

Estimating behavioral parameters in animal movement models using a state-augmented particle filter

MICHAEL DOWD^{1,3} AND RUTH JOY^{2,4}

¹Department of Mathematics and Statistics, Dalhousie University, Halifax, Nova Scotia B3H 3J5 Canada

²Marine Mammal Research Unit, University of British Columbia, Vancouver, British Columbia V6T 1Z4 Canada

Abstract. Data on fine-scale animal movement are being collected worldwide, with the number of species being tagged and the resolution of data rapidly increasing. In this study, a general methodology is proposed to understand the patterns in these high-resolution movement time series that relate to marine animal behavior. The approach is illustrated with dive data from a northern fur seal (*Callorhinus ursinus*) tagged on the Pribilof Islands, Alaska, USA. We apply a state-space model composed of a movement model and corresponding high-resolution vertical movement data. The central goal is to estimate parameters of this movement model, particularly their variation on appropriate time scales, thereby providing a direct link to behavior. A particle filter with state augmentation is used to jointly estimate the movement parameters and the state. A multiple iterated filter using overlapping data segments is implemented to match the parameter time scale with the behavioral inference. The time variation in the auto-covariance function facilitates identification of a movement model, allows separation of observation and process noise, and provides for validation of results. The analysis yields fitted parameters that show distinct time-evolving changes in fur seal behavior over time, matching well what is observed in the original data set.

Key words: behavioral inference; movement model; northern fur seal; parameter estimation; particle filter; state-space model; time series analysis; vertical velocity.

INTRODUCTION

Fine-scale archival, data-logging technology has given rise to a rapidly growing body of information on the movement of many marine animal species. It is widely recognized that these technological developments have far outpaced the analysis methods available for extracting meaningful biological information from these high-resolution and complex data types (Schick et al. 2008). To date, many applications of animal movement data have focused on reconstructing tracks from sparse and noisy fixes of geographical position (Jonsen et al. 2005), but as positional information improves and motion sensors are incorporated (Ropert-Coudert and Wilson 2005), there is strong interest in inferring behavior from much higher resolution data (Polansky et al. 2010). These data are fundamentally time series whose salient character is their autocorrelated structure and the nonstationarity (Gurarie et al. 2009). State-space models provide a flexible framework for a unified treatment of tag time series and animal movement models, and are recognized as a promising way forward (Patterson et al. 2008). Here, we propose a modeling approach that uses

high-resolution movement data to estimate continuously varying behavioral parameters of movement models on the appropriate time scale.

Animal movement models are formulated as stochastic differential (or difference) equations, such as the correlated random walk (Morales et al. 2004, Codling et al. 2008). These models can describe a wide variety of movement patterns depending on their parameter values. Our main premise is that as the character of the movement observations change over the course of the record, so do the corresponding parameters of the animal movement model that embody the behavioral information. These parameters can be estimated directly using a state-space model formulation. One possibility is to use a state space behavioral switching model (Jonsen et al. 2007, Patterson et al. 2009), but this requires a pre-definition of a small set of behavioral modes, and the probability of transitioning between them. We propose an alternative approach that directly estimates the time-varying behavioral parameters of a movement model using state space methods.

We develop and illustrate the approach using a high resolution data set from an at-sea foraging track of a northern fur seal (*Callorhinus ursinus*). The analysis uses vertical velocity (dive) data recorded at a two second time interval. These data are of a much higher resolution (and volume) than considered by other state-space models that infer behavior using episodic horizontal position fixes (e.g., Jonsen et al. 2007). Moreover, these

Manuscript received 24 March 2010; revised 6 July 2010; accepted 13 September 2010. Corresponding Editor: D. C. Speirs.

³ E-mail: mdowd@mathstat.dal.ca

⁴ Present address: Department of Statistics, Simon Fraser University, Burnaby, British Columbia V6B 5K3 Canada.

data serve to illustrate the nontrivial issues that arise in application of the state-space model for a realistic example. We emphasize, however, that the approach is a general one, and applicable to most horizontal and/or vertical movement models and high-resolution tag data.

