speed. For biological processes with a distinct separation of timescales between the running and reorientation phases, it is often possible to assume that reorientations between successive runs occur instantaneously and therefore neglect the reorientation phase in a model of the process [2, 5]. Repeated simulation from this model can be used to generate trajectories of individual particles (see Fig 1, for example).

In the simplest case, where we assume that the running time distribution is memoryless, that is, exponentially distributed, with $f_r(t) = \lambda \exp(-\lambda t)$, the long time behaviour of the mean squared displacement scales linearly with time, t [2]. Further moments of the motion of individuals displaying such VJP behaviour have been characterised by making certain closure assumptions [7]. For the case of a more general running time distribution, with finite mean and variance, in the large time limit the probability of the particle being at position x at time t follows a diffusion equation [5].

In practice, VJP models are often parameterised by obtaining measurements of the effective diffusion coefficient or the mean squared displacement, and using these data to estimate the parameters of a specific running distribution [2, 5]. Rosser et al. [4] parameterised a HMM via maximum likelihood estimation, whilst Nicosia et al. [21] fitted a hidden state random walk model to animal movement data using an expectation-maximization algorithm. These frequentist approaches can provide useful point estimates of the VJP parameters, and quantify the error in these estimates. However, Bayesian approaches can propagate uncertainty from both process noise and measurement noise, which can be crucial when dealing with noisy biological data (see Fig 1). In addition, they provide the added benefit that we can interpret the results of a Bayesian analysis as probabilistic statements about the model parameters. This quantification of uncertainty enables the generation of predictions of further biological

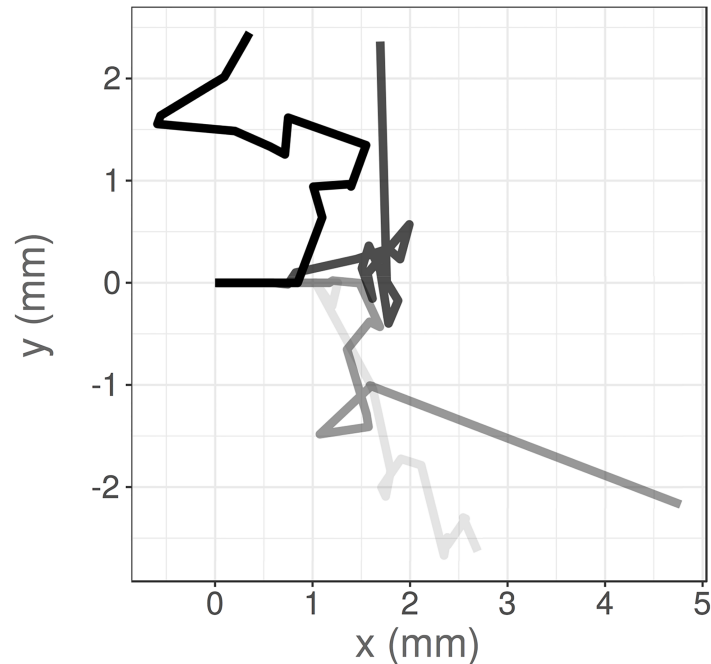


Fig 1. Four example VJP trajectories, simulated with a uniform reorientation kernel, running speed $c = 50\mu\text{ms}^{-1}$, exponential running time distribution, $f_{\tau}(t) = \lambda e^{-\lambda t}$, with reorientation frequency $\lambda = 0.2\text{ s}^{-1}$ and no waiting time between runs, for a duration of $T_{\text{final}} = 64\text{s}$. These trajectories start from the origin, and are orientated initially parallel to the positive x -axis.

<https://doi.org/10.1371/journal.pcbi.1006235.g001>

behaviour using the model. In addition, we can consider the effects of noisy data measurements upon the accuracy of the inferred parameters distributions, something which has previously been difficult to deal with in practice, or has been neglected.

For simplicity, in this work we will assume that there is a separation of timescales between the lengths of the running and reorientation phases, such that reorientations can be considered instantaneous. In addition, we assume that the running time distribution, f_{τ} , is exponentially distributed, that particles run at a fixed speed, c , and the reorientation kernel is a uniform distribution on $[-\pi, \pi)$. That is $f_{\tau}(t) = \lambda \exp(-\lambda t)$ and $f_{\phi}(\phi) = 1/2\pi \mathbb{1}_{[-\pi, \pi)}(\phi)$. We follow the trajectory of a single, motile individual, and take, as experimental measurements, the change in angle between successive observed positions (calculated relative to the previous observed position, see Fig 2), subject to measurement noise drawn from a wrapped Normal distribution, $N(0, \sigma^2)$, where σ is the magnitude of the noise. Our Bayesian inference approach will target estimation of the reorientation rate, λ , and the magnitude of the noise, σ .

Inference for velocity jump process models via particle Markov Chain Monte Carlo

Inference for partially observed Markov processes can be performed using particle Markov Chain Monte Carlo (pMCMC), as developed by Andrieu and Roberts [22] and Andrieu et al. [23]. pMCMC provides a Bayesian framework for parameter estimation by allowing samples to be drawn from the posterior distribution of the model parameters, given observed data, without needing to evaluate the likelihood function directly. For partially observed Markov process models, the model structure makes directly evaluating the likelihood difficult or

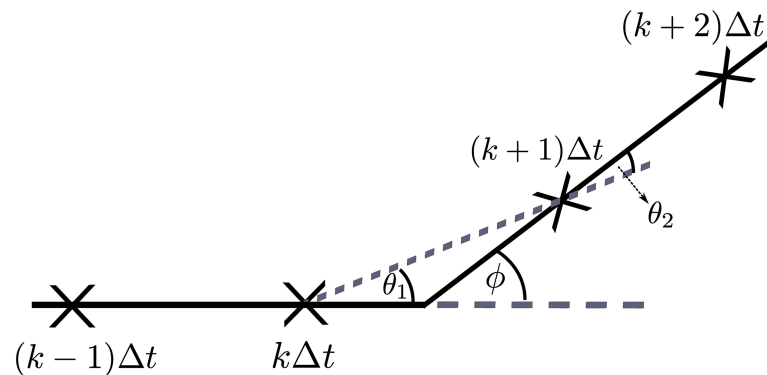


Fig 2. The motile individual is observed at discrete times $t = k\Delta t$ for $k \in \{0, 1, \dots, T\}$ (marked by crosses). In each time interval between observations, we measure the angle change between the observed direction of travel between successive observations. We assign a hidden state to each time interval according to whether a reorientation event occurred in that time interval or not. In this example, which excludes measurement noise, there is a reorientation, through angle ϕ , between observations at times $t = k\Delta t$ and $t = (k + 1)\Delta t$; however, the measured angle change between these observations is θ_1 . Conversely, there is no reorientation event between observations at $t = (k + 1)\Delta t$ and $t = (k + 2)\Delta t$ and an observed angle change of θ_2 .

<https://doi.org/10.1371/journal.pcbi.1006235.g002>

expensive. (We note that VJP models can be viewed in this form by introducing hidden states, as explained in Section ‘Methods’.) Instead of evaluating the likelihood directly, we can use (unbiased) estimates of the likelihood within an MCMC algorithm. Estimating the likelihood of the observed data given certain parameters can be achieved with a particle filter (also known as a sequential Monte Carlo scheme) [24] for a fixed, finite, number of particles. The results of Andrieu et al. [23] demonstrate that, even when a finite number of particles are used in the filter to estimate the likelihood, the MCMC algorithm will still target the correct posterior distribution. These methods have been applied in the context of modelling epidemics [25] and biochemical reaction networks [26, 27]. However, pMCMC methods for parameter inference have not previously been applied to spatial agent-based models, such as the VJP model considered here. The details of the pMCMC algorithm used in this work are given in Section ‘Methods’.

Methods

We will demonstrate how to exploit pMCMC methods to obtain posterior parameter estimates for the VJP outlined in the Introduction by formulating the model within an appropriate framework that incorporates hidden states. This formulation additionally allows us to incorporate a model for measurement noise in our observations, as well as explicitly accounting for the temporal discretisation of the data. The hidden states (also known as latent variables) in our model will describe whether or not a reorientation event occurred between observations of the VJP. This is not a variable that we can observe directly, since we only measure the observed angle change. For example, in Fig 2 there was a reorientation event between observations at $t = k\Delta t$ and $t = (k + 1)\Delta t$, and an observed angle change of θ_1 . On the other hand, there was no reorientation event between observations at $t = (k + 1)\Delta t$ and $t = (k + 2)\Delta t$ but still a non-zero observed angle change of θ_2 . Note that the true reorientation angle for the reorientation event that takes place between observations at $t = k\Delta t$ and $t = (k + 1)\Delta t$ is ϕ .

We assume that we observe the system at times $\{k\Delta t: k \in \{0, 1, \dots, T\}\}$, up until a final time $T_{\text{final}} = T\Delta t$, by measuring the position of the individual of interest and recording the observed

angle change. We define the hidden variable, X_k , as follows:

$$X_k = \begin{cases} 1, & \text{if a reorientation event occurred during } [k\Delta t, (k+1)\Delta t), \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The observed state is the observed angle change, θ_k , obtained as the difference between the observed direction of travel during $[(k-1)\Delta t, k\Delta t)$ and $[k\Delta t, (k+1)\Delta t)$, $k \geq 1$, as illustrated in Fig 2. For example, if we directly observe as data the positions $\{(x_k, y_k): k \in \{0, 1, \dots, T\}\}$, then we can calculate the observed angle change as

$$\left\{ \theta_k = \arctan\left(\frac{y_{k+1} - y_k}{x_{k+1} - x_k}\right) : k \in \{0, 1, \dots, T-1\} \right\}.$$

We illustrate the sequence of hidden and observed states in Fig 3, with dependence between these states given by the transition and emission probabilities, denoted β_i and p_{ij} , respectively. The hidden variable, X_k , evolves according to the VJP model. Here, since we have an exponential distribution for the running time, $P(\text{reorientation event in } [(k-1)\Delta t, k\Delta t)) = 1 - \exp(-\lambda\Delta t)$.

We assume that reorientation events are rare relative to the sampling rate such that $\lambda\Delta t \ll 1$. This assumption allows us to neglect multiple reorientations in a single time interval, enabling us to greatly simplify the problem, so that it is tractable via the binary hidden variables X_k . We highlight that the model structure as shown in Fig 3 is an approximation. In reality, observed data will not be a function of only the current and previous hidden states, X_k and X_{k-1} , but may depend on other previous states also. The assumption that $\lambda\Delta t \ll 1$ enables us to neglect dependence on earlier states.

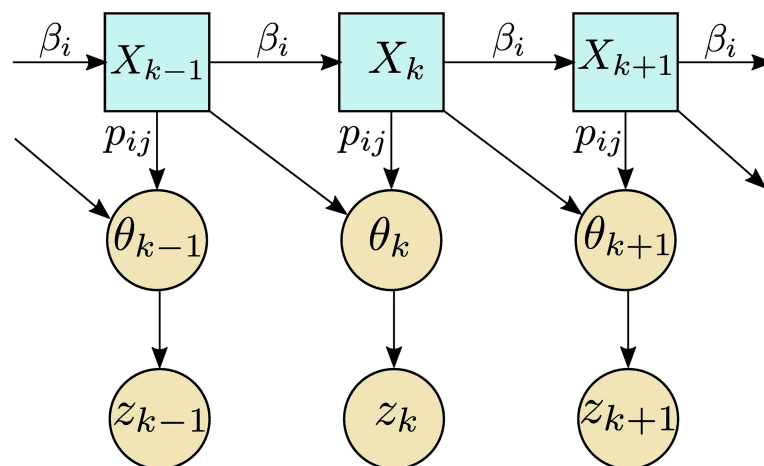


Fig 3. Hidden and observed states in a partially observed Markov process model. We observe an angle change, θ_k , which is dependent on a hidden state X_k defined in Eq (1). This dependency is shown by the arrows between states. Here, we require dependencies on the previous hidden state, X_{k-1} , since observation times will not coincide with reorientation events in general. Additionally, we introduce an extra layer of state dependencies to capture the measurement noise in this noisy biological system. We suppose that θ_k is the true observed angle change and our noisy observed version of this is z_k , with dependencies shown by the arrows.

<https://doi.org/10.1371/journal.pcbi.1006235.g003>

Model without measurement noise

We initially make progress by simplifying the problem and its exposition via the assumption of zero measurement error. To relate the unobserved hidden state to the observed angle change, which we can measure, we derive probability distributions for the angle change given the hidden state. Note that there is additional dependence not only on the current hidden state, but also on the previous hidden states, as shown in Fig 3. We can obtain expressions for the emission probabilities by considering the path that is taken under different sequences of hidden states (Fig 4); we will outline simplifying assumptions as they are made.

