

Methods in Ecology and Evolution

State-switching continuous-time correlated random walks

Théo Michelot^{1*}, Paul G. Blackwell¹

¹*School of Mathematics and Statistics, University of Sheffield, UK*

Abstract

1. Continuous-time models have been developed to capture features of animal movement across temporal scales. In particular, one popular model is the continuous-time correlated random walk, in which the velocity of an animal is formulated as an Ornstein-Uhlenbeck process, to capture the autocorrelation in the speed and direction of its movement. In telemetry analyses, discrete-time state-switching models (such as hidden Markov models) have been increasingly popular to identify behavioural phases from animal tracking data.
2. We propose a multistate formulation of the continuous-time correlated random walk, with an underlying Markov process used as a proxy for the animal's behavioural state process. We present a Markov chain Monte Carlo algorithm to carry out Bayesian inference for this multistate continuous-time model.
3. Posterior samples of the hidden state sequence, of the state transition rates, and of the state-dependent movement parameters can be obtained. We investigate the performance of the method in a simulation study, and we illustrate its use in a case study of grey seal (*Halichoerus grypus*) tracking data.
4. The method we present makes use of the state-space model formulation of the continuous-time correlated random walk, and can accommodate irregular sampling frequency and measurement error. It will facilitate the use of continuous-time models to estimate movement characteristics and infer behavioural states from animal telemetry data.

Keywords: animal movement, continuous time, Ornstein-Uhlenbeck process, state-space model, multistate model, random walk

*Corresponding author: tmichelot1@sheffield.ac.uk

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi:

10.1111/2041-210X.13154

This article is protected by copyright. All rights reserved.

Accepted Article

1 Introduction

The collection of large high-resolution animal tracking data sets has motivated the development of a wide range of statistical methods (Patterson et al., 2017; Hooten et al., 2017). For the analysis of animal movement, an important conceptual modelling choice is the time formulation, i.e. the choice between discrete-time and continuous-time models (McClintock et al., 2014). Although animals move in continuous time, their location may only be observed at discrete intervals (e.g. every minute or every hour). Discrete-time approaches are based on the assumption that the underlying movement process can be appropriately modelled at the time scale of the observations. Most often, movement is described by the “step lengths” (distances between successive locations) and “turning angles” (angles between successive directions), derived from the location data (Morales et al., 2004; Jonsen et al., 2005). However, the distributions of these metrics of movement strongly depend on the sampling rate, such that the resulting inference is tied to a specific temporal scale (Codling and Hill, 2005; Schlägel and Lewis, 2016). One of the consequences is that discrete-time methods require locations to be collected at regular time intervals through the period of the study. Many telemetry data sets are collected at irregular time intervals, for example with marine mammals which may only be observed when they surface. To use a discrete-time model in such cases, it is then necessary to interpolate the data points on a regular time grid, introducing approximation uncertainty.

On the other hand, continuous-time models consider that telemetry observations arise from a continuous movement process. As such, they can naturally accommodate different temporal scales, and irregular sampling rates (Patterson et al., 2017). Various approaches have been used to model animal movement in continuous time, most of them based on diffusion processes. These include Ornstein-Uhlenbeck processes (Dunn and Gipson, 1977; Blackwell, 1997, 2003), Brownian bridges (Horne et al., 2007), and potential functions (Preisler et al., 2013). A very popular model is the continuous-time correlated random walk (CTCRW) introduced by Johnson et al. (2008), in which the velocity of the animal is formulated as an Ornstein-Uhlenbeck process. Through the velocity, this model incorporates autocorrelation into both the speed and the direction of the movement, similarly to discrete-time correlated random walks based on step lengths and turning angles. Johnson et al. (2008) formulated the CTCRW as a state-space model, making fast inference possible through the Kalman filter, and made it available in the R package *crawl* (Johnson and London, 2018). Fleming et al. (2017) extended this implementation to a wider family of diffusion processes, including their “OUF” model of correlated movement around a centre of attraction.

Random walks have been used as “building blocks” for more complex, multistate, models. These state-switching models describe animal movements as the outcome of several distinct behaviours, e.g. “foraging”, “resting”, “exploring”, based on the notion that the behavioural states of the animal differ noticeably in

terms of some metrics of the movement, e.g. speed or sinuosity (Morales et al., 2004). The advantage of multistate time series models over simpler clustering methods is that they account for the temporal autocorrelation in the movement behaviours (Edelhoff et al., 2016). Although this idea has received a lot of attention for discrete-time models, with the growing popularity of hidden Markov models (Patterson et al., 2009; Langrock et al., 2012), it has been underutilised in continuous-time approaches. Blackwell (1997) introduced a multistate movement model, where the location of the animal is modelled with an Ornstein-Uhlenbeck process. That model does not directly capture the movement persistence in speed and direction, which makes its application limited for high-frequency tracking data. More recently, Parton and Blackwell (2017) described a multistate approach in which the speed and the bearing of the animal are modelled with diffusion processes, analogously to discrete-time models based on step lengths and turning angles. However, their method requires computationally-costly numerical approximation to reconstruct the movement path at a fine time scale, a disadvantage in dealing with large tracking data sets. McClintock et al. (2014) presented a multistate analysis based on the CTCRW, but they constrained the state process to be constant over each time interval between two observations. Therefore, they do not carry out exact inference from the continuous-time model. Gurarie et al. (2017) recently reviewed the use of the CTCRW model for the analysis of animal tracking data. They proposed a method based on change point analysis to segment movement tracks into behavioural phases. Although they can be a useful tool of classification, change point approaches do not provide a mechanistic understanding of the behavioural state process.

Alternatively, to circumvent the limitations of discrete-time models, McClintock (2017) suggested a two-stage approach based on multiple imputation methods. A one-state continuous-time model (such as the CTCRW) is fitted to the data, and a large number m of possible realisations of the movement process are simulated from the model on a regular time grid. Then, a hidden Markov model is fitted to each realisation, to investigate the state-switching dynamics. The m sets of estimates are pooled, such that the resulting model takes into account the uncertainty in the locations. Note that, since the realisations are generated without taking into account the possible behaviours, this is not fully equivalent to fitting a multistate CTCRW model. In particular, if the one-state model fails to capture the behavioural heterogeneity in the movement, the simulated realisations may not correctly reflect the uncertainty in the continuous trajectory.

Here, we extend the framework of Johnson et al. (2008) to incorporate behavioural states directly into the CTCRW framework, with an underlying continuous-time Markov process. The state-switching CTCRW offers a rigorous approach to model behavioural heterogeneity and movement persistence, two common features in modern telemetry data collected at high frequency over long periods of time. It can be used with irregularly-sampled movement data without the need to interpolate the locations, and can incorporate measurement error. We present the model formulation and describe a Bayesian estimation method to infer

hidden states and movement parameters in this framework. We investigate the performance of the method in a simulation study, and demonstrate that it can be used to recover estimates of the states and movement parameters, from irregular location data. We analyse a trajectory of grey seal (*Halichoerus grypus*) with a 2-state CTCRW model, and obtain posterior samples of the state-dependent movement parameters and of the unobserved state sequence. We explain how the movement parameters differ in the two states, and how they can be interpreted as measures of the animal's speed and movement persistence.

2 Model formulation

2.1 Continuous-time correlated random walk

The continuous-time correlated random walk (CTCRW) was introduced as a model of animal movement by Johnson et al. (2008). The underlying stochastic process was originally developed by Uhlenbeck and Ornstein (1930) to describe the movement of a particle under friction.