METHODS

Our approach fits a stochastic animal movement model to noisy position or velocity data using a state-space model, simultaneously estimating its parameters through a state augmentation procedure. The first part of a state-space model is the state evolution equation, or movement model, which describes the movement process, and allows it to evolve forward in time

$$\mathbf{x}_t = \mathbf{f}(\mathbf{x}_{t-1}, \boldsymbol{\theta}_t) + \mathbf{n}_t. \quad (1)$$

Here, \mathbf{x}_t is a column vector that represents the system state (either a position or velocity) at time t . The movement model is embodied in the operator \mathbf{f} , which depends on the state at the previous state, \mathbf{x}_{t-1} , and a vector of movement, or behavioral, parameters, $\boldsymbol{\theta}_t$, that can change through time. The system noise is given by \mathbf{n}_t .

The second part of the state-space model is the observation equation in which the data observations are assimilated

$$\mathbf{y}_t = \mathbf{H}\mathbf{x}_t + \mathbf{e}_t \quad (2)$$

where \mathbf{y}_t is an observation of vertical velocity at time t . It is related to the state, \mathbf{x}_t , through the matrix \mathbf{H} which allows for conversion between measured and modeled variables (say, observing position and modeling velocity), and partial observation of the state. The observation error process, \mathbf{e}_t , follows a specified probability distribution. There may also be parameters associated with \mathbf{e}_t or \mathbf{H} , but these are not considered in this study.

The state-space model (Eqs. 1 and 2) is a very general formulation. It is characterized by Markovian dynamics, conditionally independent observations, and mutually independent system and observation noise. It allows for nonlinear models and non-Gaussian error processes. The usual goal is to estimate the state, \mathbf{x}_t , over time (or strictly speaking, its probability density function). Sequential Monte Carlo methods, such as particle filters, provide standard solution techniques. However, the usual approaches do not provide estimates for the parameters, $\boldsymbol{\theta}_t$ (Ristic et al. 2004).

To estimate parameters in a state-space model, two main approaches are available. A likelihood function can be computed using a particle filter; however, Monte Carlo variation means the likelihood surface may be difficult to maximize (Doucet and Tadić 2003). Another approach, and the one that is the focus of this paper, is state augmentation (Kitagawa 1998). Here, we append the parameters of interest to the state vector, forming an augmented state vector. Standard sequential Monte Carlo methods are then used to determine parameter estimates.

By making the parameters in the movement model (Eq. 1) follow a simple random walk, the augmented state evolution equation becomes

$$\begin{pmatrix} \mathbf{x}_t \\ \boldsymbol{\theta}_t \end{pmatrix} = \begin{pmatrix} \mathbf{f}(\mathbf{x}_{t-1}, \boldsymbol{\theta}_t) \\ \boldsymbol{\theta}_{t-1} \end{pmatrix} + \begin{pmatrix} \mathbf{n}_t \\ \mathbf{v}_t \end{pmatrix} \quad (3)$$

or

$$\tilde{\mathbf{x}}_t = \mathbf{g}(\tilde{\mathbf{x}}_{t-1}) + \tilde{\mathbf{n}}_t. \quad (4)$$

The augmented state is given by $\tilde{\mathbf{x}}_t$, and now includes the parameters. The system noise $\tilde{\mathbf{n}}_t$ now includes both the system noise, \mathbf{n}_t , as well as the disturbance term, \mathbf{v}_t , for the simple random walk of $\boldsymbol{\theta}_t$. The operator \mathbf{g} includes both the movement dynamics \mathbf{f} and the random walk for its parameters. The augmented observation Eq. 2 is now $\mathbf{y}_t = \mathbf{H}\tilde{\mathbf{x}}_t$, where $\mathbf{0}$ is a vector of zeros and implies that we cannot observe the parameters, $\boldsymbol{\theta}_t$, directly. The augmented state-space model can be solved using the same sequential Monte Carlo methods as for regular state-space models, since this is nothing more than a reformulation of the state. Hence, dynamic, or continuously time-varying, parameters are straightforward to estimate using state augmentation. Static, or slowly time-varying, parameters are more complex however.