Suppose we observe the system at times $(k - 1)\Delta t$, $k\Delta t$, and $(k + 1)\Delta t$ for some k , as shown in Fig 4. The sequence of hidden states i, j , corresponds to whether there were reorientation events in the time intervals $[(k - 1)\Delta t, k\Delta t)$ and $[k\Delta t, (k + 1)\Delta t)$, respectively. Let $p_{ij}(\theta)$ be the pdf of observing a reorientation angle of θ given the sequence of hidden states i, j . We assume angle change θ was observed over the time interval $[k\Delta t, (k + 1)\Delta t)$ corresponding to the hidden state j . In the case where no reorientation occurs in either of the time intervals (corresponding to the situation in Fig 4a)), then, assuming no noise, we would observe zero angle

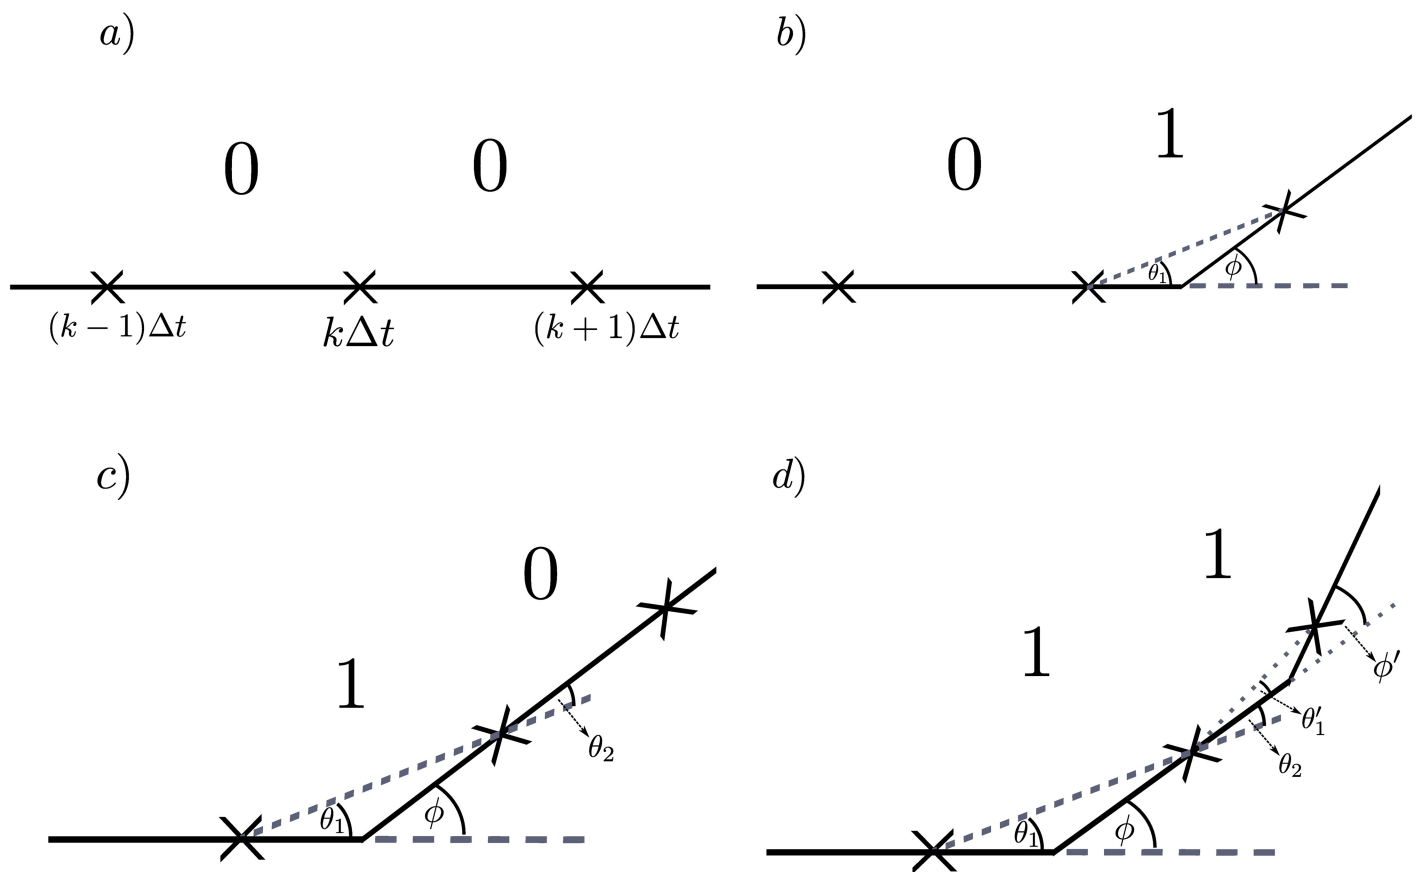


Fig 4. Example paths for each pattern of hidden states. The particle of interest moves from left to right, and is observed at times $(k - 1)\Delta t$, $k\Delta t$, and $(k + 1)\Delta t$. In a), there are no reorientation events, so the particle continues along a straight trajectory. In b), there is a single reorientation in the time interval $[k\Delta t, (k + 1)\Delta t)$, where the particle turns through an angle ϕ , but the observed angle change is θ_1 due to the discrete nature of the observations. In c), there is a single reorientation in the time interval $[(k - 1)\Delta t, k\Delta t)$. The particle turns through an angle ϕ and continues on its new trajectory. We observe an angle change θ_1 for the time interval $[(k - 1)\Delta t, k\Delta t)$, and observe an angle change θ_2 for the next time interval $[k\Delta t, (k + 1)\Delta t)$ even though there was no reorientation during this time interval. Similarly, in d), there was a reorientation of true angle change ϕ during $[(k - 1)\Delta t, k\Delta t)$ followed by another reorientation of true angle change ϕ' in the time interval $[k\Delta t, (k + 1)\Delta t)$. In this case, we observe angle change θ_1 for the time interval $[(k - 1)\Delta t, k\Delta t)$, and observe an angle change of $\theta_2 + \theta'_1$ for the time interval $[k\Delta t, (k + 1)\Delta t)$.

<https://doi.org/10.1371/journal.pcbi.1006235.g004>

change. That is, we have $p_{00}(\theta) = \delta_0(\theta)$. If a reorientation occurs in the time interval $[k\Delta t, (k + 1)\Delta t)$, but not in the preceding time interval, $[(k - 1)\Delta t, k\Delta t)$, (as shown in Fig 4b), then the observed reorientation angle is θ_1 , as labelled in Fig 4b). This gives $p_{01}(\theta) := p_{\Theta_1}(\theta)$. The marginal distribution of Θ_1 is derived in the Section ‘Derivation of emission probabilities’, and depends on the running time distribution and the reorientation kernel.

If, immediately after a reorientation, we have no reorientation in the following time interval, then we may still observe a nonzero angle change since our discrete observation times do not in general coincide with the reorientation events. In this case, we have a reorientation event during $[(k - 1)\Delta t, k\Delta t)$, and no reorientation during $[k\Delta t, (k + 1)\Delta t)$. Such a situation with a pattern of hidden states 1,0 is shown in Fig 4c), and the observed angle change during the time interval $[k\Delta t, (k + 1)\Delta t)$ corresponds to θ_2 in the diagram. We note that, by geometric arguments, we have

$$\theta_1 + \theta_2 = \phi, \tag{2}$$

where ϕ is the true angle change. Therefore, given a pattern of hidden states 1, 0 the angles θ_1 and θ_2 are not independent, and we have $p_{10}(\theta_1, \theta_2) = p_{\Theta_2|\Theta_1=\theta_1}(\theta_2)$. Note that throughout we will assume that we can observe the previous angle change directly, which is equivalent to assuming that the pattern of hidden states is in fact 0, 1, 0.

The remaining cases involve reorientation events in successive time intervals. Suppose we have the case where we have two successive reorientation events, giving a pattern of hidden states 1, 1, which is shown graphically in Fig 4d). This case is similar to the case with hidden states 1, 0 (see Fig 4c)), in that we observe a contribution to the reorientation angle, θ_2 , from the correction for the previous reorientation event, and also a contribution from the new reorientation event, θ'_1 . As can be seen in Fig 4d), these contributions sum to give an observed angle change for the time interval $[k\Delta t, (k + 1)\Delta t)$ of $\theta = \theta_2 + \theta'_1$. The probability density for sums of random variables can be expressed as a convolution [28]. Hence, $p_{11}(\theta_1, \theta) = (p_{\Theta_1} * p_{\Theta_2|\Theta_1=\theta_1})(\theta)$, where $*$ is the convolution operator. Any further cases involve multiple reorientation events within a single time interval which occurs rarely, with probability on the order of $\mathcal{O}((\lambda\Delta t)^2)$. We can safely neglect these provided $\lambda\Delta t \ll 1$.

To summarise therefore, the emission probabilities (that is the probability of observing a certain angle change given the sequence of hidden states) in the case without noise can be given as follows:

$$\begin{aligned} p_{00}(\theta) &= \delta_0(\theta); \\ p_{01}(\theta) &= p_{\Theta_1}(\theta); \\ p_{10}(\theta) &= p_{\Theta_2|\Theta_1}(\theta); \\ p_{11}(\theta) &= (p_{\Theta_2|\Theta_1} * p_{\Theta'})(\theta). \end{aligned}$$

Model with measurement noise

To account for noise, we introduce another layer of states in our diagram of state dependencies, as shown in Fig 3. Let z_k be the noisy observed angle change at time $k\Delta t$, which is a noisy observation of θ_k , such that for a noise model K we have $z_k \sim K(\cdot, \theta_k)$. Alternatively, we can write the noisy observed angle as

$$z_k = \theta_k + \epsilon_k, \tag{3}$$

where $\epsilon_k \sim K(\cdot, 0)$. Under the assumption of a wrapped normal noise model, $\epsilon_k \sim N(0, \sigma^2)$,

where σ is the magnitude of the measurement noise. Since Eq (3) represents z_k as a sum of random variables, it becomes clear that we can obtain the noisy emission probabilities, q_{ij} , corresponding to a pattern of hidden states i, j via the following convolution:

$$q_{ij}(\theta) = (p_{ij} * K)(\theta) = \int_{\Theta} p_{ij}(\theta - x)K(x, 0) dx.$$

To extend our previous results to the noisy case, we therefore need to compute these convolutions, a task that must be carried out numerically.

Derivation of emission probabilities

In previous work [29], marginal distributions for the observed angle change, Θ_1 , were obtained and we summarise the arguments here. Fig 5 shows the true angle change, Φ , and the observed angle change, Θ_1 , based on discrete time observations of the VJP for the case of hidden states 0, 1.

By changing co-ordinates from the displacement, L , (which is given by the running time distribution f_τ) and true angle change, Φ , to Cartesian co-ordinates X and Y , and then to polar co-ordinates, R and Θ_1 , we can obtain a joint distribution for R and Θ_1 of

$$f_{R,\Theta_1}(r, \theta_1) = \begin{cases} \frac{r}{c\Delta t(c\Delta t - r \cos \theta_1)} f_\Phi \left(\arccos \left[\frac{r^2 \sin^2 \theta_1 - (c\Delta t - r \cos \theta_1)^2}{r^2 \sin^2 \theta_1 + (c\Delta t - r \cos \theta_1)^2} \right] \right), & \text{if } (r, \theta_1) \in C, \\ 0, & \text{otherwise,} \end{cases}$$

where $C = \{r \leq c\Delta t, \theta \in [-\pi, \pi)\}$ is the set of permissible values for r and θ , Δt is the discretisation in time and c is the running speed. Integrating $f_{R,\Theta_1}(r, \theta_1)$ over r allows us to obtain the marginal distribution for θ_1 via

$$f_{\Theta_1}(\theta_1) = \int_0^\infty f_{R,\Theta_1}(r, \theta_1) dr.$$

In the case where we assume a uniform reorientation kernel, $f_\Phi(\phi) = 1/2\pi \mathbb{1}_{[-\pi,\pi)}(\phi)$, then this

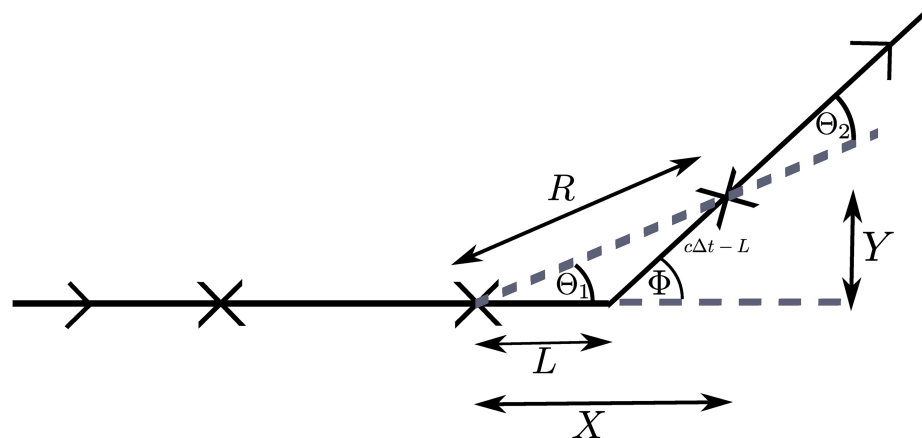


Fig 5. True and observed angle changes (Φ and Θ_1 , respectively) based on discrete time observations of a VJP without measurement error. Cartesian co-ordinates X and Y are shown in addition to polar co-ordinates R and Θ_1 . The particle is moving along the trajectory from left to right as shown by the arrows.

<https://doi.org/10.1371/journal.pcbi.1006235.g005>

integral becomes

$$\begin{aligned}
 f_{\Theta_1}(\theta_1) &= \int_0^{c\Delta t} \frac{r}{2\pi c\Delta t(c\Delta t - r \cos \theta_1)} dr \\
 &= \frac{-1}{2\pi c\Delta t \cos(\theta_1)} \int_0^{c\Delta t} \left(1 + \frac{\frac{c\Delta t}{\cos \theta_1}}{r - \frac{c\Delta t}{\cos \theta_1}} \right) dr \\
 &= \frac{-1}{2\pi c\Delta t \cos(\theta_1)} \left[r + \frac{c\Delta t}{\cos \theta_1} \log \left| r - \frac{c\Delta t}{\cos \theta_1} \right| \right]_0^{c\Delta t} \\
 &= \frac{-1}{2\pi \cos^2 \theta_1} (\cos \theta_1 + \log(1 - \cos \theta_1)).
 \end{aligned}$$

We note that the marginal distribution for θ_1 does not depend on the speed, c , or the time discretisation, Δt . This is intuitive because f_{Θ_1} is the pdf of a certain observed angle change, θ_1 , given there was a reorientation in that interval, irrespective of the length of that interval.