Denote $\mathbf{z}_t = (x_t, y_t)^\top$ the location of the animal at time t , and $\mathbf{v}_t = (v_t^x, v_t^y)^\top$ its velocity, linked by the equation

$$d\mathbf{z}_t = \mathbf{v}_t dt. \quad (1)$$

In the CTCRW model, the velocity of the animal is modelled by an Ornstein-Uhlenbeck process, defined as the solution of the stochastic differential equation

$$d\mathbf{v}_t = \beta(\boldsymbol{\gamma} - \mathbf{z}_t)dt + \sigma d\mathbf{w}_t, \quad (2)$$

where \mathbf{w}_t denotes a Wiener process, $\boldsymbol{\gamma} \in \mathbb{R}^2$ is the mean velocity, $\beta > 0$ measures the reversion of the velocity to its mean, and $\sigma > 0$ measures the spread of the velocity around its mean. In practice, the mean velocity parameter $\boldsymbol{\gamma}$ is often taken to be zero, corresponding to the case where there is no systematic drift in the animal's movement (although see Johnson et al., 2008, for an example of analysis with drift). Here, for simplicity, we consider that β and σ are scalar parameters, corresponding to the isotropic case, but they could be taken as matrices for a more general formulation (Blackwell, 2003; Gurarie et al., 2017). We will sometimes refer to the location process \mathbf{z}_t as an integrated Ornstein-Uhlenbeck process, to indicate that its derivative (with respect to time) is an Ornstein-Uhlenbeck process.

This formulation is very convenient because it is possible to derive the transition densities of the velocity process \mathbf{v}_t and of the position process \mathbf{z}_t analytically. In the following, we assume $\boldsymbol{\gamma} = \mathbf{0}$. Using Itô calculus,

it can be shown that, under Equations 1 and 2, we have

$$\mathbf{v}_{t+\delta} = e^{-\beta\delta}\mathbf{v}_t + \boldsymbol{\zeta}(\delta),$$

and

$$\mathbf{z}_{t+\delta} = \mathbf{z}_t + \left(\frac{1 - e^{-\beta\delta}}{\beta}\right)\mathbf{v}_t + \boldsymbol{\xi}(\delta),$$

for any time interval $\delta > 0$ and, for either dimension $c \in \{x, y\}$,

$$\zeta_c(\delta) \sim N\left[0, \frac{\sigma^2}{2\beta}(1 - e^{-2\beta\delta})\right], \quad (3)$$

$$\xi_c(\delta) \sim N\left[0, \left(\frac{\sigma}{\beta}\right)^2 \left(\delta + \frac{1 - e^{-2\beta\delta}}{2\beta} - \frac{2(1 - e^{-\beta\delta})}{\beta}\right)\right]. \quad (4)$$

The steps of the calculation are detailed in Appendix A.

Johnson et al. (2008) developed a method to estimate the movement parameters β and σ from observed telemetry data. The parameters can be linked to the speed and sinuosity of the animal's movement, and therefore used for the biological interpretation of tracking data. Indeed, Gurarie et al. (2017) presented an alternative parametrisation of the CTCRW, defined with $\tau = 1/\beta$ and $\nu = \sqrt{\pi}\sigma/(2\sqrt{\beta})$. In this formulation, the parameter $\tau > 0$ is the time interval over which the autocorrelation function of the velocity process decreases by a factor e , and it is sometimes called the “relaxation time” of the process (Gillespie, 1996). Larger values of τ (corresponding to smaller values of β) indicate longer-term persistence in the speed and direction of the animal's movement. The autocorrelation function of the velocity process decreases to 0.05 over a time interval of length 3τ , because $e^{-3} \approx 0.05$. For most practical purposes, \mathbf{v}_t and $\mathbf{v}_{t+3\tau}$ can therefore be regarded as approximately independent (Johnson et al., 2008). The parameter $\nu > 0$ is the long-term mean of the speed of movement of the animal. With this convenient formulation, the model offers a very useful framework to quantify movement characteristics of animals.

2.2 Multistate model

In this paper, we use the CTCRW as a building block for more complex and realistic movement models. Multistate models of animal movement have been developed to account for behavioural heterogeneity. In the most common formulation, a (discrete- or continuous-time) Markov process models switches between discrete “behavioural” states, on which depend the parameters of the movement process (Blackwell, 1997; Morales et al., 2004). Following this idea, we introduce a N -state continuous-time Markov process $(s_t)_{t \geq 0}$,

characterised by its infinitesimal generator matrix,

$$\mathbf{\Lambda} = \begin{pmatrix} -\lambda_1 & \lambda_{12} & \cdots & \lambda_{1N} \\ \lambda_{21} & -\lambda_2 & \cdots & \lambda_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{N1} & \lambda_{N2} & \cdots & -\lambda_N \end{pmatrix}, \quad (5)$$

where $\forall i \in 1, \dots, N$, $\lambda_i = \sum_{j \neq i} \lambda_{ij}$. At any time $t \geq 0$, the discrete state s_t takes one of N values $\{1, \dots, N\}$, typically used as proxies for the behavioural states of the animal (e.g. “foraging”, “exploring”).

The generator matrix is the continuous-time analogue of the transition probability matrix used in hidden Markov models, and its entries determine the state-switching dynamics. In particular, as a consequence of the Markov property, the dwell times in state i follow an exponential distribution with rate λ_i .

We now consider that the parameters of the CTCRW model (β and σ in Equation 2) are state-dependent, so that each behavioural state can be associated with a different type of movement. Using the notation introduced in Section 2.1, the multistate CTCRW model is defined by

$$\begin{cases} dz_t = \mathbf{v}_t dt, \\ d\mathbf{v}_t = -\beta_{s_t} \mathbf{v}_t dt + \sigma_{s_t} d\mathbf{w}_t. \end{cases}$$

This can be viewed as a higher-dimension continuous-time Markov process composed of a continuous component (the location and velocity processes) and a discrete component (the discrete state process), as described e.g. by Berman (1994). In the following, we develop a method of Bayesian inference for the multistate CTCRW model.

3 Bayesian inference

3.1 Likelihood evaluation with Kalman filter

We present a method to evaluate the likelihood of the multistate CTCRW model, given a reconstruction of the underlying state process. Johnson et al. (2008) implemented a Kalman filter to obtain the likelihood of movement trajectories in the single-state CTCRW model. With only minor changes, it can be extended to evaluate the likelihood of the multistate model, conditionally on the behavioural state sequence. The behaviours of the animal are typically not known, and will thus need to be imputed. One method of reconstructing the state sequence is to sample over the state process in a MCMC algorithm, which we

describe in Section 3.2.

We consider a data set of observed locations, augmented with the times of the reconstructed state transitions. The locations associated with the transitions are generally not observed, and they are thus treated as missing data. We denote by $\{\tilde{z}_1, \dots, \tilde{z}_n\}$ the augmented sequence of locations, possibly observed with measurement error, and $\{t_1, \dots, t_n\}$ the associated times. We denote by s_i the (imputed) behavioural state between t_i and t_{i+1} .

Following Johnson et al. (2008), the model can be written as a state-space model, where the observation process is the animal's observed location (possibly including measurement error), and the continuous state process is composed of both the true location and the velocity. We denote by $\boldsymbol{\omega}_i = (x_i, v_i^x, y_i, v_i^y)^\top$ the continuous state vector at time t_i , $\Delta_i = t_{i+1} - t_i$ the time intervals, and $\boldsymbol{\xi}_i = \boldsymbol{\xi}(\Delta_i)$ and $\boldsymbol{\zeta}_i = \boldsymbol{\zeta}(\Delta_i)$ the stochastic terms of the transition densities for the location and the velocity, respectively (Equations 3 and 4, substituting β_{s_i} for β and σ_{s_i} for σ).

Both s_i and $\boldsymbol{\omega}_i$ are referred to as “state”, in state-switching models and state-space models, respectively. Here, we combine both and, to avoid confusion, we will refer to $\boldsymbol{\omega}_i$ as the “continuous” state of the process (as opposed to the “discrete” behavioural state s_i).