The central goal in this study is to estimate the time-varying movement parameters that have a direct link to fur seal behavior. Hence, identification of the appropriate time scale for this estimation is a key element. Standard application of the augmented state-space model (Eq. 4) with a parameter random walk having a constant variance, σ_v^2 (c.f. Kitagawa 1998), means that the estimated time-varying movement parameters simply follow short period fluctuations in the data, which is not always desirable for analysis and inference. For example, the vertical dive data in this study (see *Data*) was sampled every two seconds, which meant that movement parameters varied within individual dives. However, the movement models under consideration here are appropriate for describing behavior on longer time scales, i.e., the variations associated with sequences, or ensembles, of dives that collectively encapsulate behavior. We reformulate the parameter estimation problem accordingly.

Parameters of the movement model are estimated as static parameters for short segments of the data record. In order to consider the entire record, the time series is broken up into short, overlapping segments (or time windows). Parameter values are then determined for each of these time windows. The resultant time sequence of parameter estimates can indicate behavioral changes over time. The length of the data segments is guided by the type of movement model considered, along with the observed time scale for behavioral changes. This idea of evolving parameters is well established for nonstationary statistics that vary over time (e.g., Priestley 2004: chapter 11), and has the advantage that the investigator can control the time scale for behavioral parameter

estimation by choosing an appropriate data segment length.

The estimation procedures for parameters that are static within an analysis time window also relies on the state augmentation procedures outlined above. The main idea is that by successively reducing the variance of the disturbance term, σ_v^2 , in the parameter random walk, the movement parameter vector, θ_t , ends up being fixed at a particular value. One way this can be accomplished is by using a single pass of a particle filter over an analysis time window and reducing σ_v^2 with increasing t , until σ_v^2 reaches a small value at the end of the analysis segment and the parameters, θ , are locked in at their estimated values (Kitagawa 1998). As a related alternative, Ionides et al. (2006) suggest using multiple passes of a state-augmented particle filter, where each pass has a successively smaller σ_v^2 ; they show this can yield the maximum likelihood value for the parameter vector θ . We make use of a hybrid approach, using the multiple iterated filtering approach of Ionides et al. (2006), but using unweighted averages at each pass of the particle filter and a simple ramp-down of the variance, σ_v^2 . This is done with a sequential time windowing procedure and permits the estimates of the movement parameter vector to slowly time-vary over successive windows. In *Data*, we examine the observed temporal variations in northern fur seal dive behavior and link it to a suitable movement model.

DATA

Our analysis uses part of a data set collected from lactating northern fur seals tagged (with Driessen and Kern dead reckoner tags) at Reef Rookery on the Pribilof Islands, Alaska in 2005 and 2006. We focus on analyzing diving data for this study. Fig. 1 shows the depth measurements and the derived (via differencing) vertical velocity from the dive record of a single fur seal on the third day (18 August 2006) of a 12.5-day at-sea foraging trip. The two second sampling interval of the archival tag yielded 43200 data points for this single day, and these data are the basis for our application.

These data clearly show a distinction between dive and non-dive periods. The active dive periods also have distinctly different signatures or behaviors. Fig. 1 shows details of three selected one hour segments. In segment A, we see rapid, regular, relatively shallow diving with large frequent vertical speed changes, which might be indicative of foraging/feeding. In segment B, the dives are becoming deeper, less frequent and irregular with more surface time and smaller vertical speeds, perhaps indicating exploring/searching. In segment C, the fur seal is at the surface with a small vertical speed, and is either transiting or resting. The statistical character of these velocity observations obviously changes over the course of the record, and our goal is to quantify these by determining the associated time-varying behavioral parameters of a suitable movement model.