Derivation of the joint distribution of Θ_1 and Θ_2 . By considering the displacement in the X and Y directions during a time step, we have

$$\begin{aligned}
 r \cos(\theta_1) &= l + (c\Delta t - l) \cos(\phi), \\
 r \sin(\theta_1) &= (c\Delta t - l) \sin(\phi).
 \end{aligned}$$

Dividing these expressions, we can relate θ_1 to l and ϕ by

$$\tan(\theta_1) = \frac{(c\Delta t - l) \sin(\phi)}{l + (c\Delta t - l) \cos(\phi)}, \tag{4}$$

and rearranging for l , we have

$$l = \frac{c\Delta t(\tan \phi - \tan \theta_1)}{\tan \phi + \tan \theta_1(\sec \phi - 1)}.$$

Differentiating with respect to θ_1 , we obtain

$$\frac{\partial l}{\partial \theta_1} = \frac{-c\Delta t \sec^2 \theta_1}{\tan \phi + \tan \theta_1(\sec \phi - 1)} - \frac{c\Delta t(\tan \phi - \tan \theta_1) \sec^2 \theta_1 (\sec \phi - 1)}{(\tan \phi + \tan \theta_1(\sec \phi - 1))^2},$$

which can be simplified to give

$$\frac{\partial l}{\partial \theta_1} = \frac{-c\Delta t \sin \phi}{(\sin \phi \cos \theta_1 + \sin \theta_1(1 - \cos \phi))^2}.$$

To transform from coordinates (L, Φ) to coordinates (Θ_1, Φ) , we can use the Jacobian $J_{L, \Phi}$, where

$$\det J_{L, \Phi} = \frac{\partial l}{\partial \theta_1} = \frac{-c\Delta t \sin \phi}{(\sin \phi \cos \theta_1 + \sin \theta_1(1 - \cos \phi))^2}. \tag{5}$$

Therefore the joint distribution of θ_1 and ϕ is

$$f_{\Theta_1, \Phi}(\theta_1, \phi) = \begin{cases} \frac{|c\Delta t \sin \phi|}{(\sin \phi \cos \theta_1 + \sin \theta_1(1 - \cos \phi))^2} f_{\phi}(\phi) \cdot \frac{1}{c\Delta t}, & \text{if } \theta_1, \phi \in C, \\ 0, & \text{otherwise,} \end{cases}$$

where $C = \{(\theta_1, \phi): \phi \in [-\pi, \pi], \theta \in (\min(0, \phi), \max(0, \phi))\}$. Under the assumption that ϕ is

uniform on $[-\pi, \pi)$, as described in the Section ‘Velocity jump process models’, such that $f_\phi(\phi) = 1/2\pi\mathbb{1}_{[-\pi,\pi)}(\phi)$, we have

$$f_{\Theta_1,\Phi}(\theta_1, \phi) = \begin{cases} \frac{|\sin \phi|}{2\pi(\sin \phi \cos \theta_1 + \sin \theta_1(1 - \cos \phi))^2}, & \text{if } \theta_1, \phi \in C, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Changing variables again to (θ_1, θ_2) , via $\phi = \theta_1 + \theta_2$, we have

$$f_{\Theta_1,\Theta_2}(\theta_1, \theta_2) = \begin{cases} \frac{|\sin(\theta_1 + \theta_2)|}{2\pi(\sin(\theta_1 + \theta_2) \cos \theta_1 + \sin \theta_1(1 - \cos(\theta_1 + \theta_2)))^2}, & \text{if } \theta_1, \theta_2 \in C, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where $C = \{(\theta_1, \theta_2): \theta_1 + \theta_2 \in [-\pi, \pi)\}$. From this joint distribution, we can then find the conditional distribution of θ_2 given θ_1 , which is required to give $p_{10}(\theta) = f_{\Theta_1,\Theta_2}(\theta_1, \theta)/f_{\Theta_1}(\theta_1)$.

Particle MCMC algorithm

The pMCMC algorithm used in this work is given in Algorithm 1 for an observed dataset $\mathbf{y} = \{y_k \mid k = 1, 2, \dots, T\}$. We use a Metropolis-Hastings MCMC algorithm [30], proposing new parameters using a proposal distribution $q(\cdot|\theta)$. In step 6 of Algorithm 1, we need to evaluate the likelihood, $p(\mathbf{y}|\theta)$, in calculating the acceptance probability for the proposed move. We replace the likelihood with an unbiased estimate of the likelihood, $\hat{p}(\mathbf{y}|\theta)$, obtained from a particle filter.

Algorithm 1 Particle MCMC

- 1: Initialise parameters, θ_0 .
- 2: Run a particle filter (see Algorithm 2) to compute an estimate of the marginalised likelihood $\hat{p}(\mathbf{y}|\theta_0)$, where \mathbf{y} is the observed data.
- 3: **for** $j = 1:N$ **do**
- 4: Draw parameters θ^* from a proposal distribution $q(\cdot|\theta_{j-1})$.
- 5: Run a particle filter to compute an estimate of the marginalised likelihood $\hat{p}(\mathbf{y}|\theta^*)$.
- 6: Accept the proposed move with probability α where

$$\alpha = \max\left(1, \frac{\pi(\theta^*)q(\theta_{j-1}|\theta^*)\hat{p}(\mathbf{y}|\theta^*)}{\pi(\theta_{j-1})q(\theta^*|\theta_{j-1})\hat{p}(\mathbf{y}|\theta_{j-1})}\right).$$

If the move is accepted, set $\theta_j = \theta^*$, otherwise set $\theta_j = \theta_{j-1}$.

7: **end for**

Algorithm 2 Bootstrap particle filter

- 1: Sample a collection of particles $\{x_1^1, \dots, x_1^M\}$ from an initial density $p(x_1)$.
- 2: **for** $i = 1:M$ **do**
- 3: Compute the weights for each particle, i , via $w_1^i = p(y_1|x_1^i, \theta)$.
- 4: Find the normalised weights

$$\tilde{w}_1^i = \frac{w_1^i}{\sum_{j=1}^M w_1^j}.$$

5: **end for**

```

6: Resample  $N$  times with replacement from the collection of particles
    $\{x_1^1, \dots, x_1^M\}$  with probabilities given by the normalised weights
    $\{w_1^1, \dots, w_1^M\}$ .
7: for  $t = 1:(T-1)$  do
8:   for  $i = 1:M$  do
9:     Evolve the current collection of particles according to the
       forward model, by drawing  $x_{t+1}^i \sim p(x_{t+1}^i | x_t^i, \theta)$ .
10:    Compute the weights for each particle,  $i$ , via  $w_{t+1}^i = p(y_{t+1} | x_{t+1}^i, \theta)$ .
11:    Find the normalised weights

```

$$\tilde{w}_{t+1}^i = \frac{w_{t+1}^i}{\sum_{j=1}^M w_{t+1}^j}.$$

```

12:   end for
13: Resample  $M$  times with replacement from the collection of
    particles  $\{x_{t+1}^1, \dots, x_{t+1}^M\}$  with probabilities given by the normalised
    weights  $\{\tilde{w}_{t+1}^1, \dots, \tilde{w}_{t+1}^M\}$ .
14: end for
15: Obtain an estimate of the marginal likelihood using the unnormalised
    weights, via

```

$$\hat{p}(\mathbf{y}|\theta) = \prod_{t=1}^T \frac{1}{M} \sum_{i=1}^M w_t^i. \quad (16)$$

```

16: return  $\hat{p}(\mathbf{y}|\theta)$ 

```

Details of the bootstrap particle filter algorithm [24] used within the pMCMC algorithm are given in Algorithm 2. A particle filter represents the state of the system via a population of weighted particles [31]. We obtain an estimate of the likelihood by successively updating the hidden state of the system (represented via the particles) and comparing this hidden state with the observed data at each observed time point, to give new weights for the particles according to how well they match the observed data. To prevent a degenerate situation where the state of the system is represented solely by a single particle, it is necessary to resample from the population of particles according to their weights.

Implementation of Markov Chain Monte Carlo methods

We use a Metropolis-Hastings algorithm to run the MCMC algorithm with a bootstrap particle filter [24] using 400 particles to provide an estimate of the likelihood. As a proposal distribution, we use the kernel $K(\cdot, \theta) \sim N(\theta, \Sigma)$, where

$$\Sigma = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.05 \end{bmatrix}.$$

We run the Markov chain for $N = 50,000$ steps with thinning of $m = 2$ which gives a minimum effective sample size of $n_{\text{eff}} = 743$ for the slowest chain to converge, corresponding to the smallest value of Δt . We demonstrate the convergence of our Markov chains in S4 Fig. We choose the number of particles in the filter by considering the variance in estimating the log likelihood using the particle filter at the true values of the parameters used to generate the synthetic data, and balancing this with the time needed to run the particle filter to obtain a single estimate. The variance and computational cost are shown in S5 Fig. To further tune the number of particles for optimal efficiency of the pMCMC algorithm, the recommendations in Sherlock et al.

[32] and Pitt et al. [33] could be employed. The prior is uniform on the log of the parameters over the intervals $[-1.70, 1.30]$ and $[-5, 1]$ for λ and σ , respectively.

Approximate Bayesian computation

An alternative method commonly used for parameter estimation for models with intractable likelihoods is approximate Bayesian computation, or ABC [34, 35]. Suppose we can simulate data from a generative model, $x \sim g(x|\theta)$. An example would be simulating a path from our VJP model. To generate samples from the posterior distribution, we repeatedly generate parameters from the prior, $\theta \sim \pi(\theta)$, and use these parameters in our model to simulate synthetic data, $x \sim g(x|\theta)$. We compare the simulated data with the true observed data and, if it is similar enough, we accept the sampled parameter as a sample from the posterior. In cases where the data are discrete, we can consider whether simulated data, x , are equal to observed data, y , and accept parameters correspondingly, which gives exact samples from the posterior. Often, though, our models provide continuous data and the acceptance rate from an exact comparison is prohibitively small. Instead we can take a distance function and consider when the distance between x and y is within a chosen tolerance ϵ , that is $d(x, y) < \epsilon$. This introduces an approximation which becomes exact only in the limit $\epsilon \rightarrow 0$.

In practice, datasets can be high dimensional, resulting in very low acceptance rates for simulated data. By using summary statistics of data instead of the full datasets, we can reduce the dimensionality and obtain higher acceptance rates. Using (insufficient) summary statistics introduces another approximation into the method, meaning we compare $d(s(x), s(y)) < \epsilon$, where $s(x)$ gives the summary statistics for dataset x . We present the ABC method in Algorithm 3, and will show a comparison between application of pMCMC and ABC for parameter estimation of a VJP model based on time series data.

Algorithm 3 ABC rejection sampling

- 1: Sample parameter value θ from the prior $\pi(\theta)$.
- 2: Generate synthetic dataset from the model $x \sim g(x|\theta)$.
- 3: Accept sample θ if $d(s(x), s(y)) < \epsilon$, for a distance function, $d(\cdot, \cdot)$, summary statistics, $s(\cdot)$, and tolerance ϵ .

In practice, the choice of summary statistics can have a strong effect on the efficiency of the ABC algorithm. Here, we perform ABC with three different choices of summary statistic. First, we consider using the full data, a vector of the noisy observed angle changes, $z_i, i = 1, \dots, T$. Second, we use a simple threshold-based summary statistic which produces a count of the number of observed angle changes with magnitude above a certain threshold, h . That is $s(\mathbf{z}) = \sum_{i=1}^T \mathbb{1}_{|z_i| > h}$. This provides an intuitive one-dimensional summary statistic. Finally, we use a transition matrix as a summary statistic of the data, an approach described by Jones et al. [13]. A transition matrix allows us to summarise time series data via a two-dimensional binning of the observed angle changes, based on the current observed angle change and the previous observed angle change. We bin the observed angle change time series data into an n by n matrix, where $n = 5$. This summary statistic provides a much lower dimensional summary of the data for small values of Δt , although for large values of Δt we may end up with higher dimensional (sparse) data compared to the full time series data, as in this case there will be fewer than n^2 observations. The distance function, $d(s(x), s(y))$, used also plays a role in the parameter estimation possible via ABC, and has been investigated in other work [13, 36–38].

In using Algorithm 3, we generate $N = 1,000,000$ synthetic datasets based on parameters sampled from the prior, calculate the Euclidean distance between these synthetic datasets and our observed data, and select the parameters corresponding to the 0.1% of the datasets closest to the observed data. As for the pMCMC approach, we take a duration of time series $T_{\text{final}} =$

64s, and a prior uniform on the log of the parameters on the intervals $[-1.70, 1.30]$ and $[-5, 1]$ for λ and σ , respectively.

Results

We are able to investigate the effects of experimental design such as the discretisation, Δt , of a time series on the estimated biophysical parameters, using the framework for inferring the parameters of a VJP described in the Section ‘Methods’. We demonstrate the effects of restrictions imposed by experimental constraints on a trade-off between discretisation in time versus measurement noise. Our results can be used to provide guidance for experimental design choices. As we vary experimental design hyperparameters, such as Δt , we will produce a separate posterior distribution for each of the model parameters, λ and σ , for each value of the hyperparameter. This approach allows us to illustrate directly the effects of the experimental design hyperparameter on the inferred posterior distributions.

Increasing Δt results in an abrupt breakdown in the posterior

We first assume that the magnitude of the noise on the observed angle, σ , is fixed. We vary the sampling frequency used to collect the data (corresponding to using different values of Δt). We generate observed (*in silico*) data by simulating one single trajectory directly from the VJP model. We discretise this trajectory to generate observed angle changes with a temporal resolution of Δt , and add independent Gaussian noise with zero mean and variance σ^2 to these observations, to represent measurement noise. Using datasets generated from the same simulated path discretised at different temporal resolutions, Δt , as shown in Fig 6, we infer the posterior distribution for the reorientation rate, λ , of the VJP model and the magnitude of the measurement noise, σ , via pMCMC.