The observation equation of the CTCRW is

$$\tilde{z}_i = \mathbf{Z}\boldsymbol{\omega}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim N(\mathbf{0}, \mathbf{H}_i),$$

where \mathbf{H}_i is the 2×2 measurement error covariance matrix, and

$$\mathbf{Z} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

That is, the observed location \tilde{z}_i is the sum of the true location $\mathbf{z}_i = (x_i, y_i)$ and an error term $\boldsymbol{\varepsilon}_i$. Using block matrix notation, the continuous state equation is

$$\boldsymbol{\omega}_{i+1} = \begin{pmatrix} \mathbf{T}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_i \end{pmatrix} \boldsymbol{\omega}_i + \boldsymbol{\eta}_i, \quad \boldsymbol{\eta}_i \sim N \left[\mathbf{0}, \begin{pmatrix} \mathbf{Q}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_i \end{pmatrix} \right]$$

where

$$\mathbf{T}_i = \begin{pmatrix} 1 & (1 - e^{-\beta_{s_i}\Delta_i})/\beta_{s_i} \\ 0 & e^{-\beta_{s_i}\Delta_i} \end{pmatrix}, \quad \mathbf{Q}_i = \begin{pmatrix} \text{Var}(\boldsymbol{\xi}_{ci}) & \text{Cov}(\boldsymbol{\xi}_{ci}, \boldsymbol{\zeta}_{ci}) \\ \text{Cov}(\boldsymbol{\xi}_{ci}, \boldsymbol{\zeta}_{ci}) & \text{Var}(\boldsymbol{\zeta}_{ci}) \end{pmatrix}.$$

The variances are given in Equation 3 and 4. The derivation of the covariance is given in Appendix A,

and yields

$$\text{Cov}(\xi_{ci}, \zeta_{ci}) = \frac{\sigma_{s_i}^2}{2\beta_{s_i}^2} (1 - 2e^{-\beta_{s_i}\Delta_i} + e^{-2\beta_{s_i}\Delta_i}).$$

Under this state-space model formulation, the Kalman filter can be used to obtain the log-likelihood of observed locations $\{\tilde{z}_1, \tilde{z}_2, \dots, \tilde{z}_n\}$. Although it is not needed to fit the model, we can also implement the so-called Kalman smoother algorithm. The Kalman smoother provides an estimate $\hat{\omega}_i$ of the continuous state at each time step, as well as the covariance matrix $\hat{\Sigma}_i$ of the estimate, conditional on all observations (Johnson et al., 2008). As in Johnson et al. (2008), the Kalman filter is here used to integrate over both measurement error and unobserved velocities simultaneously. Appendix B gives the Kalman filter and smoother equations for the model, and the expression of the log-likelihood. For more details on state-space models and the Kalman filter, see e.g. Durbin and Koopman (2012).

3.2 MCMC algorithm

We describe a method to infer the parameters of the multistate CTCRW model introduced in Section 2.2. In Section 3.1, we explained how the Kalman filter can be used to compute the likelihood of the model, conditionally on the behavioural state sequence. However, in most applications, the behaviours of the animal are not observed. We propose to estimate the unobserved states by sampling over all possible sequences in a Markov chain Monte Carlo (MCMC) algorithm, following the Metropolis-within-Gibbs approach introduced by Blackwell (2003). It relies on the successive updates of the three components of the model: the reconstructed behavioural state process, the movement parameters, and the transition rates.

We denote by $p(\tilde{z}|\boldsymbol{\theta}, \mathcal{S})$ the likelihood of a sequence of observed locations \tilde{z} , given the movement parameters $\boldsymbol{\theta}$ and the reconstructed state sequence \mathcal{S} , as obtained from the Kalman filter presented in Section 3.1. We denote by $\mathbf{\Lambda}$ the generator matrix of the behaviour process, as defined in Equation 5. We initialise the state sequence to $\mathcal{S}^{(0)}$, the movement parameters to $\boldsymbol{\theta}^{(0)}$, and the generator matrix to $\mathbf{\Lambda}^{(0)}$. Then, for K iterations ($k = 1, \dots, K$), we run the three following steps alternately, to sample $\mathcal{S}^{(k)}$, $\boldsymbol{\theta}^{(k)}$, and $\mathbf{\Lambda}^{(k)}$.

Update of the behavioural state process The evaluation of the likelihood of the model, described in Section 3.1, is conditional on the sequence of underlying states. In practice, the states are generally not observed, such that we need to impute them. The sequence of states \mathcal{S} is composed of the times of the state transitions, and the values of the states. At each iteration, an updated state sequence \mathcal{S}^* is proposed as follows. We choose an interval $[t_a, t_b]$, where a and b are two integers such that $t_1 \leq t_a < t_b \leq t_n$. We simulate the state process s_t from t_a to t_b , conditional on s_{t_a} and s_{t_b} , e.g. using the endpoint-conditioned continuous-time Markov process simulation methods from Hobolth and Stone (2009). The proposed sequence

of states \mathcal{S}^* remains identical to $\mathcal{S}^{(k-1)}$ outside $[t_a, t_b]$. The acceptance ratio for \mathcal{S}^* is

$$r = \min \left\{ 1, \frac{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^{(k-1)}, \mathcal{S}^*)}{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^{(k-1)}, \mathcal{S}^{(k-1)})} \right\}.$$

The proposed state process reconstruction is accepted with probability r . The length of the interval $[t_a, t_b]$ over which the state sequence is updated is a tuning parameter of the sampler, and updates over longer intervals are generally less likely to be accepted. Note that the state sequence is generated from a continuous-time Markov process, and transitions can therefore occur at any point in (continuous) time. In particular, they are not constrained to happen at the times of the observations.

Update of the movement parameters Denote $\boldsymbol{\theta}$ the vector of parameters of the movement process, i.e. $\boldsymbol{\theta} = (\beta_1, \dots, \beta_N, \sigma_1, \dots, \sigma_N)$ for a N -state model. We use a Metropolis-Hastings step to update the movement parameters. At iteration k , we propose new movement parameters $\boldsymbol{\theta}^*$, from a proposal density $q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(k-1)})$, and the acceptance ratio is

$$r = \min \left\{ 1, \frac{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^*, \mathcal{S}^{(k)})p(\boldsymbol{\theta}^*)q(\boldsymbol{\theta}^{(k-1)}|\boldsymbol{\theta}^*)}{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^{(k-1)}, \mathcal{S}^{(k)})p(\boldsymbol{\theta}^{(k-1)})q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(k-1)})} \right\},$$

where $p(\boldsymbol{\theta})$ denotes the prior distribution on the movement parameters $\boldsymbol{\theta}$. The parameters are updated to $\boldsymbol{\theta}^*$ with probability r . Note that, if the proposal distribution is symmetric such that $\forall \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, q(\boldsymbol{\theta}_1|\boldsymbol{\theta}_2) = q(\boldsymbol{\theta}_2|\boldsymbol{\theta}_1)$, then r simplifies to

$$r = \min \left\{ 1, \frac{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^*, \mathcal{S}^{(k)})p(\boldsymbol{\theta}^*)}{p(\tilde{\mathbf{z}}|\boldsymbol{\theta}^{(k-1)}, \mathcal{S}^{(k)})p(\boldsymbol{\theta}^{(k-1)})} \right\}.$$

In practice, a standard choice is to use a multivariate normal proposal distribution on the working scale of the parameters (in this case, on the log scale). Its variance can be tuned to obtain different acceptance rates, and covariance structure can be added to explore the parameter space more efficiently.

Update of the transition rates Following Blackwell (2003), using conjugate priors, the transition rates can be directly sampled from their posterior distribution, which is known conditionally on the reconstructed state sequence $\mathcal{S}^{(k)}$. We find it convenient to parametrise the generator matrix as

$$\boldsymbol{\Lambda} = \begin{pmatrix} -\lambda_1 & \lambda_1 p_{12} & \cdots & \lambda_1 p_{1N} \\ \lambda_2 p_{21} & -\lambda_2 & \cdots & \lambda_2 p_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_N p_{N1} & \lambda_N p_{N2} & \cdots & -\lambda_N \end{pmatrix}$$

where $\lambda_i > 0$ is the rate of transition out of state i , and the $p_{ij} \in [0, 1]$ are the transition probabilities out of state i . For each state i , they satisfy $\sum_j p_{ij} = 1$. The transition rates and the transition probabilities can be sampled separately.