A nonstationary time series analysis was undertaken in the form of the sample evolutionary auto-covariance function (ACVF) of the vertical velocity data. This is a way to account for auto-correlation that varies over the data record, analogous to time-frequency analysis (e.g., Wittemyer et al. 2008, Polansky et al. 2010). Specifically, the ACVF was computed for 110 sliding time windows, each with a length of 26 minutes, and overlapping the previous window by 13 minutes, covering the length of our selected 24-hour period. The time window of 26 minutes (or 780 data points) was chosen so that these movement statistics would be approximately stationary within the data window, and also covered enough individual dives so that the emergent behavior of the fur seal could be determined.

Fig. 2a shows the evolutionary ACVF for observed vertical velocity, y_t . It indicates an obvious cycling between dive periods and quiescent non-dive periods. Within the dive period (e.g., between 00:00 and 06:00 hours, including segment A), we see an oscillating and decaying ACVF with a relatively large variance, suggesting a periodic process. During the non-dive periods (e.g., 17:00–22:00, including segment C), the ACVF has low variance and decays to zero for lags one and beyond, suggesting a purely random process. Behavioral distinctions between the dive and non-dive periods are evident, especially in terms of the shifting in the magnitude of the velocity variance. The periods of diving show a dominance of temporally coherent motions with periods reflecting the dive length.

A key aspect of state-space modeling is to identify a suitable movement model. We chose an auto-regressive model of order two, or AR(2), as the most suitable model

$$z_t = a_1 z_{t-1} + a_2 z_{t-2} + \varepsilon_t. \quad (5)$$

Here, z_t represents the vertical velocity at time t , and a_1 and a_2 are coefficients that multiply z lagged by 1 and 2 time units, respectively. (Note that z_t will become a portion of the augmented state \tilde{x}_t). The system noise ε_t will be taken as a purely random Gaussian process. The model is the most parsimonious description that can explain the ACVF, i.e., oscillating and decaying during dives, and cutting off after lag 0 during non-dives. (The choice of this model was also validated with formal statistical time series model fitting procedures).

This AR(2) movement model (Eq. 5) corresponds, in continuous time, to a second-order stochastic ordinary differential equation. Its coefficients are interpretable in terms of fur seal dive behavior. The variance of ε_t scales the magnitude of the dive velocity. If, say, $a_1 = a_2 = 0$ then the model corresponds to a purely random process which might describe incoherent vertical velocity signal associated with surface swimming. If $a_1 = 1$, it is a simple random walk. Generally, however, a_1 and a_2 will be non-zero. If $a_2 < -a_1^2/4$ movement is pseudo-periodic and could describe a regular and repeating set of dives (Priestley 2004: section 3.5.3). The vertical motion is

aperiodic when $a_2 > -a_1^2/4$. The model in Eq. 5 is therefore a flexible description for fur seal vertical motion and its parameters values can be linked to behavior.

APPLICATION

The state augmentation approach estimates the parameters, a_1 and a_2 , of the movement model (Eq. 5). However the model must first be in the Markovian form (Eq. 1 or 4), i.e., with a dependence on only one time lag. Hence, we rewrite Eq. 5 as follows:

$$\begin{pmatrix} z_t \\ \zeta_t \end{pmatrix} = \begin{pmatrix} a_1 & a_2 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} z_{t-1} \\ \zeta_{t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_t \\ 0 \end{pmatrix} \quad (6)$$

where the dummy variable ζ has been introduced to transform a non-Markovian model into a Markovian one.

The state vector (the left-hand side of Eq. 6) is next augmented with the parameters a_1 and a_2 . These each follow a simple random walk:

$$\begin{aligned} a_{1,t} &= a_{1,t-1} + v_{1,t} & v_{1,t} &\sim \mathcal{N}(0, \sigma_v^2) \\ a_{2,t} &= a_{2,t-1} + v_{2,t} & v_{2,t} &\sim \mathcal{N}(0, \sigma_v^2) \end{aligned} \quad (7)$$

where the variance of the disturbance term, σ_v^2 , must be appropriately specified to allow for estimation of the movement parameters. An algorithm for doing this is given later.