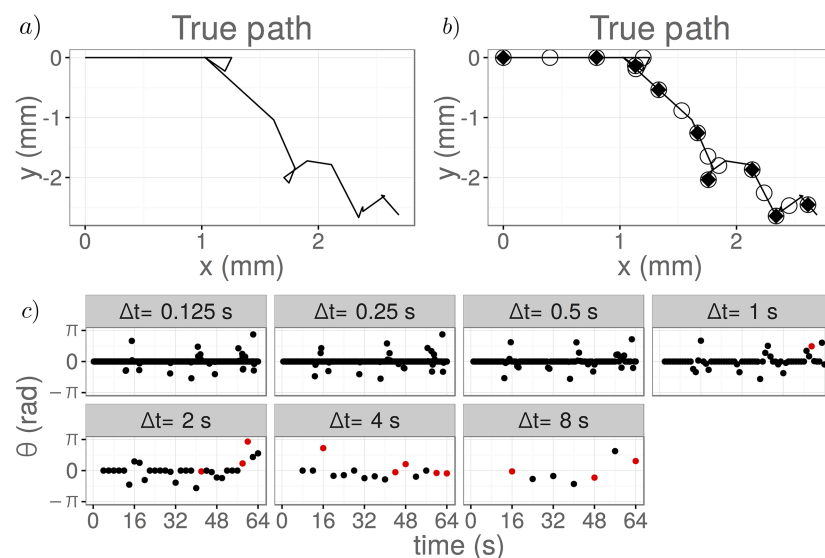


Fig 6. The data used in the pMCMC inference discretised at differing resolutions. The true trajectory is shown without discretisation in a). The same path is shown in b) with circle markers to show the observed positions with $\Delta t = 4$ s and filled triangular markers to show observed positions with $\Delta t = 8$ s. The corresponding observed angle changes are shown in c) for different values of Δt . Observations corresponding to multiple reorientations in a single time interval are highlighted in red. Parameters used in generating these data were a reorientation rate of $\lambda = 0.2 \text{ s}^{-1}$, a run speed of $c = 50 \mu\text{ms}^{-1}$ and a total duration of observation of $T_{\text{final}} = 64$ s. A circular uniform distribution was used for the reorientation kernel. Observations are shown without measurement noise.

<https://doi.org/10.1371/journal.pcbi.1006235.g006>

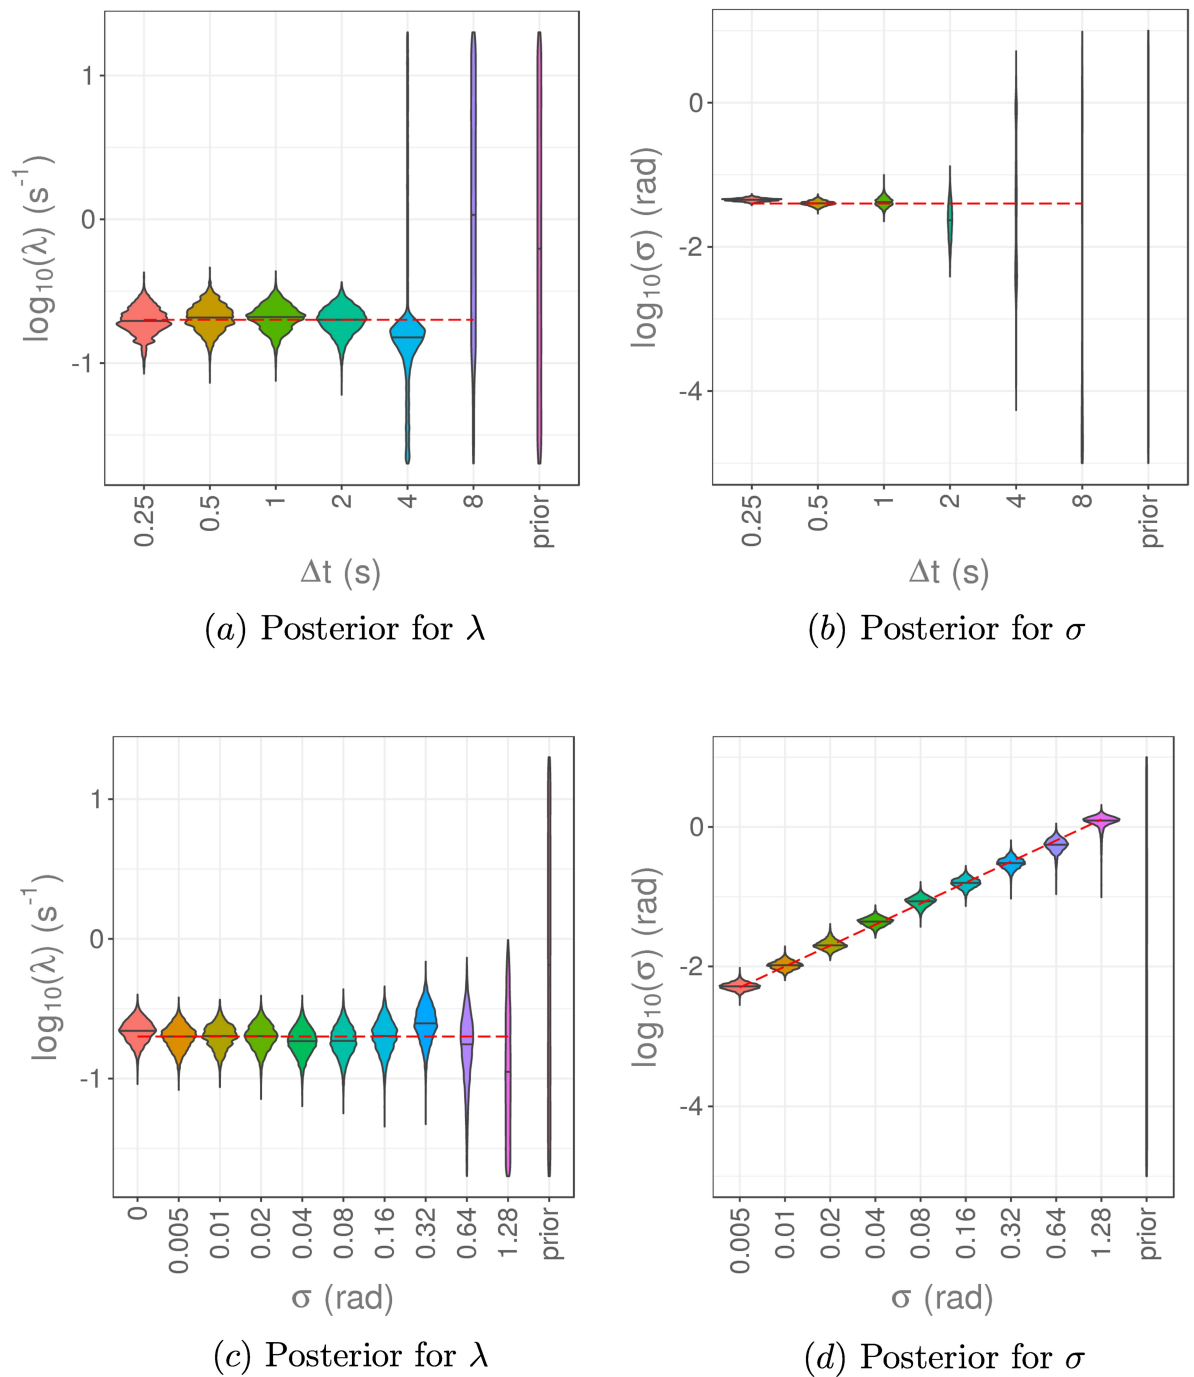


Fig 7. Results of parameter estimation via pMCMC for the reorientation rate, λ , and measurement noise, σ . Data collected with different values of Δt and fixed measurement noise, $\sigma = 0.04$ rad, were used in a) and b). Data generated with a fixed value of $\Delta t = 1$ s and different values of the measurement noise, σ , were used in c) and d). The marginal posterior distributions for λ are given in a) and c), while the marginal posterior distributions for σ are shown in b) and d). The red dashed lines indicate the true values of parameters used in simulation of the datasets. A reorientation rate of $\lambda = 0.2 \text{ s}^{-1}$ was used throughout.

<https://doi.org/10.1371/journal.pcbi.1006235.g007>

The results of using pMCMC to infer model parameters are shown in Fig 7 for $\Delta t = 1/8, 1/4, 1/2, 1, 2, 4, 8$ seconds, with noise of unknown magnitude σ , where $\sigma = 0.04$ rad. The estimated posterior distributions are very similar for small Δt , but we observe an abrupt breakdown in the quality of the posterior distributions obtained as Δt is increased. This breakdown

in posterior quality arises for values of $\Delta t \geq 2s$, as multiple reorientation events in a time interval start to become more common (see also S2 Fig). Provided that Δt and λ are such that the probability of multiple reorientations in a time interval is small, we obtain accurate estimates of the joint posterior distribution for the model parameters.

Increasing σ increases the variance of the posterior for λ

To investigate sensitivity of the posterior to measurement noise, we now fix the discretisation, Δt , and vary the measurement noise amplitude, σ , used to create each dataset. Applying the same analysis as in the previous subsection, for a fixed value of $\Delta t = 1s$, we obtain posterior distributions for the reorientation rate, λ , and measurement noise, σ , as shown in Fig 7. We find that increasing the noise amplitude, σ , increases the variance in the posterior distribution obtained for the reorientation rate, λ . In addition, we are able to accurately estimate the value of σ used to generate the datasets.

We note that the presence of noise in the datasets results in a bias towards smaller estimates of the reorientation rate, λ . This can be explained intuitively by reasoning that some reorientations leading to very small observed angle changes are mistaken for noise as σ increases.

More data provides better estimates

Let T_{final} be the total duration of our observations of the system. For a fixed value of the time discretisation, Δt , varying T_{final} is equivalent to gathering a bigger dataset. We fix $\Delta t = 1s$ and allow T_{final} to vary so that we can consider the effects of collecting more data. In practice, collecting more data experimentally may come at a cost. Quantifying the benefits of collecting more data with an *in silico* model can aid decisions about whether it is worthwhile to collect a larger dataset. Posteriors for the reorientation rate, λ , are shown in Fig 8.

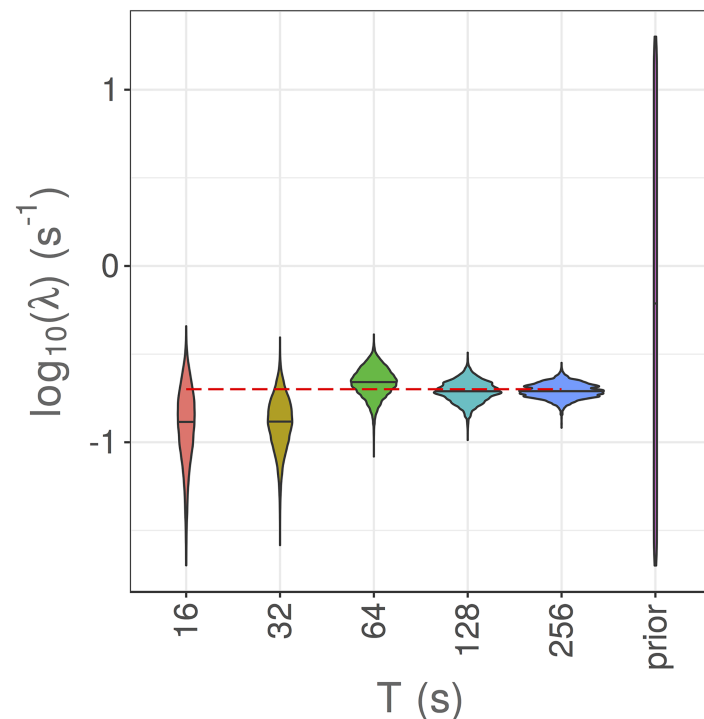


Fig 8. Results of parameter estimation via pMCMC for the reorientation rate, λ , with $\Delta t = 1s$ and $\sigma = 0rad$, whilst varying T_{final} to give different sized datasets.

<https://doi.org/10.1371/journal.pcbi.1006235.g008>

It is clear that larger datasets result in less bias and less variance for estimates of the reorientation rate, λ . However, we note that running the particle filter within the pMCMC algorithm becomes much more computationally intensive as the size of the dataset increases, scaling as $\mathcal{O}(T)$ as the size of the dataset increases.

Experimental constraints

We investigate the implications of these results for experimental design by considering an imaging experiment to observe the position of a particle of interest (for example, a bacterium) at regular time points. The behaviour of our system of interest happens over a certain timescale inherent to the biological process, so we fix the total duration for the imaging experiment, T_{final} , based on this timescale. We consider how best to choose the time between successive observations, Δt , given the restriction of a fixed photon budget. That is, we assume that the biological sample can only be exposed to a fixed number of photons before phototoxicity or photobleaching significantly reduce the quality of, or destroy, any further potential data. We assume that the sample is only exposed to photons during imaging. The results of Zhao et al. [39] suggest that the signal to noise ratio (SNR) is proportional to the square root of the time between successive frames, Δt , giving a relationship of the form

$$\text{SNR} = \kappa\sqrt{\Delta t}, \quad (8)$$

where κ is a constant that depends on the imaging set up, but not other experimental design choices. Assuming that for our model the noise is of magnitude σ , we find an inverse square relationship between σ and Δt , such that

$$\sigma = \frac{K}{\sqrt{\Delta t}}, \quad (9)$$

for a new constant K which is κ times the average angle change.

To investigate how to choose Δt given a fixed photon budget, we set a value of the proportionality constant and vary σ and Δt according to this relationship. We take proportionality constants $K = 0.08$ and $K = 0.8$ in Fig 9. A larger value of the proportionality constant, K , corresponds to worse imaging conditions, in that for a fixed value of Δt , the noise in the images obtained will be greater. Therefore we expect our inference method to perform worse for a larger value of K . We then ask the question, for a given value of K , how should we choose Δt to improve the parameter estimation process? Our results in Fig 9 suggest that the value of Δt should be taken as small as possible, even if this increases the noise present in the data.