For each $i \in \{1, \dots, N\}$, we denote n_i the number of time intervals spent in state i , and $(d_i^{(1)}, d_i^{(2)}, \dots, d_i^{(n_i)})$ their lengths. These dwell times are exponentially distributed with rate λ_i . The conjugate prior of the exponential distribution is the gamma distribution such that, with the prior

$$\lambda_i \sim \text{gamma}(\alpha_1, \alpha_2),$$

the transition rates are sampled from the posterior distribution

$$\lambda_i^{(k)} | \mathcal{S}^{(k)} \sim \text{gamma} \left(\alpha_1 + n_i, \alpha_2 + \sum_{j=1}^{n_i} d_i^{(j)} \right).$$

For $i \in \{1, \dots, N\}$ and $j \in \{1, \dots, N\}$ such that $i \neq j$, we denote n_{ij} the number of transitions from state i to state j , and $n_i = \sum_j n_{ij}$ the number of transitions out of state i . Then,

$$n_{i1}, n_{i2}, \dots, n_{iN} \sim \text{multinom}(n_i, p_{i1}, p_{i2}, \dots, p_{iN}).$$

The conjugate prior of the multinomial distribution is the Dirichlet distribution such that, with the prior

$$\mathbf{p}_i \sim \text{Dir}(\kappa_{i1}, \kappa_{i2}, \dots, \kappa_{iN}),$$

the posterior distribution of the transition probabilities is

$$\mathbf{p}_i^{(k)} | \mathcal{S}^{(k)} \sim \text{Dir}(\kappa_{i1} + n_{i1}, \kappa_{i2} + n_{i2}, \dots, \kappa_{iN} + n_{iN}).$$

where $\mathbf{p}_i = (p_{i1}, p_{i2}, \dots, p_{iN})$ is the vector of transition probabilities out of state i .

4 Motivation

The method described in Sections 2 and 3 is widely applicable to many types of telemetry data. The flexible formulation of the state-switching CTCRW model can accommodate data with the following features: (1) movement persistence, (2) behavioural heterogeneity, (3) irregular time intervals, and (4) measurement error.

High-resolution movement data often display strong movement persistence, i.e. autocorrelation in the

Accepted Article

speed and direction of movement. The recent development of discrete- and continuous-time correlated random walks arises from the need to account for this autocorrelation (Morales et al., 2004; Jonsen et al., 2005; Johnson et al., 2008). In the CTCRW model, this persistence is modelled by an autocorrelated velocity process. In an alternative continuous-time formulation, proposed by Dunn and Gipson (1977) and Blackwell (1997), the location of an animal – rather than its velocity – is modelled with an Ornstein-Uhlenbeck process. That model does not describe smooth movement trajectories, and it is therefore less applicable to modern tracking data. The more recent approach of Parton and Blackwell (2017) addresses this persistence, but at substantial computational cost. The collection of high-resolution telemetry data will keep increasing in coming years, and rapid fitting of models of persistent movement will be compelling for the analysis of data from many taxa.

State-switching models provide a useful framework to account for behavioural heterogeneity in animal movement (Blackwell, 1997; Morales et al., 2004). Movement tracks are routinely collected over periods of several weeks or several months. Over such long periods of time, an animal may display different types of movement, which we capture with the underlying “behavioural” state process. The combination of movement persistence and behavioural states in the state-switching CTCRW model allows for a wide range of realistic movement patterns.

It is common for tracking data to be irregular in time. It can be accidental, for example for marine mammals which can only be monitored when they surface, or it can be by design, for example if the sampling frequency depends on the level of activity of the animal (Brown et al., 2012). Most currently-used movement models are formulated in discrete time, and they cannot be used to analyse irregular data, because their parameters are defined for a fixed (regular) time scale. Continuous-time models, such as the state-switching CTCRW, can deal with location data collected irregularly, and they can also be used to compare studies based on different sampling intervals. We illustrate this important difference between the scaling properties of discrete-time and continuous-time models using simulations in Section 5.2.

The presence of measurement error in telemetry data has been an important issue, in particular for radio-tracking and Argos devices (e.g. Springer, 1979; Hays et al., 2001). Many studies use a two-stage approach, in which the track is first filtered to estimate the true locations, and the filtered track is then analysed (e.g. to estimate behavioural states, or habitat selection). In two-stage methods, the uncertainty associated with the measurement error is generally not propagated to the final results. It is preferable to model the movement process and the observation process jointly in a state-space model framework, like the one we described for the state-switching CTCRW model (Jonsen et al., 2003; Patterson et al., 2008). In the method described in this paper, the final inferences into the movement and behaviour of an animal include the uncertainty of measurement error. We illustrate the application of the method to noisy observations in

a simulation study in Appendix C of the supplementary material.

In this paper, we focus on the estimation of the hidden state sequence and of movement parameters, for the state-switching CTCRW model. However, this modelling framework offers other possibilities. In particular, it can be used to predict the location of the animal between observations (with associated uncertainty). This has been one of the main applications of the single-state CTCRW, as implemented in the R package *crawl* (Johnson and London, 2018), to obtain smooth estimates of a trajectory from noisy and irregular telemetry data (e.g. Robinson et al., 2012; Baylis et al., 2015; Rode et al., 2015). Predictions from the state-switching CTCRW model account for behavioural heterogeneity in the movement patterns, and are therefore more susceptible to provide realistic location and uncertainty estimates. More generally, all functionalities of the single-state CTCRW can be implemented in the state-switching case, based on the Kalman filter and smoother algorithms presented in Section 3.1 (and detailed in Appendix B).

5 Simulation study

We used simulations to investigate the performance of the MCMC algorithm described in Section 3 to infer the hidden state sequence and the movement parameters from irregular movement data, and to compare this method to analogous discrete-time approaches. The R code used in the simulation study is provided in the supplementary material.

5.1 Estimation from irregular data

We simulated 10,000 locations from a 2-state model at a fine time scale (every 0.1 time unit), and thinned them by keeping 10% of the points at random (i.e. 1000 irregularly-spaced locations), to emulate real movement data. The time intervals in the resulting data set ranged between 0.1 and 8 time units.

The movement parameters and switching rates of the simulated process were chosen as

$$(\beta_1, \beta_2) = (1, 0.3), \quad (\sigma_1, \sigma_2) = (1, 3), \quad \mathbf{\Lambda} = \begin{pmatrix} -0.03 & 0.03 \\ 0.03 & -0.03 \end{pmatrix}.$$

In state 1, the variance was smaller and the reversion to the mean larger, which resulted in slower and more sinuous movement (perhaps analogous to “area-restricted search” behaviour). State 2 corresponded to faster and more directed movement (analogous to “transit”). This can be seen from the time scale of autocorrelation τ and mean speed ν for each state, as defined in Section 2.1. In this simulation, we have $(\tau_1, \tau_2) = (1, 3.33)$ and $(\nu_1, \nu_2) = (0.89, 4.85)$, i.e. more persistent and faster movement in State 2. The

transition rates were chosen such that the process would on average stay about 30 time units in a state before switching to the other state. The simulated track (after thinning) is shown in Figure 1.




Figure 1: Track simulated from a 2-state CTCRW model.

We initialised the reconstructed state sequence by classifying each observation randomly as being in state 1 or state 2, with probability 0.5 each. At each iteration, the state process was updated over a randomly-selected interval of five time steps. We used independent normal proposal distributions (on the working log scale) to update the movement parameters. The proposal variances were tuned based on initial test runs, to obtain near-optimal acceptance rates. We chose normal prior distributions on the working scale for the movement parameters, centred on the true values of the parameters, and with large variances.

We ran 10^5 MCMC iterations, which took around 20 min on a 2GHz i5 CPU, and discarded the first 5×10^4 as burn-in. Figure 2 shows the posterior probabilities of being in state 2 at the times of the observations, to compare with the “true” simulated state sequence. For each $i = 1, \dots, n$, we calculated the posterior probability as the proportion of reconstructions of the state process in which \mathbf{x}_i was classified in state 2. (It would have been equivalent to consider the posterior probability of being in state 1; the two probabilities sum to 1.) We considered that a step was misclassified if the posterior probability of being in the true state was less than 0.5. The states were correctly estimated in the vast majority of time steps, with only 2% of misclassified steps. The posterior probability of being in the true state was less than 0.9 for about 4.5% of the locations.