The augmented state vector is therefore $\tilde{\mathbf{x}}_t = (z \ \zeta \ a_1 \ a_2)'_t$, and the state evolution equation is defined by Eqs. 6 and 7. The augmented system noise term is $\tilde{\mathbf{n}}_t = (\varepsilon \ 0 \ v_1 \ v_2)'_t$. The vertical velocity observations (Fig. 1b) are contained in \mathbf{y}_t , and it is assumed that the observation error $\mathbf{e}_t \sim \mathcal{N}(0, \sigma_e^2)$. The augmented observation operator in Eq. 2 is a row vector $\tilde{\mathbf{H}} = (1 \ 0 \ 0 \ 0)$ which multiplies the augmented state $\tilde{\mathbf{x}}_t$, and indicates we observe only its first element.

The statistics of the observation error, \mathbf{e}_t , and the system noise, ε_t , are difficult to specify since stochastic variations seen in the observations are due to both measurement noise, as well as fluctuations due to the animal movement process itself. Their variances, σ_e^2 and σ_v^2 , can, however, be separated and estimated using the ACVF and the principle that the observation error is uncorrelated over time, whereas animal movement is time correlated (see Appendix A). An advantage of determining the system and observation error variances offline is that it minimizes the number of identifiable parameters that need to be estimated.

As part of the implementation of the state-space model, the system noise takes the form of a normal mixture process

$$\varepsilon_t \sim c_1 \mathcal{N}(0, \sigma_e^2) + c_2 \mathcal{N}(0, \delta \sigma_e^2). \quad (8)$$

The system noise has a overall variance corresponding to the offline estimated σ_e^2 , and will vary in time. We choose $c_1 = 0.9$, $c_2 = 0.1$, and $\delta = 10$ which allow the




FIG. 1. (a) Observed time series of depth (blue line) and (b) derived vertical velocity (black line) from a northern fur seal collected on 18 August 2006, with hour of day indicated. Panels (c)–(e) show depth (blue line) and vertical velocity (black line) for three selected one-hour-long data segments, for time periods as indicated [the gray shaded areas labeled A, B, and C in panels (a) and (b)].

system noise variance to have occasional large values. This mixture form for the system error allows the state estimates to closely follow the observations when abrupt changes in velocity occurred, such as those found at the onset of a dive.

The state z_t (and ζ_t) and the parameters a_1 and a_2 are jointly estimated using a state-augmented particle filter based on sequential importance resampling (for details, see Appendix B). This estimation is carried out for each of the 110 analysis time windows (each 26 minutes long, and overlapping by 13 minutes). The slow time variation in these parameters across the 110 windows corresponds to the changes in fur seal behavior over the data record.

To estimate the static parameters a_1 and a_2 within a single time window, we use the following algorithm:

- 1) Run the particle filter with the data \mathbf{y}_t using the state-augmented model (Eqs. 6 and 7), and with system noise and observation noise as specified above. Set the initial random walk variance for the parameters at $\sigma_v^2(0)$ (the bracketed value indicates the iteration number), as well as initial values for $a_{1,t=0}(0)$ and $a_{2,t=0}(0)$. Take the mean estimated by the particle filter across all time steps of $\hat{a}_{1,t}(0)$ and $\hat{a}_{2,t}(0)$ to be the parameter values obtained from iteration zero, i.e., $\hat{a}_1(0)$ and $\hat{a}_2(0)$.
- 2) For iterations $k = 1, \dots, m$. Let $\sigma_v^2(k) = \alpha \sigma_v^2(k - 1)$, where α is a discounting factor used to control the

reduction in the random walk variance at each iteration. Use as initial values for $a_{1,t=0}(k) = \hat{a}(k-1)$ and $a_{2,t=0}(k) = \hat{a}_2(k-1)$. Take the mean of $a_{1,t}(k)$ and $a_{2,t}(k)$ from the particle filter for estimates for $\hat{a}_1(k)$ and $\hat{a}_2(k)$ for iteration k .