Comparison to approximate Bayesian computation

A common approach for mathematical and computational models where the likelihood is intractable is to apply ABC for parameter inference [34, 35]. Although ABC produces samples from an approximate posterior, rather than the exact posterior distribution, it is intuitively simple to understand and implement. ABC methods have been applied to parameter estimation for biased, persistent random walk models very similar to our VJP model [12, 13, 40, 41]. For these methods, the choice of which summary statistics to use has notable effects on the approximate posterior distributions obtained via ABC. We consider three different summary statistic choices: the full time series of observed angle changes, a threshold-based summary statistic, and a transition matrix summary statistic.

We compare the quality of resulting posterior distributions obtainable with ABC to those from pMCMC. Our results, shown in Fig 10, indicate that the choice of summary statistic has a substantial effect on the inferred posterior distribution. When the dimensionality of the data

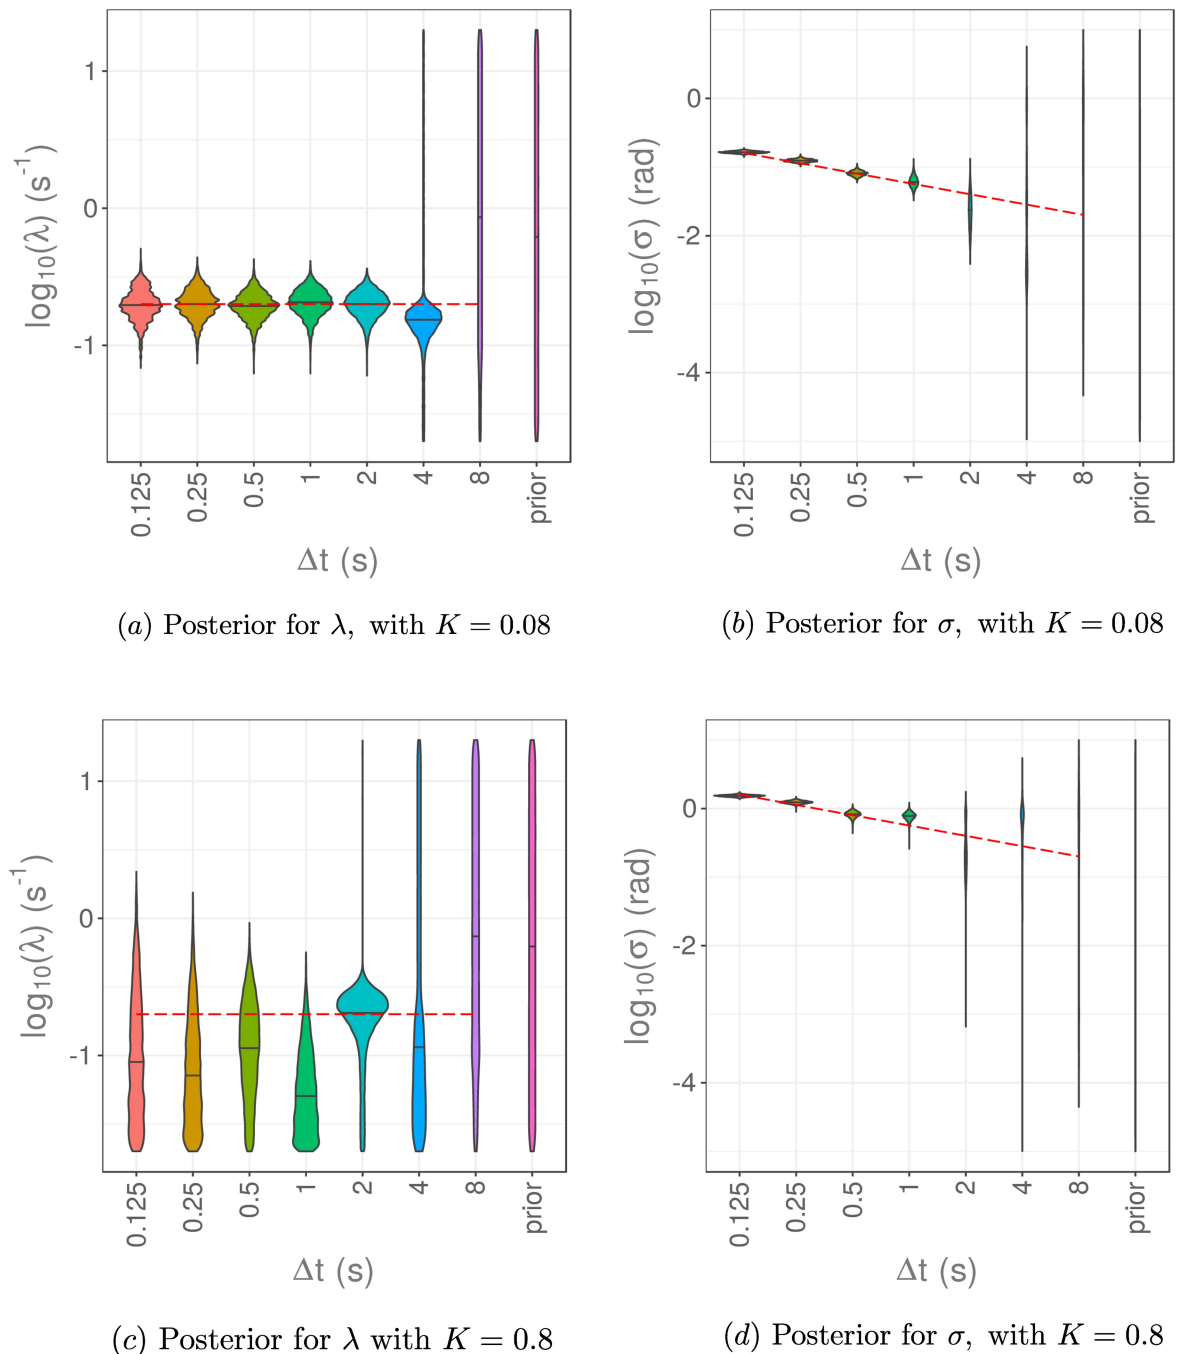


Fig 9. Results of parameter estimation via pMCMC for the reorientation rate, λ , and measurement noise, σ , using data generated with a fixed value of Δt and different values of the measurement noise, σ . We have varied both σ and Δt with an inverse square root relationship between these, as in Eq (9). The posterior distributions for λ are given in a) and c), while the posterior distributions for σ are shown in b) and d). The proportionality constant is $K = 0.08$ for a) and b), and $K = 0.8$ for c) and d).

<https://doi.org/10.1371/journal.pcbi.1006235.g009>

observed is high, ABC performs poorly at approximating the posterior for the parameters of our VJP, which we were able to sample exactly using pMCMC. For the full time series summary statistic, the dimensionality is high (up to 511 dimensions for $\Delta t = 0.125s$), meaning that the approximation we obtain to the posterior is very poor (see Fig 10a) and 10b)). Since

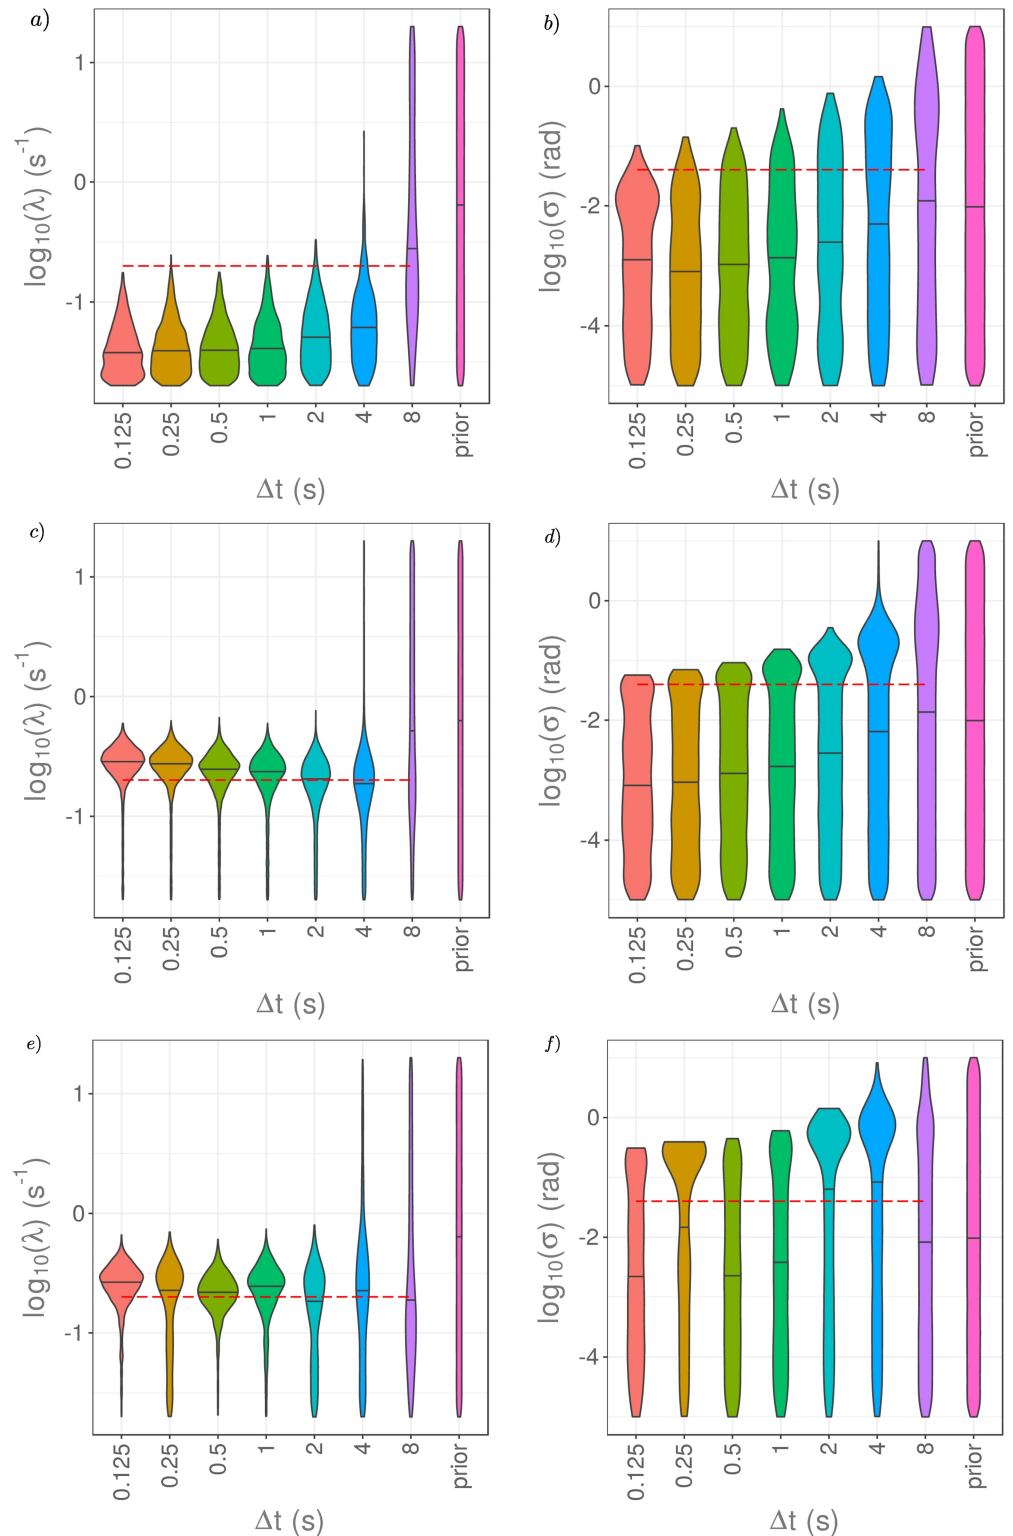


Fig 10. Parameter estimation via ABC rejection sampling for the reorientation rate, λ , in a), c), and e) and for the measurement noise, σ , in b), d), and f) using different summary statistics and data collected with noise of magnitude $\sigma = 0.04\text{rad}$, for different values of Δt . $N = 1,000,000$ datasets were generated and parameter samples corresponding to the closest 0.1% of the datasets were retained to give the approximate posterior. As summary statistics, we take the full time series of observed angle changes in a) and b). In c) and d), we use a simple threshold

summary statistic counting the number of observed angle changes with magnitude above a threshold, here taken as $h = 0.1$ rad. In d) and e), we use a transition matrix summary statistic using $n = 5$ bins to discretise the observed angle change.

<https://doi.org/10.1371/journal.pcbi.1006235.g010>

reorientation events are rare, particularly for small Δt , a greater fraction of the distance between simulated datasets and observed data is accounted for by the observation noise. ABC excludes part of the parameter space considered in the prior, but does poorly at identifying the reorientation rate, λ , since reorientations are rare events.

The threshold summary statistic is effective at identifying the number of reorientation events and provides a reasonable approximation of the marginal posterior for λ . However, it offers very little information about the observation noise, σ , except in relation to the threshold chosen. As a result the approximate marginal posteriors for σ are poor (see Fig 10c) and 10d)). Another summary statistic more informative about σ could be used in addition here to improve results.

When using the transition matrix summary statistic, the approximate posteriors for small Δt provide better estimates of the true reorientation rate and are much closer to the true posterior, as shown in Fig 10e) and 10f). The transition matrix is able to capture the distribution of the angle changes and dependence on recent history, whilst also reducing the dimensionality of the data. However, the transition matrix is unable to distinguish angle changes smaller than the resolution of the discretisation, $2\pi/n$, and so offers limited information about the measurement noise, σ . This issue could be mitigated by increasing the number of bins used, n , to give a finer discretisation of the observed angle changes (which would require sufficient data), or by targeting the noise with an additional summary statistic.