Figure 3 displays posterior samples for the state-dependent movement parameters, β_1 , β_2 , σ_1 , and σ_2 , as well as the true parameter values used in the simulation. The posterior distributions seem to appropriately




Figure 2: Posterior probabilities of being in state 2 at the times of the observations. The true (simulated) states are shown by the colours (red: state 1, blue: state 2).

estimate all movement parameters. Although there appears to be some possible bias, replications (not shown here) reveal that it is only due to randomness, and not consistent across simulations.




Figure 3: Posterior samples of the movement parameters in state 1 (left) and in state 2 (right), in the simulation study. The black dots are the true values of the parameters, used in the simulation. The black lines show contours of a kernel density estimate of the posterior samples. The samples are thinned to every tenth value, for visualisation purposes.

We were able to recover the values of the state process and of the state-dependent movement parameters from a simulated track thinned to irregular intervals. This demonstrates the ability of the method to work across temporal scales, and to cope with irregular sampling.

In Appendix C of the supplementary material, we repeat this simulation experiment, after introducing some “observation error” to the simulated locations. The movement parameters and state process are estimated well in that scenario, although there is some indication that large measurement error (compared

to the scale of movement) leads to larger uncertainty in the estimates, and slower mixing of the MCMC algorithm. There, we also show how the Kalman smoother can be used to estimate the true location process at a fine time scale, from irregular and noisy observations.

5.2 Comparison to discrete-time model

We simulated two tracks from the state-switching CTCRW model, at a time resolution of 0.1, with the same parameters used in the simulations of Section 5.1. We then thinned them to obtain two data sets: (1) a track of 1000 locations, at a regular time resolution of $\Delta = 0.5$, and (2) a track of 1000 locations, at a regular time resolution of $\Delta = 5$. Both tracks arise from the same underlying process (the true state-switching CTCRW), and they emulate data sets observed at different time resolutions. We fitted a discrete-time state-switching model, and the state-switching CTCRW model, to the two data sets, to illustrate the differences in their scaling properties. In this experiment, we thinned the tracks at regular time intervals (rather than irregularly, at random) because discrete-time models cannot be applied to irregular location data.

We fitted a 2-state discrete-time hidden Markov model to each data set, using the R package `moveHMM` v1.6 (Michelot et al., 2016). The standard hidden Markov model, described by Patterson et al. (2009) and Langrock et al. (2012), is formulated in terms of step lengths (distances between successive locations) and turning angles (angles between successive locations). We modelled the step lengths with gamma distributions, and the turning angles with von Mises distributions. There were six estimated movement parameters: the step length means (μ_1, μ_2) , the step length standard deviations (s_1, s_2) , and the turning angle concentrations (κ_1, κ_2) . The mean of the distribution of turning angles was fixed to zero, which is a standard choice to model movement persistence. See Michelot et al. (2016) for more detail about the hidden Markov model formulation implemented in `moveHMM`. The parameters of the step length distribution are related to the speed of movement, and the parameters of the turning angle distribution are related to the directional persistence. There were two estimated parameters of the state process: the transition probabilities γ_{12} and γ_{21} , where γ_{ij} is the probability of a transition from state i to state j over one time interval.

We also fitted the state-switching CTCRW model to each thinned data set separately. Like in the simulation study of Section 5, we ran 10^5 MCMC iterations, and discarded the first 5×10^4 as burn-in. We estimated four movement parameters: $(\beta_1, \beta_2, \sigma_1, \sigma_2)$. There were two estimated parameters of the state process: the transition rates λ_1 and λ_2 , as defined in Equation 5. The parameter estimates for the hidden Markov models and for the state-switching CTCRW models are given in Table 1.

The parameters of the continuous-time model (state-switching CTCRW) were very similar in both analyses. Although the two thinned data sets have very different sampling intervals ($\Delta = 0.5$ and $\Delta = 5$), the

Parameter	True value	$\Delta = 0.5$		$\Delta = 5$		
		Estimate	95% CI	Estimate	95% CI	
Discrete time						
Movement parameters	μ_1	—	0.52	(0.48,0.56)	2.83	(2.66,3.02)
	μ_2	—	2.60	(2.50,2.71)	18.95	(17.98,19.98)
	s_1	—	0.33	(0.30,0.37)	1.69	(1.53,1.87)
	s_2	—	1.01	(0.94,1.09)	10.07	(9.27,10.94)
	κ_1	—	1.38	(1.22,1.55)	0.21	(0.12,0.39)
	κ_2	—	7.03	(5.94,8.32)	0.75	(0.62,0.90)
Transition probabilities	γ_{12}	—	0.063	(0.045,0.089)	0.132	(0.104,0.167)
	γ_{21}	—	0.068	(0.048,0.097)	0.134	(0.104,0.172)
Continuous time						
Movement parameters	β_1	1	1.12	(0.95,1.30)	0.93	(0.69,1.23)
	β_2	0.3	0.32	(0.25,0.39)	0.27	(0.23,0.32)
	σ_1	1	1.04	(0.98,1.11)	0.95	(0.74,1.22)
	σ_2	3	3.03	(2.90,3.17)	2.84	(2.58,3.14)
Transition rates	λ_1	0.03	0.026	(0.013,0.045)	0.035	(0.026,0.044)
	λ_2	0.03	0.034	(0.016,0.058)	0.033	(0.025,0.042)

Table 1: Parameter estimates in the comparison of hidden Markov models (“discrete time”) and state-switching CTCRW models (“continuous time”) over two different temporal scales. For the hidden Markov model parameters, maximum likelihood estimates and 95% confidence intervals are shown (obtained with the package *moveHMM*). For the CTCRW parameters, mean posterior estimates and 95% credible intervals are shown.

parameters of the CTCRW are independent of the time intervals of observations.

However, the parameters of the discrete-time model (hidden Markov model) were very different in the two analyses. This should be expected, because the model is formulated in terms of scale-dependent metrics: the step lengths and turning angles. The step lengths increase with the time interval of observation, and the estimates reflect this. Indeed, the mean and standard deviation of the distribution of step lengths were smaller in the analysis with $\Delta = 0.5$ than when $\Delta = 5$. It should be noted that there is no clear scaling rule for those parameters. In particular, the mean step length is generally *not* 10 times longer over $\Delta = 5$ than over $\Delta = 0.5$. This highlights a problem that often arises in analyses of irregular tracking data with discrete-time models. A method that has been proposed to deal with irregular data is to derive a “movement rate” for each time step, by dividing the step length by the time interval. Then, it is assumed that the movement rates do not depend on the time interval, and represent a measure of the animal’s speed of movement. However, we can see from this simulated example that the movement rates do, in fact, depend on the time intervals of observation. Indeed, the mean movement rate can be obtained as the mean step length divided by the time interval. In the first data set, the mean movement rates in the two states are $0.52/0.5 = 1.04$ and $2.60/0.5 = 5.2$. In the second data set, the mean movement rates are $2.83/5 = 0.57$ and $18.95/5 = 3.79$. In general, the sparser the data, the more the speed of movement is underestimated by this method. A similar

problem arises when linear interpolation is used to obtain locations at regular intervals in time. Linear interpolation draws straight lines between observed locations, and will therefore tend to underestimate the speed of movement, and overestimate the persistence in direction.

The turning angles tend to be less concentrated around 0 as the time interval of observation increases, because movement persistence is less clearly visible at a coarse time resolution. This can be seen in the parameter estimates, as κ_1 and κ_2 are much larger in the analysis with $\Delta = 0.5$ than in the analysis with $\Delta = 5$. However, similarly to the mean step length, it is not clear how the concentration parameter scales with the time interval of observations.

This dependence on the time scale of observations must be taken into account in the interpretation of the parameters of discrete-time movement models. For example, it is difficult to interpret the mean step length as a measure of the animal's speed, because its definition is tied to the time interval of observation. It will tend to increase with the time interval, but generally not linearly (because animals do not move in straight lines between observed locations). Crucially, the results of different discrete-time analyses cannot be compared if the time interval of observation is different, as illustrated in this simulation experiment.