3) Stop after m iterations. Use $\hat{a}_1(m)$ and $\hat{a}_2(m)$ as final estimates for a_1 and a_2 .

This is a simplified (unweighted) version of the multiple iterated filter of Ionides et al. (2006). It was found that the following values yielded good results: (1) an initial $\sigma_v^2(0) = 0.1^2$; (2) a discounting factor of $\alpha = 0.5$, and (3) $m = 10$ iterations providing a stopping criteria ($\sigma_v^2(10)$ was by then small enough that the parameters were effectively fixed at a constant value). This algorithm was applied for each of the 110 time analysis windows to estimate the time variation of a_1 and a_2 .

To test the accuracy and stability of this algorithm, we ran multiple runs with synthetic data generated using the movement model and known parameters. These tests indicated that the above algorithm provided a good recovery of the movement parameters. In the Supplement, R code is given for implementing the entire procedure.

RESULTS

State estimates for the vertical velocity are obtained as part of the analysis but were not of interest in our behavioral inference. The estimated state (Fig. 3b, c, in gray) conforms well to the velocity observations (Fig. 1b). In fact, it is simply a low-pass filtered version, which identifies the underlying velocity signal. It is not discussed further.

Fig. 3a shows the offline estimates of the system noise variance and observation error variance, and how they evolve through time for each of the analysis time windows. Both scale proportionally to vertical speed, and are large during periods of active diving, dropping toward zero during non-diving phases. These variances, σ_o^2 and σ_e^2 , were used as inputs to the state-space model.

Fig. 3 also shows the estimated slowly time varying movement parameters from the analysis, i.e., the autoregressive coefficients a_1 (Fig. 3b) and a_2 (Fig. 3c). During the initial part of the day (00:00–06:00, and highlighted in segment A in Fig. 1c, and segment A in Fig. 3a–c), the fur seal exhibits rapid, frequent, regular dives and $a_1 \approx 1.5$, $a_2 \approx -0.6$, and the system noise variance, σ_e^2 , is high (0.15–0.2). Just past midday, there is a 4-hour period of occasional deep dives with large velocities (highlighted in segment B in Fig. 1d and segment B in Fig. 3a–c). The system noise variance increases to 0.05, and $a_1 \approx 0.8$ and $a_2 \approx 0.1$. There are also periods where the fur seal remains at the surface with small vertical speeds (highlighted in segment C in Fig. 1e and segment C in Fig. 3a–c) in which the a_1 and a_2 parameters and system noise are near zero. The behavioral signatures in these vertical velocity data are however not always discrete modes, but have time variation and mixes of behaviors. These are captured

well by the continuously time-varying movement parameters (Fig. 3b, c).

To further validate the results, we address the question of how well the estimated parameters can reconstruct the emergent statistical properties of the data. The evolutionary ACVF can be predicted from knowledge of the time evolution of a_1 , a_2 , and σ_e^2 (Priestley 2004: chapters 3 and 4). The predicted ACVF can then be compared to its corresponding sample version computed from the data in Fig. 2b. Note that the sample ACVF does not separate out stochastic fluctuations in the movement process from the observations, as does the state-space model. Also, since the movement model is a simplified representation of reality, it can only capture a portion of the variability in the data. Hence we do not expect exact comparability or reproduction. It is nevertheless a useful comparison and validation.