We still notice a deterioration in the quality of the approximate posterior distributions as Δt increases when using ABC for each choice of summary statistic even though there are no explicit assumptions made about multiple reorientations within a time interval, as was the case when using pMCMC. This suggests that the lack of information content about the parameters may be a property of the data, rather than the inference method. For large values of Δt , the data contains limited information about the model parameters, particularly the reorientation rate, λ .

To obtain potentially improved results for inference with ABC for this type of data (time series observations of angle changes), we could consider a more systematic choice of summary statistics [42–44] or a more efficient version of the ABC algorithm, such as ABC-SMC [45, 46], which applies sequential Monte Carlo methods to generate a sequence of approximations to the posterior distribution, using the previous posterior distribution to propose parameters for the next approximation. Other improvements could be possible by applying a regression adjustment, via linear regression [35] or using nonlinear regression techniques such as with a neural network [47].

Model misspecification

In general, in applying a model to a real world dataset, any model that we choose will be an approximation of the true data generating process. Before applying our inference methods to real world data, where we will fit to a model that is but an approximation to the true biological process, it is important to check the robustness of our method to fit a misspecified model. We investigate this robustness computationally by considering a misspecified model which should be no longer misspecified in an appropriate limit. It has previously been shown by Frazier et al. [48] that inference with ABC on a misspecified model can concentrate posterior mass around different pseudo-true parameter values depending on the version of ABC used.

Here, we demonstrate the robustness of our inference framework, using the pMCMC algorithm within a hidden states framework, to give relevant estimates for a misspecified model. We generate synthetic datasets using a wrapped normal reorientation kernel with dispersion parameter γ , but choose to estimate the model parameters with a model different to the data generating process by assuming a uniform reorientation kernel. As previously, we attempt to infer posterior distributions for the reorientation rate, λ , and the measurement noise, σ . The effect of misspecifying the model will be that some of the transmission probabilities used in the particle filter estimate of the likelihood of model parameters given the data will be wrong (see S3 Fig). This will give rise to some bias in our estimates of the likelihood within the particle filter (as demonstrated by Fig 11), and hence to some bias in our final approximation of the posterior. We will consider the effect of varying the dispersion parameter, γ , in the reorientation kernel used to generate the synthetic dataset, on the resulting posterior distribution for the model parameters. In the limit $\gamma \rightarrow \infty$, the wrapped normal distribution converges to the uniform distribution on $(-\pi, \pi]$, meaning that the model is no longer misspecified. The misspecification of the model is most pronounced for smaller values of γ .

We find, for a fixed value of the measurement noise, σ , that a reasonable approximation to the true posterior can be obtained for values of the dispersion, γ , larger than the measurement noise, σ . The approximate posterior distributions obtained are shown in Fig 12. When the dispersion, γ , is a similar magnitude to the measurement noise, which here is $\sigma = 0.04\text{rad}$, then reorientation events are frequently missed resulting in low estimates of the reorientation rate. For larger values of the dispersion parameter, we are able to robustly sample from the posterior distribution for λ and σ , despite notable differences in the misspecified reorientation kernel compared to the assumed uniform reorientation kernel. In this case, when a reorientation occurs, the observed angle changes are sufficiently different to the background measurement noise, that often it is interpreted as a reorientation event by the particle filter, despite error in the emission probabilities due to the model misspecification.

Application to RNA transport dataset

To demonstrate the effectiveness of our inference framework, we apply it here to a dataset of tracks obtained from imaging the transport of RNA-protein (RNP) complexes in a *Drosophila* oocyte. These complexes move on microtubules via molecular motors [1]. Occasionally, the complexes fall off a microtubule, and reattach on a different microtubule moving in a different direction. We neglect the diffusive stationary phase between falling off and reattaching on microtubules, meaning that we assume a running phase is followed immediately by another running phase. We apply our modelling framework, assuming an exponentially distributed running time distribution, $f_r(t) = \lambda \exp(-\lambda t)$, with constant running speed, c , and a uniform reorientation kernel, $f_\phi(\phi) = 1/2\pi \mathbb{1}_{(-\pi, \pi]}(\phi)$.

We take 10 tracks of separate complexes moving in *Drosophila* nurse cells obtained from an *in vivo* imaging dataset of movements of stau protein available in Zimyanin et al. [49]. Each track is short since the dataset is imaged in a single plane in the z direction, meaning that complexes move out of the frame of view frequently. The tracks used are shown in Fig 13a). We infer a subposterior distribution for each individual track and combine these together to produce a single posterior distribution combining data from different tracks. To combine the subposteriors, we use the consensus Monte Carlo algorithm of Scott et al. [50] developed for running MCMC on large datasets, via the implementation in the R [51] package `parallelMCMCcombine` [52]. This produces an approximate posterior distribution for the turning rate, λ , and measurement noise, σ , as shown in Fig 13b). We find 95% credible intervals for the turning rate λ of $[1.06, 1.58] \text{ s}^{-1}$ and for the measurement noise σ of $[0.27,$

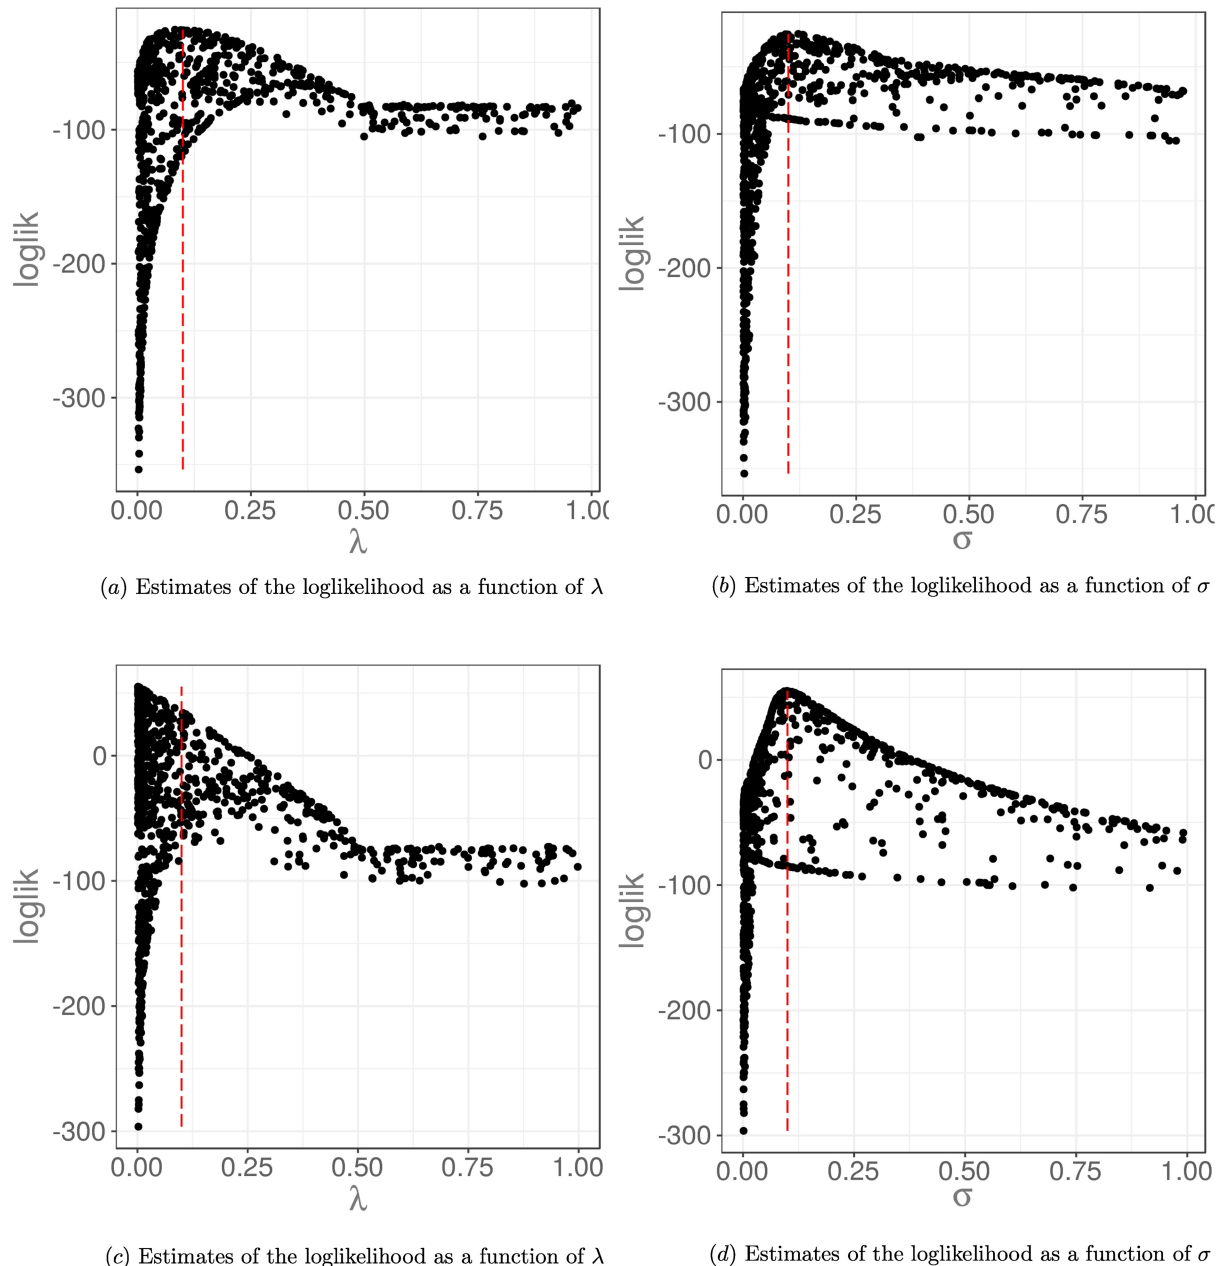


Fig 11. Estimates of the log likelihood via a particle filter for data generated with the true model and under a misspecified model. In a) and b), we show estimates of the log likelihood for different values of λ and σ based on a dataset simulated with a uniform reorientation kernel, and assuming this same reorientation kernel in estimating the log likelihood via the particle filter. The peak of the log likelihood coincides with the true values of the parameters shown by the red dashed line. In c) and d), we show estimates of the log likelihood for different values of λ and σ based on a dataset simulated with a wrapped normal reorientation kernel with dispersion $\gamma = 0.1$. However, we assume a uniform reorientation kernel in the particle filter. This model misspecification results in a slight shift in the peak of the log likelihood compared to the true parameter values used shown by the red dashed line. The log likelihood estimates will be biased in the case of the misspecified model.

<https://doi.org/10.1371/journal.pcbi.1006235.g011>

0.47] rad, respectively. Based on a VJP model corresponding to Brownian motion, the relation $\lambda = c^2/(2D)$ holds for a constant running speed c and diffusion constant D . The running speed in active transport for staufer RNPs is of the order $0.5\mu\text{ms}^{-1}$ [49], and based on the size of the staufer protein, we can estimate its diffusion coefficient as $1\mu\text{m}^2\text{s}^{-1}$ [53], which gives a