The same scale dependence arises for the parameters of the state process. The transition probabilities of the hidden Markov models are defined over a given (fixed) time interval, and can only be interpreted over that time interval. For example, in the first data set, we estimated that the probability of a transition from state 1 to state 2 was $\hat{\gamma}_{12} = 0.06$. This means that, if the animal is in state 1, there is a probability 0.06 of switching to state 2 over each time step of length $\Delta = 0.5$. On the other hand, the transition rates of the continuous-time model are defined independently of any specific time interval. Regardless of the time interval of observation, the estimated transition rate can be interpreted as the “mean number of transitions per hour”.

The objective of this simulation study is not to compare how well the discrete-time and continuous-time models captured the true underlying process, or how well they recovered the true parameter values. This would be unfair, because the true process used to simulate the tracks is the state-switching CTCRW. Nevertheless, we believe that the results presented here are useful to illustrate the fundamental difference in the formulations of discrete-time and continuous-time models. In particular, a considerable practical advantage of continuous-time approaches is that the estimated parameters, and therefore the interpretation and inference, do not depend on the time interval of observation.

6 Grey seal case study

We illustrate the use of the method described in Section 3 for the analysis of a grey seal (*Halichoerus grypus*) movement track. We considered a trajectory of 2535 observations, collected in the North Sea between April and December 2008 (McConnell, 2019), and previously described by Russell et al. (2015). The base sampling frequency was of one location every 30 minutes, but many fixes were missed, and the resulting time grid was highly irregular ($P_{0.025} = 27$ min, $P_{0.975} = 10$ hours). Note that the CTCRW model describes movement on a plane, and thus requires that the longitude-latitude locations be projected to UTM coordinates for the analysis.

We considered a 2-state CTCRW model, with four movement parameters to estimate: $\beta_1, \beta_2, \sigma_1, \sigma_2$. Similarly to the simulation study, we initialised the state reconstruction to a random sequence of 1s and 2s (with probability 0.5 each). We used independent normal proposal and prior distributions on the working scale of the movement parameters. We selected the proposal variances based on test runs, and used weakly informative prior distributions. We ran 2 million MCMC iterations, discarding the first half as burn-in, which took about 14 hours on a 2GHz i5 CPU. We only saved every 100th reconstructed state sequence, because of memory limitations.

Figure 4 shows a map of the track, coloured by posterior state probabilities, and Figure 5 shows posterior samples for the four movement parameters $(\beta_1, \beta_2, \sigma_1, \sigma_2)$. State 2 captured very directed movements, corresponding to periods of transit between areas of interest, and state 1 captured more tortuous phases of the track. This can be seen in Figure 5: the posterior distribution of β_1 covers much larger values than that of β_2 (posterior means of 1.73 and 0.06, respectively), indicating stronger reversion to the mean in state 1, and thus less movement persistence. There were no signs of label switching in the posterior samples; if there were, a straightforward solution would be to constrain (β_1, β_2) and (σ_1, σ_2) to be ordered (McClintock et al., 2014). We derived effective sample sizes (ESS) for the posterior samples of the movement parameters, with the R package coda (Plummer et al., 2006). We found $ESS(\beta_1) = 1493$, $ESS(\beta_2) = 2413$, $ESS(\sigma_1) = 3150$, and $ESS(\sigma_2) = 1234$. There was no indication of mixing issues in the MCMC algorithm.

We transformed the posterior samples of movement parameters to obtain estimates of τ and ν in both states (as defined in Section 2.1). In the following, we report posterior mean estimates, and histograms of the posterior samples for the transformed parameters can be found in Appendix D. The posterior means were $\hat{\tau}_1 = 0.55$ h and $\hat{\tau}_2 = 17.21$ h for the time scales of autocorrelation, and $\hat{\nu}_1 = 2.63$ km/h and $\hat{\nu}_2 = 2.90$ km/h for the mean speeds. This indicates that the two states are very similar in terms of the speed of movement, but that the autocorrelation function of the velocity drops much faster in state 1 than in state 2.

The posterior samples of the transition rates can be used to derive mean dwell times in each state, and

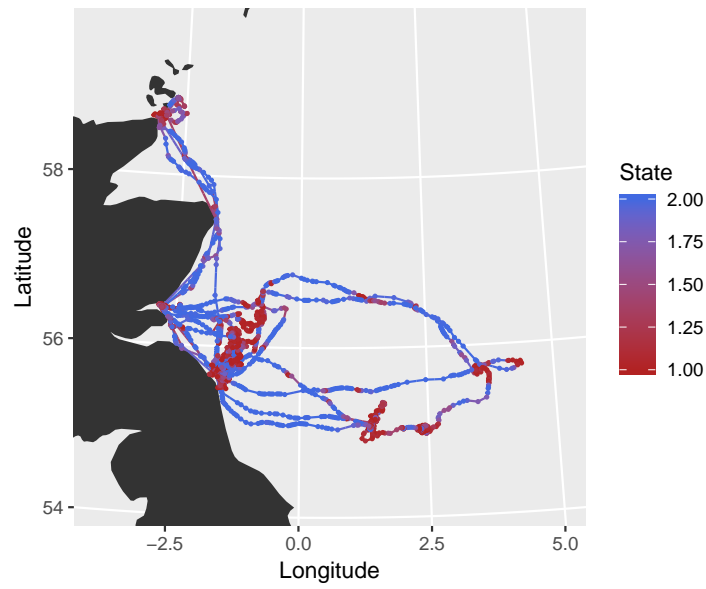


Figure 4: Grey seal track, off the East coast of Great Britain, coloured by posterior state probabilities.




Figure 5: Posterior samples of the movement parameters in state 1 (left) and in state 2 (right), in the grey seal case study. The black lines show contours of a kernel density estimate of the posterior samples. The samples are thinned to every 100th value, for visualisation purposes.

long-term activity budgets. The dwell times in state i follow an exponential distribution with rate λ_i (the rate of transition out of state i). The mean dwell time can thus be derived as $d_i = 1/\lambda_i$. In the grey seal analysis, the posterior means for the mean dwell time were $\hat{d}_1 = 7.5\text{h}$ in state 1, and $\hat{d}_2 = 7.8\text{h}$ in state 2, indicating similar dwell times in both states. Activity budgets refer to the proportion of time spent by an animal in each of its behavioural states (Russell et al., 2015). In a time-homogeneous state-switching model, an estimate of the long-term activity budget can be calculated as the stationary distribution of the underlying Markov process. The stationary distribution of a N -state Markov process is the vector $\boldsymbol{\pi} = (\pi_1, \dots, \pi_N)$ which satisfies $\boldsymbol{\pi}\boldsymbol{\Lambda} = \mathbf{0}$, subject to the constraint $\sum_{i=1}^N \pi_i = 1$, where $\boldsymbol{\Lambda}$ is the generator matrix defined in Equation 5. In the 2-state case, solving the equation yields $\pi_1 = \lambda_2/(\lambda_1 + \lambda_2)$ and $\pi_2 = \lambda_1/(\lambda_1 + \lambda_2)$. The posterior mean estimate for the stationary distribution was $(\hat{\pi}_1, \hat{\pi}_2) = (0.48, 0.52)$, i.e. the seal will tend to spend roughly the same proportion of time in both states, in the long term. Histograms of posterior draws for the dwell times and stationary distribution are displayed in Appendix D.

The Kalman filter and smoother recursions given in Appendix B can be used to compute estimated velocities at the times of the observations. The velocities obtained with the mean posterior movement parameters are displayed in Figure 6, and split by posterior state estimates. The strong movement autocorrelation in state 2 can best be seen in the outer rim of the plot, where the velocity sometimes persists with little variation over many time steps. Interestingly, a cluster of very small velocities (close to zero) were also classified in state 2. This is because, as seen in the estimates of the movement parameters, the main difference between the two states are not in the speed, but in the velocity persistence. State 2 therefore captures both fast persistent and slow persistent movements of the seal. On the other hand, state 1 captures less persistent movement, characterised by a weaker autocorrelation in the velocity process. The Kalman algorithm can also provide estimates of the locations (and possibly velocities) of the animal – and associated standard errors – on any time grid, e.g. on a finer time grid than that of the observations.