The main features of the predicted ACVF (Fig. 2b) compare very well with the sample ACVF (Fig. 2a), especially at the important small time lags. In particular, the oscillation and decay of the ACVF with increasing lag is predicted for the initial foraging period (segment A, Fig. 2c). In fact, the major differences for this period are the added noise at large lags in the sample ACVF, and that the sample ACVF has deeper negative values at larger lags than does the predicted one. This latter discrepancy occurs since the observed velocity series does not have an exact periodicity even within a single time analysis segment. That is, the dives are quite regular, but not exactly repeating. In Fig. 2d, the sample and predicted ACVF are compared for segment B, characterized by irregular diving; the magnitude and decay rate of the ACVF is well captured in both sample and predicted ACVF. Finally, in Fig. 2e, the low velocities and cutoff after lag zero characterizing segment C are seen in both the sample and predicted ACVF.

CONCLUDING REMARKS

We have explored the idea of estimating behavior parameters from marine animal archival tag data using movement models. The central idea is that by estimating the time variation of parameters for a suitable movement model, researchers can then objectively and quantitatively infer animal activity (and its behavioral state). Here, we have offered a statistical-dynamical approach suitable for extracting behavioral information from high-resolution data, and one that is widely applicable to a variety of movement models and observation types. It offers an alternative to behavioral switching state-space models (Jonsen et al. 2007, Patterson et al. 2009), and is flexible enough to obtain solutions even for nonlinear and non-Gaussian cases. Our approach builds upon the state augmentation procedures of Kitagawa (1998) and Ionides et al. (2006).

Our application focused on analysis and interpretation of the vertical movement data from a tagged northern fur seal. The analysis provides slowly time-