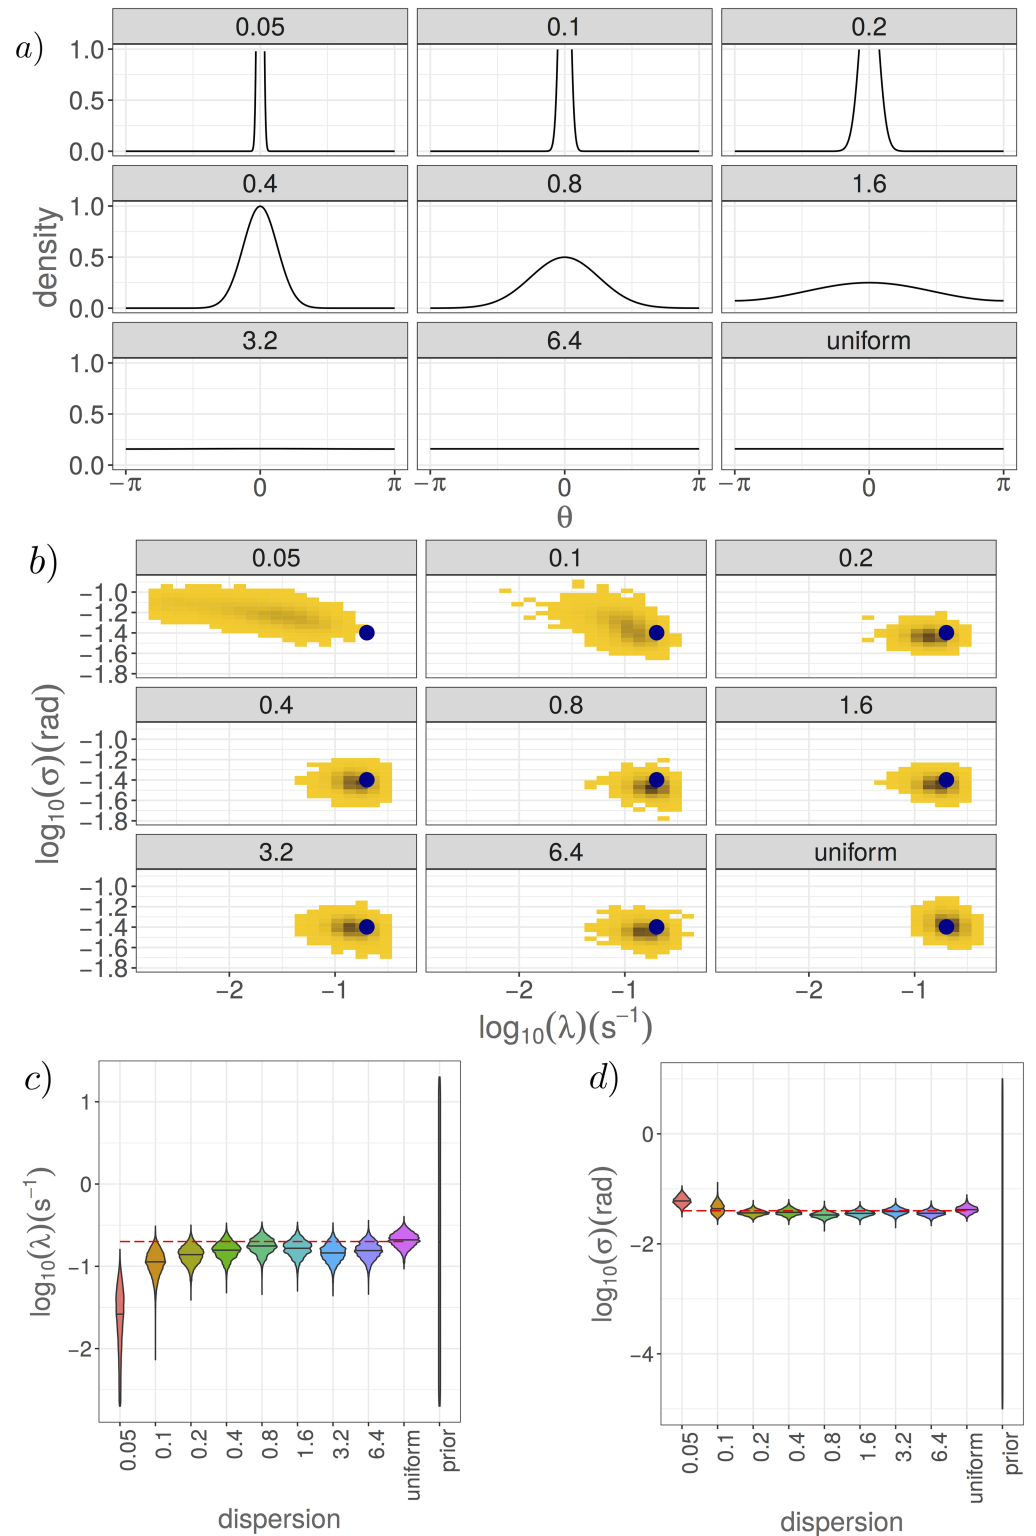


Fig 12. The posterior distribution obtained under model misspecification converges to the distribution obtained with the true model as the misspecified model approaches the true model. We assume a wrapped normal reorientation kernel with dispersion parameter γ and vary this dispersion parameter. For large γ , this tends in distribution towards a uniform reorientation kernel, as shown in a). The posterior distributions obtained from performing inference with data generated for different values of γ are shown in b), with the corresponding marginal

distributions shown in c) and d). In b), the true parameter values used are shown by the dark blue circle. In c) and d), the true parameter values are given by the red dashed line. Parameters used to generate the synthetic datasets are $\lambda = 0.2 \text{ s}^{-1}$, $\sigma = 0.04 \text{ rad}$, $\Delta t = 0.25 \text{ s}$, $c = 50 \mu\text{ms}^{-1}$, and the details of the pMCMC are as in Fig 7.

<https://doi.org/10.1371/journal.pcbi.1006235.g012>

similar order of magnitude for λ as the estimate obtained here. Evidently the dataset used here is very noisy, but nonetheless we are able to obtain estimates for the turning rate in a model of RNA transport.

Discussion

In this work, we have considered parameter estimation of a VJP model for biological transport and how insights from this parameter inference process can inform experimental design. We generated estimates of the posterior distribution for parameters of a VJP model based on noisy datasets collected at varying temporal resolutions. To perform this parameter inference, we used pMCMC and derived the appropriate emission probabilities. We observed an abrupt breakdown in the quality of the posteriors obtained when decreasing the temporal resolution. This transition corresponds to a breakdown in modelling assumptions underlying the derivation of the emission probabilities. For example, as Δt increases we see multiple reorientations in a time interval. These assumptions are necessary in the development of our inference framework which relies on hidden states consisting of binary variables indicating whether or not a reorientation occurred in a time interval.

Increasing the magnitude of the noise in the data slowly decreased the quality of the median posterior estimates and increased the posterior variance. In general, better estimates were obtained when more data was available, either by decreasing Δt , or by increasing the total duration of the experimental observations, T_{final} . These results are intuitive, but the real benefit here is in quantifying the effects of choices of these experimental hyperparameters to enable researchers to make decisions about experimental design.

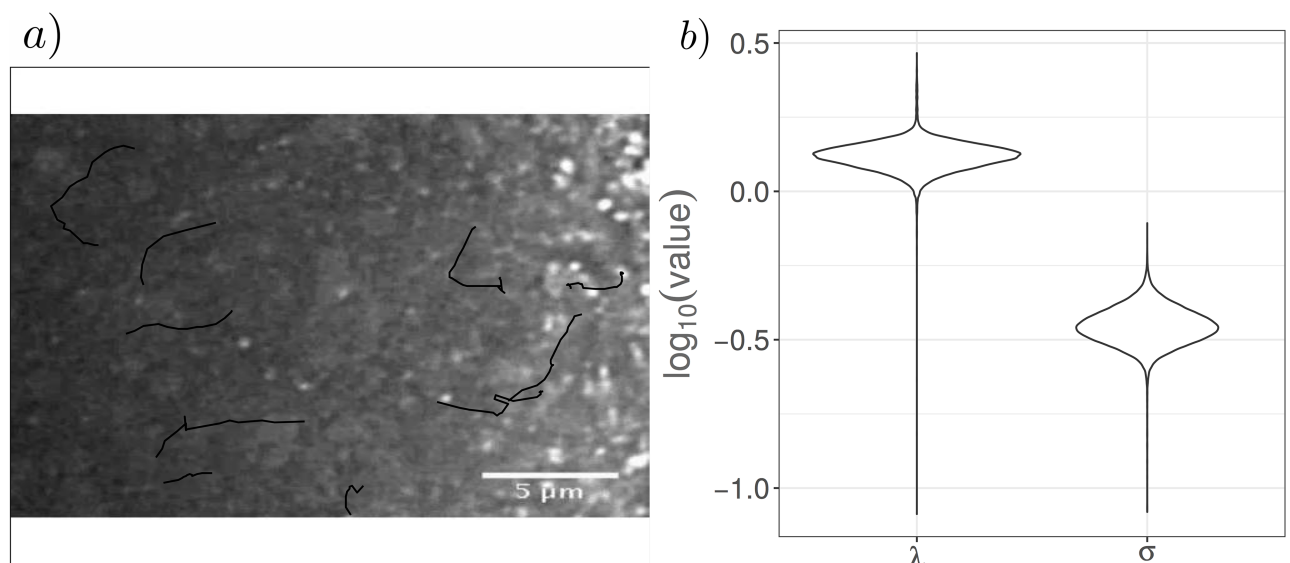


Fig 13. Posterior distribution for model parameters λ and σ for RNA transport in a *Drosophila* oocyte. The results are based on 10 tracks labelled manually and sampled at a resolution of $\Delta t = 0.478 \text{ s}$, with subposteriors obtained for each track and combined to give a single posterior via Consensus Monte Carlo. The tracks are shown drawn on the first frame in a), although not all tracks start from the same time point. The marginal posterior distributions for the reorientation rate, λ , and the measurement noise, σ , are shown in b).

<https://doi.org/10.1371/journal.pcbi.1006235.g013>

We also compared parameter estimation with pMCMC to that with ABC, and results suggest that different methods are appropriate in different situations, dependent on the data available. In particular, pMCMC performs well provided the assumptions made in deriving the emission probabilities hold, which is true for small Δt . For larger values of Δt , neither ABC nor pMCMC provide accurate samples from the full posterior distribution. The information about the model parameters present in the data for large Δt is limited. In terms of computational cost comparisons between ABC and pMCMC, an ABC rejection sampler is highly parallelisable over independent parameter samples [54] whereas pMCMC is less amenable to simple parallelisation. Making use of this parallelisation, our implementation of ABC rejection sampling was more computationally efficient than pMCMC. We show a quantitative comparison between the methods in S6 Fig. We note that the computational cost is strongly problem and implementation dependent.

Through comparison between inference with pMCMC and ABC, our results highlight a weakness of inference with ABC, in that it performs poorly for high dimensional data, and also a weakness of our application of pMCMC, which is that it relies upon assumptions about the number of reorientation events in a time step; when this fails our posterior estimates are no longer accurate. We note, additionally, that pMCMC allows us to sample from the exact posterior for parameters from a model (given that model is appropriate to describe the data), whereas via ABC we obtain approximate posteriors for a fixed tolerance, ϵ , which only become exact in the limit as $\epsilon \rightarrow 0$.

In Section ‘Methods’, we described a framework for inference using pMCMC, and made certain assumptions about the VJP model, such as a separation in timescales between runs and reorientations, and a memoryless exponential distribution for the running time. Although many of these are standard assumptions, it would be possible to perform the same analysis in a more general model. For instance, if a running time distribution was chosen that does not satisfy the memoryless property, we could introduce an extra hidden variable, s , for the time since the last reorientation. In addition, in this work we have described parameter estimation via a hidden states formulation of the VJP model using dependence on hidden states from two time intervals. This can be extended to hidden states from three time intervals, which can allow rare consecutive reorientation events to be handled more accurately, albeit at greater computational cost.

Given a fixed photon budget and a trade-off between temporal sampling frequency and measurement noise, our results in the Section ‘Experimental constraints’ indicate that a small value of Δt should be used for the discretisation in time i.e. that motile individuals should be imaged as frequently as possible. In practice, there may be disadvantages to this choice of Δt . Computationally, conducting parameter inference via pMCMC will be significantly more expensive, although the computational run time may still be small in comparison to the duration of an experimental protocol. Datasets with higher noise present may also be much harder to interpret. Here, we have assumed that the noise present in the data is applied to the observed angle change. In reality there is noise on each pixel, which may contribute to an uncertainty in identifying the observed position of the object of interest.

Additionally, we considered using our particle MCMC inference framework to fit data from a misspecified model; our results indicate that the framework is robust to moderate amounts of model misspecification. This allowed us to be confident in applying our framework to a real dataset of RNA tracks to consider the motility of RNA-protein complexes moving on the cytoskeleton, demonstrating the ability of our framework to obtain parameter estimates in challenging conditions with very noisy data.

Supporting information

S1 Appendix. Comparison with simulations. Description of simulations to verify the analytic form for the emission probabilities.

(PDF)

S1 Fig. Comparison with simulations. Comparison between simulated results and theoretical predictions of the observed angle change for hidden states of the form 0, 1 in a), 1, 0 in b), and 1, 1 in c). The simulated results are shown by the histogram and the theoretical prediction for the observed angle change distribution is shown as the red dashed line. For both b) and c), we have conditioned on an observed angle change in the previous time interval of 0.1 rad, and for c) we also conditioned on an observed angle change prior to that of -1.0 rad. To generate these results, we used $N = 10^7$ simulated trajectories with running speed $c = 50\mu\text{ms}^{-1}$, uniform reorientation kernel, reorientation rate $\lambda = 0.2\text{ s}^{-1}$ and time discretization $\Delta t = 1\text{ s}$.

(TIF)

S2 Fig. Probability of hidden state sequences. The probability of sequences of hidden states as Δt varies with reorientation rate $\lambda = 0.2\text{ s}^{-1}$. For large values of Δt , the assumptions of the model start to break down as multiple reorientations appear within a single time interval.

(TIF)

S3 Fig. Emission probabilities for misspecified model. Comparison between assumed distribution of observed angle changes and simulated distributions when using a misspecified model for the reorientation kernel. We simulate $N = 10^7$ angle changes using a wrapped normal reorientation kernel with dispersion parameter γ given a certain sequence of hidden states (0, 1 in a), d), g); 1, 0 in b), e), h); 1, 1 in c), f), i)) and show a grey histogram of the simulated observed angle changes. To demonstrate the misspecification in the emission probabilities, we plot the assumed theoretical distribution of the observed angle changes as a red dashed line, based on assuming that the reorientation kernel is a uniform distribution. The model is more misspecified for a smaller value of the dispersion parameter γ . As in S1 Fig, we have conditioned on an observed angle change in the previous time interval of 0.1 rad for b), c), e), f), h), and i). For c), f), and i), we also conditioned on an observed angle change prior to that of -1.0 rad. We used a dispersion parameter, γ , in the reorientation kernel of $\gamma = 0.4$ for a), b), and c), $\gamma = 1$ for d), e), and f), and $\gamma = 6.4$ for g), h), and i).

(TIF)

S4 Fig. Traceplots and autocorrelation functions for MCMC chains. Traceplots and autocorrelation functions for MCMC chains to analyse convergence are shown for data generated with parameters $\lambda = 0.2\text{ s}^{-1}$, $\Delta t = 0.25\text{ s}$, $\sigma = 0.04\text{ rad}$, as for the first posterior shown in Fig 7a). Similar results are seen in sampling for other posterior distributions shown.

(TIF)

S5 Fig. Analysis of number of particles in the particle filter. The number of particles used in the particle filter affects the variance of estimates of the log-likelihood. However, a higher computational cost is required to use more particles. In a), we show how the distribution of the log-likelihood estimates varies (provided the filter does not become degenerate) as we change the number of particles. We use 1000 runs of the particle filter with the specified number of particles and estimate the log-likelihood at the true value of the parameters ($\lambda = 0.2\text{ s}^{-1}$ and $\sigma = 0.08\text{ rad}$) used to generate a synthetic dataset. In b), we illustrate the variance in the (nondegenerate) log-likelihood estimates, and the mean time to obtain a single estimate. A moderate increase in the compute time to run the particle filter offers substantial decrease in

the variance of the log-likelihood estimates. To strike a reasonable balance, we use 400 particles to generate the results presented in this work.