7 Discussion

We presented a Bayesian framework to infer discrete behavioural states and movement parameters from a multistate continuous-time model of animal movement. The continuous-time formulation can accommodate irregular time intervals, and is consistent across temporal scales. The conditional likelihood of the model, used in the MCMC algorithm, is implemented using the Kalman filter, making it relatively fast and allowing for the inclusion of measurement error.

The MCMC algorithm of Section 3.2 closely resembles that developed by Blackwell (1997, 2003). However, they modelled the location of an animal (rather than its velocity) with an Ornstein-Uhlenbeck process, which




Figure 6: Predicted velocities from the grey seal example, obtained with the mean posterior movement parameter estimates, for time steps classified in state 1 (left) and in state 2 (right). The grey segments link consecutive velocities that were classified in the same state.

could not capture strong movement persistence. Here, we adapted the approach for the CTCRW process, to model autocorrelation in the velocity of an animal, i.e. persistence in the speed and direction of movement (which is often present in high-frequency telemetry data). However, the CTCRW may not be adequate to analyse movement data that do not display this type of persistence. In particular, data collected on a coarse time grid may not exhibit autocorrelated velocities, and could instead be analysed with the state-switching Ornstein-Uhlenbeck model of Blackwell (1997, 2003).

The inferential approach introduced in this paper could in principle be used to implement a state-switching version of the OUF model described by Fleming et al. (2014). The OUF process is a generalisation of the CTCRW used in this paper, and of the Ornstein-Uhlenbeck location process used e.g. by Blackwell (1997, 2003). It features both persistence in velocity and long-term attraction towards a point in space, making it a very flexible model of animal movement. Like the CTCRW, it can be written as a state-space model, and the Kalman filter can be used to derive the likelihood of the model (Fleming et al., 2017). The MCMC algorithm described in Section 3.2 could then be used to fit a multistate OUF model to animal movement. However, the OUF process has five parameters (against 2 only for the CTCRW), which could make estimation more challenging. More generally, this methodology could be applied to a model switching between processes with different formulations (e.g. a 2-state model switching between a CTCRW and OUF process). These complex multistate models could for example capture the structured trips of central place foragers (similarly to the discrete-time models of Michelot et al., 2017).

Analyses of animal movement and behaviour often combine telemetry and environmental data. In state-switching models, the effect of environmental (or other) covariates on the transition probabilities is of particular interest, and is used to uncover the drivers of animal behavioural and movement decisions (Patterson et al., 2009; Bestley et al., 2012; Blackwell et al., 2016). Blackwell et al. (2016) described a method of inference for the state-switching Ornstein-Uhlenbeck movement model. They allow the transition rates to be functions of spatial covariates (i.e. that can be evaluated at any point of the study region), or to be functions of the time of day (to analyse circadian cycles in the behaviour of an animal). The MCMC algorithm of Section 3.2 could be extended, following Blackwell et al. (2016), to allow for the inclusion of covariates in the state-switching dynamics.

Although we refer to the states of the Markov process as “behavioural states”, it is important to note that they really are *statistical* states, that capture the temporal autocorrelation in the velocity process. They should be interpreted with caution, and may not exactly correspond to separate behaviours (Patterson et al., 2017). In particular, there can be greater uncertainty in the partitioning of a track if the states are not very distinct, i.e. if they do not clearly differ in terms of the animal’s velocity process (Beyer et al., 2013). In the simulation study and grey seal case study, we focused on the 2-state model formulation, because the interpretation becomes more difficult in models with more states. Pohle et al. (2017) discussed this problem in hidden Markov models, which are the discrete-time analogue of the model presented in this work. A possible solution is to use auxiliary data, if available, to identify behavioural states. For example, if the behaviour of the animal is known for some of the observations, the corresponding states can be fixed throughout the algorithm, in a semi-supervised framework (Leos-Barajas et al., 2017). We could also include observation variables to the state-space model, in addition to the locations, to inform the behavioural states. For example, information about vertical movement has been used to identify behaviours in marine mammals (DeRuiter et al., 2017; McClintock et al., 2017). In our framework, the additional observation variables would have to be modelled with a normal transition density, to be integrated into the Kalman filter likelihood computations.

The computational cost of the method presented in this paper greatly depends on the number of MCMC iterations needed to obtain reliable estimates. As for any MCMC algorithm, convergence should be checked, for example using trace plots of the posterior samples (Gelman et al., 2013). We can improve the mixing speed of the algorithm with the choice of the tuning parameters, i.e. the variances of the proposal distributions for the movement parameters, and the length of the time interval over which the state sequence should be updated at each iteration. In the applications of Sections 5 and 6, we chose the tuning parameters to obtain acceptance rates that are close to the optimal value (i.e. around 23%; see Roberts et al., 1997). As in other frameworks, the parameters of the model may not all be identifiable if the behavioural states are

not clearly distinct, or if the measurement error is large compared with the scale of the movement. The simulation study shown in Appendix C suggests that such issues can be detected by monitoring the mixing speed of the MCMC algorithm, e.g. with the effective sample size of posterior samples. Slow mixing (and large uncertainty) indicates that the corresponding parameter estimates should be interpreted with caution.

Acknowledgements

TM was supported by the Centre for Advanced Biological Modelling at the University of Sheffield, funded by the Leverhulme Trust, award number DS-2014-081. We thank Debbie Russell and Esther Jones for very helpful advice about the grey seal data set. The grey seal telemetry data were provided by the Sea Mammal Research Unit (SMRU), University of St Andrews; their collection was conducted under UK Home Office Licence (60/3303) and supported by funding from the Natural Environment Research Council to SMRU. We are grateful to Guillaume Péron, Ian Jonsen, and to the associate editor, for their helpful comments that greatly improved the presentation of this work.

Authors' contribution

TM and PGB conceived the ideas and designed methodology; TM led the writing of the manuscript. Both authors contributed critically to the drafts and gave final approval for publication.

Data accessibility

The code for the simulations of Section 5 and the analysis of Section 6 is available on Zenodo, doi:10.5281/zenodo.2537668. The data used in this study are available on Movebank (movebank.org, study name “GreySeal_McConnell-UK”), and are published in the Movebank Data Repository (McConnell, 2019).

References

- Baylis, A., Orben, R., Arnould, J., Peters, K., Knox, T., Costa, D., and Staniland, I. (2015). Diving deeper into individual foraging specializations of a large marine predator, the southern sea lion. *Oecologia*, 179(4):1053–1065.
- Berman, S. M. (1994). A bivariate Markov process with diffusion and discrete components. *Stochastic Models*, 10(2):271–308.
- Bestley, S., Jonsen, I. D., Hindell, M. A., Guinet, C., and Charrassin, J.-B. (2012). Integrative modelling

of animal movement: incorporating in situ habitat and behavioural information for a migratory marine predator. *Proceedings of the Royal Society of London B: Biological Sciences*, page rspb20122262.

- Beyer, H. L., Morales, J. M., Murray, D., and Fortin, M.-J. (2013). The effectiveness of bayesian state-space models for estimating behavioural states from movement paths. *Methods in Ecology and Evolution*, 4(5):433–441.
- Blackwell, P. G. (1997). Random diffusion models for animal movement. *Ecological Modelling*, 100(1-3):87–102.
- Blackwell, P. G. (2003). Bayesian inference for Markov processes with diffusion and discrete components. *Biometrika*, 90(3):613–627.
- Blackwell, P. G., Niu, M., Lambert, M. S., and LaPoint, S. D. (2016). Exact Bayesian inference for animal movement in continuous time. *Methods in Ecology and Evolution*, 7(2):184–195.
- Brown, D. D., LaPoint, S., Kays, R., Heidrich, W., Kümme, F., and Wikelski, M. (2012). Accelerometer-informed gps telemetry: reducing the trade-off between resolution and longevity. *Wildlife Society Bulletin*, 36(1):139–146.
- Codling, E. and Hill, N. (2005). Sampling rate effects on measurements of correlated and biased random walks. *Journal of Theoretical Biology*, 233(4):573–588.
- DeRuiter, S. L., Langrock, R., Skirbutas, T., Goldbogen, J. A., Calambokidis, J., Friedlaender, A. S., Southall, B. L., et al. (2017). A multivariate mixed hidden Markov model for blue whale behaviour and responses to sound exposure. *The Annals of Applied Statistics*, 11(1):362–392.
- Dunn, J. E. and Gipson, P. S. (1977). Analysis of radio telemetry data in studies of home range. *Biometrics*, pages 85–101.
- Durbin, J. and Koopman, S. J. (2012). *Time series analysis by state space methods*, volume 38. OUP Oxford.
- Edelhoff, H., Signer, J., and Balkenhol, N. (2016). Path segmentation for beginners: an overview of current methods for detecting changes in animal movement patterns. *Movement ecology*, 4(1):21.
- Fleming, C. H., Calabrese, J. M., Mueller, T., Olson, K. A., Leimgruber, P., and Fagan, W. F. (2014). From fine-scale foraging to home ranges: a semivariance approach to identifying movement modes across spatiotemporal scales. *The American Naturalist*, 183(5):E154–E167.
- Fleming, C. H., Sheldon, D., Gurarie, E., Fagan, W. F., LaPoint, S., and Calabrese, J. M. (2017). Kálmán filters for continuous-time movement models. *Ecological Informatics*, 40:8–21.

- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.
- Gillespie, D. T. (1996). Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral. *Physical review E*, 54(2):2084.
- Gurarie, E., Fleming, C. H., Fagan, W. F., Laidre, K. L., Hernández-Pliego, J., and Ovaskainen, O. (2017). Correlated velocity models as a fundamental unit of animal movement: synthesis and applications. *Movement Ecology*, 5(1):13.
- Hays, G., Akesson, S., Godley, B., Luschi, P., and Santidrian, P. (2001). The implications of location accuracy for the interpretation of satellite-tracking data. *Animal behaviour*, 61(5):1035–1040.
- Hobolth, A. and Stone, E. A. (2009). Simulation from endpoint-conditioned, continuous-time Markov chains on a finite state space, with applications to molecular evolution. *The annals of applied statistics*, 3(3):1204.
- Hooten, M. B., Johnson, D. S., McClintock, B. T., and Morales, J. M. (2017). *Animal movement: statistical models for telemetry data*. CRC Press.
- Horne, J. S., Garton, E. O., Krone, S. M., and Lewis, J. S. (2007). Analyzing animal movements using Brownian bridges. *Ecology*, 88(9):2354–2363.
- Johnson, D. S. and London, J. M. (2018). crawl: an R package for fitting continuous-time correlated random walk models to animal movement data.
- Johnson, D. S., London, J. M., Lea, M.-A., and Durban, J. W. (2008). Continuous-time correlated random walk model for animal telemetry data. *Ecology*, 89(5):1208–1215.
- Jonsen, I. D., Flemming, J. M., and Myers, R. A. (2005). Robust state–space modeling of animal movement data. *Ecology*, 86(11):2874–2880.
- Jonsen, I. D., Myers, R. A., and Flemming, J. M. (2003). Meta-analysis of animal movement using state-space models. *Ecology*, 84(11):3055–3063.
- Langrock, R., King, R., Matthiopoulos, J., Thomas, L., Fortin, D., and Morales, J. M. (2012). Flexible and practical modeling of animal telemetry data: hidden Markov models and extensions. *Ecology*, 93(11):2336–2342.
- Leos-Barajas, V., Photopoulou, T., Langrock, R., Patterson, T. A., Watanabe, Y. Y., Murgatroyd, M., and Papastamatiou, Y. P. (2017). Analysis of animal accelerometer data using hidden Markov models. *Methods in Ecology and Evolution*, 8(2):161–173.

- McClintock, B. T. (2017). Incorporating telemetry error into hidden Markov models of animal movement using multiple imputation. *Journal of Agricultural, Biological and Environmental Statistics*, 22(3):249–269.
- McClintock, B. T., Johnson, D. S., Hooten, M. B., Ver Hoef, J. M., and Morales, J. M. (2014). When to be discrete: the importance of time formulation in understanding animal movement. *Movement ecology*, 2(1):21.
- McClintock, B. T., London, J. M., Cameron, M. F., and Boveng, P. L. (2017). Bridging the gaps in animal movement: hidden behaviors and ecological relationships revealed by integrated data streams. *Ecosphere*, 8(3):e01751.
- McConnell, B. J. (2019). Data from: State-switching continuous-time correlated random walks. Movebank data repository. doi:10.5441/001/1.m7j2263r.
- Michelot, T., Langrock, R., Bestley, S., Jonsen, I. D., Photopoulou, T., and Patterson, T. A. (2017). Estimation and simulation of foraging trips in land-based marine predators. *Ecology*, 98(7):1932–1944.
- Michelot, T., Langrock, R., and Patterson, T. A. (2016). moveHMM: An R package for the statistical modelling of animal movement data using hidden Markov models. *Methods in Ecology and Evolution*, 7(11):1308–1315.
- Morales, J. M., Haydon, D. T., Frair, J., Holsinger, K. E., and Fryxell, J. M. (2004). Extracting more out of relocation data: building movement models as mixtures of random walks. *Ecology*, 85(9):2436–2445.
- Parton, A. and Blackwell, P. G. (2017). Bayesian inference for multistate ‘step and turn’ animal movement in continuous time. *Journal of Agricultural, Biological and Environmental Statistics*, 22(3):373–392.
- Patterson, T. A., Basson, M., Bravington, M. V., and Gunn, J. S. (2009). Classifying movement behaviour in relation to environmental conditions using hidden Markov models. *Journal of Animal Ecology*, 78(6):1113–1123.
- Patterson, T. A., Parton, A., Langrock, R., Blackwell, P. G., Thomas, L., and King, R. (2017). Statistical modelling of individual animal movement: an overview of key methods and a discussion of practical challenges. *AStA Advances in Statistical Analysis*, 101(4):399–438.
- Patterson, T. A., Thomas, L., Wilcox, C., Ovaskainen, O., and Matthiopoulos, J. (2008). State-space models of individual animal movement. *Trends in ecology & evolution*, 23(2):87–94.
- Plummer, M., Best, N., Cowles, K., and Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11.

- Pohle, J., Langrock, R., van Beest, F. M., and Schmidt, N. M. (2017). Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement. *Journal of Agricultural, Biological and Environmental Statistics*, 22(3):270–293.
- Preisler, H. K., Ager, A. A., and Wisdom, M. J. (2013). Analyzing animal movement patterns using potential functions. *Ecosphere*, 4(3):1–13.
- Roberts, G. O., Gelman, A., Gilks, W. R., et al. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, 7(1):110–120.
- Robinson, P. W., Costa, D. P., Crocker, D. E., Gallo-Reynoso, J. P., Champagne, C. D., Fowler, M. A., Goetsch, C., Goetz, K. T., Hassrick, J. L., Hückstädt, L. A., et al. (2012). Foraging behavior and success of a mesopelagic predator in the northeast Pacific Ocean: insights from a data-rich species, the northern elephant seal. *PloS one*, 7(5):e36728.
- Rode, K. D., Wilson, R. R., Regehr, E. V., Martin, M. S., Douglas, D. C., and Olson, J. (2015). Increased land use by Chukchi sea polar bears in relation to changing sea ice conditions. *PloS one*, 10(11):e0142213.
- Russell, D. J., McClintock, B. T., Matthiopoulos, J., Thompson, P. M., Thompson, D., Hammond, P. S., Jones, E. L., MacKenzie, M. L., Moss, S., and McConnell, B. J. (2015). Intrinsic and extrinsic drivers of activity budgets in sympatric grey and harbour seals. *Oikos*, 124(11):1462–1472.
- Schlägel, U. E. and Lewis, M. A. (2016). Robustness of movement models: can models bridge the gap between temporal scales of data sets and behavioural processes? *Journal of mathematical biology*, 73(6-7):1691–1726.
- Springer, J. T. (1979). Some sources of bias and sampling error in radio triangulation. *The Journal of Wildlife Management*, pages 926–935.
- Uhlenbeck, G. E. and Ornstein, L. S. (1930). On the theory of the Brownian motion. *Physical review*, 36(5):823.