(TIF)

S6 Fig. The computational cost of pMCMC and ABC per effective sample size. The computational cost of parameter estimation with pMCMC and ABC per effective sample size depends on Δt . We quantify here a comparison between the computational cost of the pMCMC and ABC methods for parameter estimation when varying Δt . Datasets were generated with $\sigma = 0.04$ rad as in Fig 7a). The cost to produce a sample via pMCMC increases as Δt decreases, as shown by the grey squares. The cost for ABC remains approximately constant with Δt as we simulate data from the model a fixed number of times, N . The acceptance rate, α , in ABC affects the sample size we produce for a fixed number of simulations, N . For the results in Fig 10, an acceptance rate of $\alpha = 0.1\%$ was used. This gives a cost per sample lower than for pMCMC, shown by the blue circles and triangles. For a smaller acceptance rate, $\alpha = 0.01\%$ (shown by red circles and triangles), the cost per sample is much higher and the plots of the posterior are unchanged compared to those for $\alpha = 0.1\%$. The results for pMCMC are given by squares, ABC with transition matrix summary statistics are shown as circles and ABC with time series summary statistics are shown as triangles. We note that the computational cost depends strongly on the problem considered and the implementation used.

(TIF)

S1 Code. Code implementing Bayesian inference via pMCMC for a VJP model. Example code to implement pMCMC for a VJP model, as described in Section “Methods”, is available in a .zip folder or at https://github.com/shug3502/pmmc_inference_for_vjps.

(ZIP)

Acknowledgments

The authors are grateful to Richard Parton and Ilan Davis for helpful discussions of the application to RNA transport.

Author Contributions

Conceptualization: Jonathan U. Harrison, Ruth E. Baker.

Data curation: Jonathan U. Harrison.

Formal analysis: Jonathan U. Harrison, Ruth E. Baker.

Investigation: Jonathan U. Harrison, Ruth E. Baker.

Methodology: Jonathan U. Harrison.

Software: Jonathan U. Harrison.

Supervision: Ruth E. Baker.

Validation: Jonathan U. Harrison, Ruth E. Baker.

Visualization: Jonathan U. Harrison, Ruth E. Baker.

Writing – original draft: Jonathan U. Harrison.

Writing – review & editing: Jonathan U. Harrison, Ruth E. Baker.

References

1. Parton R M, Davidson A, Davis I, and Weil T T. Subcellular mRNA localisation at a glance. *Journal of Cell Science*, 127(10):2127–2133, 2014. <https://doi.org/10.1242/jcs.114272> PMID: 24833669
2. Othmer H G, Dunbar S R, and Alt W. Models of dispersal in biological systems. *Journal of Mathematical Biology*, 26(3):263–298, 1988. <https://doi.org/10.1007/BF00277392> PMID: 3411255
3. Berg H C and Brown D A. Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking. *Nature*, 239(5374):500–504, 1972. <https://doi.org/10.1038/239500a0> PMID: 4563019
4. Rosser G, Fletcher A G, Wilkinson D A, de Beyer J A, Yates C A, Armitage J P, Maini P K, and Baker R E. Novel methods for analysing bacterial tracks reveal persistence in *Rhodobacter sphaeroides*. *PLOS Computational Biology*, 9(10):1:18, 2013. <https://doi.org/10.1371/journal.pcbi.1003276> PMID: 24204227
5. Taylor-King J P, van Loon E, Rosser G, and Chapman S J. From birds to bacteria: generalised velocity jump processes with resting states. *Bulletin of Mathematical Biology*, 77(7):1213–1236, 2015. <https://doi.org/10.1007/s11538-015-0083-7> PMID: 26060098
6. Othmer H G and Hillen T. The diffusion limit of transport equations derived from velocity-jump processes. *SIAM Journal on Applied Mathematics*, 61(3):751–775, 2000. <https://doi.org/10.1137/S0036139999358167>
7. Codling E A and Hill N A. Calculating spatial statistics for velocity jump processes with experimentally observed reorientation parameters. *Journal of Mathematical Biology*, 51(5):527–556, 2005. <https://doi.org/10.1007/s00285-005-0317-7> PMID: 15868200
8. Codling E A, Plank M J, and Benhamou S. Random walk models in biology. *Journal of the Royal Society Interface*, 5(25):813–834, 2008. <https://doi.org/10.1098/rsif.2008.0014>
9. Kareiva P M and Shigesada N. Analyzing insect movement as a correlated random walk. *Oecologia*, 56(2):234–238, 1983. <https://doi.org/10.1007/BF00379695> PMID: 28310199
10. Bovet P and Benhamou S. Spatial analysis of animals' movements using a correlated random walk model. *Journal of Theoretical Biology*, 131(4):419–433, 1988. [https://doi.org/10.1016/S0022-5193\(88\)80038-9](https://doi.org/10.1016/S0022-5193(88)80038-9)
11. Jonsen I D, Flemming J M, and Myers R A. Robust state-space modeling of animal movement data. *Ecology*, 86(11):2874–2880, 2005. <https://doi.org/10.1890/04-1852>
12. Liepe J, Taylor H, Barnes C, Huvet M, Bugeon L, Thorne T, Lamb J, Dallman M, and Stumpf M P H. Calibrating spatio-temporal models of leukocyte dynamics against in vivo live-imaging data using approximate Bayesian computation. *Integrative Biology*, 4(3):335–345, 2012. <https://doi.org/10.1039/c2ib00175f> PMID: 22327539
13. Jones P, Sim A, Taylor H, Bugeon L, Dallman M, Pereira B, Stumpf M P H, and Liepe J. Inference of random walk models to describe leukocyte migration. *Physical Biology*, 12(6):66001–66012, 2015. <https://doi.org/10.1088/1478-3975/12/6/066001>
14. Viswanathan G M, Afanasyev V, Buldyrev S V, Murphy E J, Prince P A, and Stanley H E. Lévy flight search patterns of wandering albatrosses. *Nature*, 381(6581):413–415, 1996. <https://doi.org/10.1038/381413a0>
15. Viswanathan G M, Afanasyev V, Buldyrev S V, Havlin S, Da Luz M G, Raposo E P, and Stanley H E. Lévy flights in random searches. *Physica A*, 282(1):1–12, 2000. [https://doi.org/10.1016/S0378-4371\(00\)00071-6](https://doi.org/10.1016/S0378-4371(00)00071-6)
16. Edwards A M, Phillips R A, Watkins N W, Freeman M P, Murphy E J, Afanasyev V, Buldyrev S V, da Luz M G, Raposo E P, and Stanley H E. Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature*, 449(7165):1044–1048, 2007. <https://doi.org/10.1038/nature06199> PMID: 17960243
17. Benhamou S. How many animals really do the Lévy walk? *Ecology*, 88(8):1962–1969, 2007. <https://doi.org/10.1890/06-1769.1> PMID: 17824427
18. Painter K J. Modelling cell migration strategies in the extracellular matrix. *Journal of Mathematical Biology*, 58(4):511–543, 2009. <https://doi.org/10.1007/s00285-008-0217-8> PMID: 18787826
19. Mandeville J T, Ghosh R N, and Maxfield F R. Intracellular calcium levels correlate with speed and persistent forward motion in migrating neutrophils. *Biophysical Journal*, 68(4):1207–1217, 1995. [https://doi.org/10.1016/S0006-3495\(95\)80336-X](https://doi.org/10.1016/S0006-3495(95)80336-X) PMID: 7787012
20. Pankov R, Endo Y, Even S-Ram, Araki M, Clark K, Cukierman E, Matsumoto K, and Yamada K M. A Rac switch regulates random versus directionally persistent cell migration. *Journal of Cell Biology*, 170(5):793–802, 2005. <https://doi.org/10.1083/jcb.200503152> PMID: 16129786

21. Nicosia A, Duchesne T, Rivest L P, and Fortin D. A general hidden state random walk model for animal movement. *Computational Statistics and Data Analysis*, 105:76–95, 2017. <https://doi.org/10.1016/j.csda.2016.07.009>
22. Andrieu C and Roberts G O. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009. <https://doi.org/10.1214/07-AOS574>
23. Andrieu C, Doucet A, and Holenstein R. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010. <https://doi.org/10.1111/j.1467-9868.2009.00736.x>
24. Gordon N J, Salmond D J, and Smith A F. Novel approach to nonlinear non-Gaussian Bayesian state estimation. In *IEE Proceedings F*, volume 140, pages 107–113. IET, 1993.
25. Rasmussen D A, Ratmann O, and Koelle K. Inference for nonlinear epidemiological models using genealogies and time series. *PLOS Computational Biology*, 7(8):e1002136, 2011. <https://doi.org/10.1371/journal.pcbi.1002136> PMID: 21901082
26. Golightly A and Wilkinson D J. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807, 2011. <https://doi.org/10.1098/rsfs.2011.0047> PMID: 23226583
27. Golightly A, Henderson D A, and Sherlock C. Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing*, 25(5):1039–1055, 2015. <https://doi.org/10.1007/s11222-014-9469-x>
28. Grimmett G and Stirzaker D. *Probability and Random Processes*. Oxford University Press, 2001.
29. Rosser G. *Mathematical modelling and analysis of aspects of planktonic bacterial motility*. PhD thesis, University of Oxford, 2012.
30. Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H, and Teller E. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. <https://doi.org/10.1063/1.1699114>
31. Doucet A, Godsill S, and Andrieu C. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000. <https://doi.org/10.1023/A:1008935410038>
32. Sherlock C, Thiery A H, Roberts G O, and Rosenthal J S. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *The Annals of Statistics*, 43(1):238–275, 2015. <https://doi.org/10.1214/14-AOS1278>
33. Pitt M K, dos Santos Silva R, Giordani P, and Kohn R. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2):134–151, 2012. <https://doi.org/10.1016/j.jeconom.2012.06.004>
34. Pritchard J K, Seielstad M T, Perez-Lezaun A, and Feldman M W. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 1999. <https://doi.org/10.1093/oxfordjournals.molbev.a026091> PMID: 10605120
35. Beaumont M A, Zhang W, and Balding D J. Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002. PMID: 12524368
36. Prangle D. Adapting the ABC distance function. *arXiv preprint arXiv:1507.00874*, 2015.
37. Harrison J U and Baker R E. An automatic adaptive method to combine summary statistics in approximate Bayesian computation. *arXiv preprint arXiv:1703.02341*, 2017.
38. E Bernton, Jacob P E, Gerber M, and Robert C P. Inference in generative models using the wasserstein distance. *arXiv preprint arXiv:1701.05146*, 2017.
39. Zhao Q, Young I T, and De Jong J G. Photon budget analysis for fluorescence lifetime imaging microscopy. *Journal of Biomedical Optics*, 16(8):086007–086007, 2011. <https://doi.org/10.1117/1.3608997> PMID: 21895319
40. Taylor H B, Liepe J, Barthen C, Bugeon L, Huvet M, Kirk P W, Brown S B, Lamb J R, Stumpf M P H, and Dallman M. P38 and JNK have opposing effects on persistence of in vivo leukocyte migration in zebrafish. *Immunology and Cell Biology*, 91(1):60–69, 2013. <https://doi.org/10.1038/icb.2012.57> PMID: 23165607
41. Weavers H, Liepe J, Sim A, Wood W, Martin P, and Stumpf M P H. Systems analysis of the dynamic inflammatory response to tissue damage reveals spatiotemporal properties of the wound attractant gradient. *Current Biology*, 26(15):1975–1989, 2016. <https://doi.org/10.1016/j.cub.2016.06.012> PMID: 27426513
42. Fearnhead P and Prangle D. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B*, 74(3):419–474, 2012. <https://doi.org/10.1111/j.1467-9868.2011.01010.x>

43. Barnes C, Filippi S, Stumpf M P H, and Thorne T. Considerate approaches to constructing summary statistics for ABC model selection. *Statistics and Computing*, 22(6):1181–1197, 2012. <https://doi.org/10.1007/s11222-012-9335-7>
44. Blum M G, Nunes M A, Prangle D, and Sisson S A. A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science*, 28(2):189–208, 2013. <https://doi.org/10.1214/12-STS406>
45. Sisson S A, Fan Y, and Tanaka M M. Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007. <https://doi.org/10.1073/pnas.0607208104>
46. Toni T, Welch D, Strelkova N, Ipsen A, and Stumpf M P H. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009. <https://doi.org/10.1098/rsif.2008.0172>
47. Blum M G and François O. Non-linear regression models for approximate Bayesian computation. *Statistics and Computing*, 20(1):63–73, 2010. <https://doi.org/10.1007/s11222-009-9116-0>
48. Frazier D T, Robert C P, and Rousseau J. Model misspecification in abc: Consequences and diagnostics. *arXiv preprint arXiv:1708.01974*, 2017.
49. Zimyanin V L, Belaya K, Pecreaux J, Gilchrist M J, Clark A, Davis I, and St Johnston D. In vivo imaging of oskar mRNA transport reveals the mechanism of posterior localization. *Cell*, 134(5):843–853, 2008. <https://doi.org/10.1016/j.cell.2008.06.053> PMID: 18775316
50. Scott S L, Blocker A W, Bonassi F V, Chipman H A, George E I, and McCulloch R E. Bayes and big data: The consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management*, 11(2):78–88, 2016. <https://doi.org/10.1080/17509653.2016.1142191>
51. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL <http://www.R-project.org/>.
52. Miroshnikov A and Conlon E M. ParallelMCMCcombine: an R package for Bayesian methods for big data and analytics. *PLOS One*, 9(9):e108425, 2014. <https://doi.org/10.1371/journal.pone.0108425> PMID: 25259608
53. Kumar M, Mommer S M and Sourjik V. Mobility of cytoplasmic, membrane, and DNA-binding proteins in Escherichia coli. *Biophysical Journal*, 98(4):552–559, 2010. <https://doi.org/10.1016/j.bpj.2009.11.002> PMID: 20159151
54. Owen J, Wilkinson D J, and Gillespie C S. Scalable inference for Markov processes with intractable likelihoods. *Statistics and Computing*, 25(1):145–156, 2015. <https://doi.org/10.1007/s11222-014-9524-7>