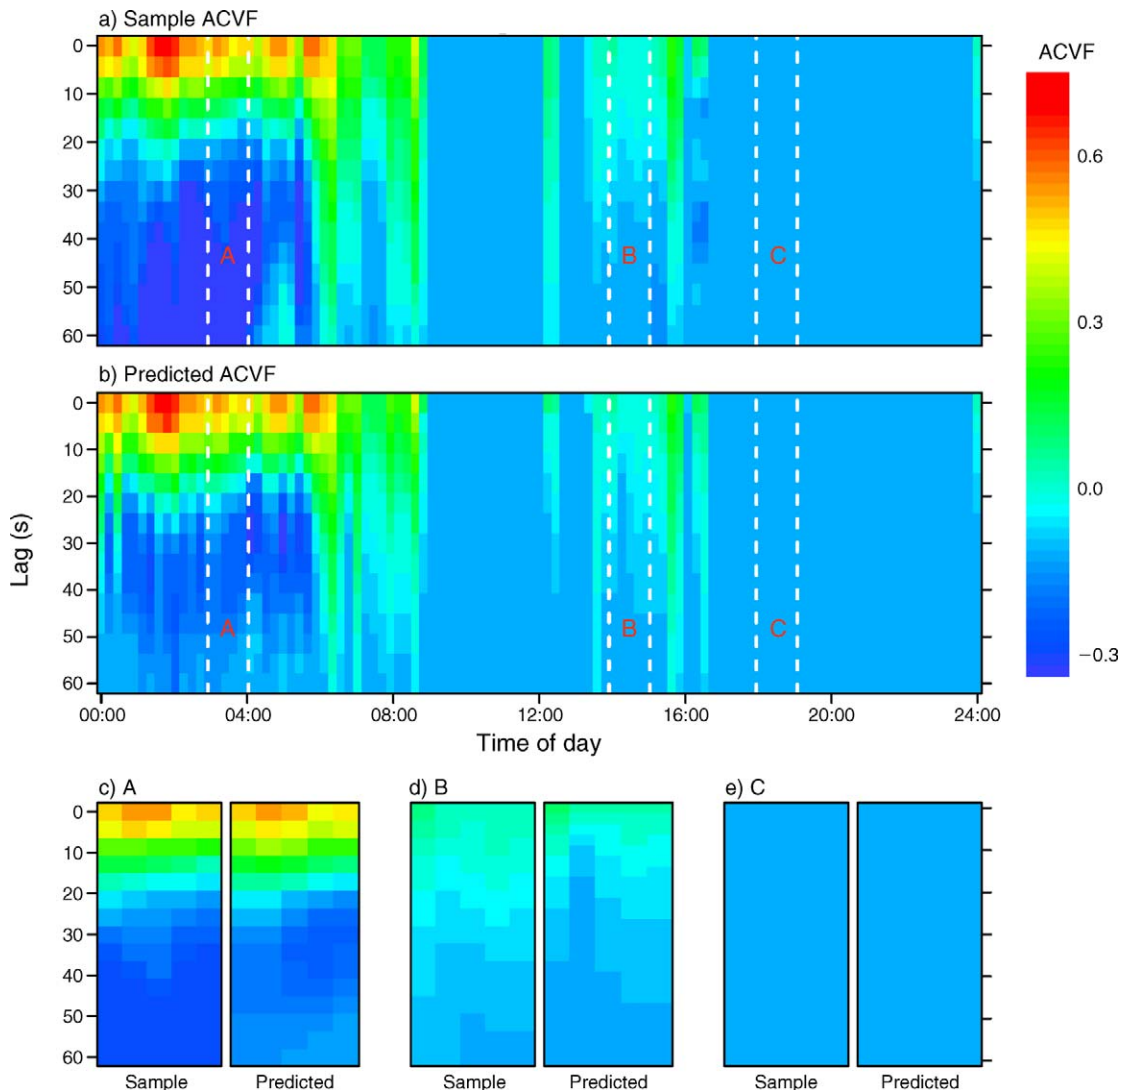


FIG. 2. Evolutionary auto-covariance function (ACVF) of vertical velocity. Panel (a) shows the sample ACVF computed from the data with time of day indicated. Panel (b) shows the corresponding ACVF predicted from the estimated behavioral parameters of the movement model. Panels (c)–(e) show the details for time segments A, B, and C [indicated by dotted lines in panels (a) and (b)] for both the sample and predicted ACVF.

varying estimates for the movement parameters. We were able to discriminate between periods of rapid shallow regular diving, episodic deep diving, and transiting/resting behavior, as well as indicate the times spent in these behavioral modes, the transitions between them, and any mixed behaviors. The application also highlighted issues that are expected to arise in any analysis of high resolution movement data, and for any movement model. These include (1) the identification of an appropriate movement model; (2) choosing the appropriate time scale for parameter estimation; and (3) the estimation of identifiable parameters. Remarks on each of these issues are offered below.

Identification of the appropriate movement model is an important issue that affects interpretation of the data,

and is key to making meaningful conclusions about behavior. In fact, we want models for which the parameters are directly interpretable as behavior. Model choice must be based both on theoretical precepts, as well as statistical features of the data. For this study, the discrete-time AR(2) model was well suited for the vertical velocity data, but continuous time formulations of movement are also possible within the state-space framework (Johnson et al. 2008). Generalizing the approach to two-dimensional and three-dimensional movement would require different types of models, such as correlated random walks (Morales et al. 2004). These would replace the process model (Eq. 1), and parameters could be estimated with the state augmentation procedure.




FIG. 3. Parameter estimates over time as obtained from the analysis. Panel (a) shows the estimated observation error variance and system noise variance. Panel (b) shows the estimated a_1 coefficient from the state-space model (black line), along with a low-pass filtered version of it (blue line). The state estimate for the vertical velocity time series is also shown (light gray). The detailed data segments (A, B, C) corresponding to Figs. 1 and 2 are indicated by gray shading. Panel (c) shows the corresponding information for the a_2 coefficient.

An issue specific to the method is the choice of the time-scale for the movement analysis. The lower limit for time window length is chosen based on the natural time scale of the movement model. Here, the AR(2) model is applicable to a time series comprised of an coherent collection of similar dives, and not a single dive. Therefore, in the northern fur seal the window width must be at least wide enough to encompass a series of dives. The upper limit on the window length is based on the time-scale at which fur seal behavioral changes take place, with the idea that the parameter estimation use a data segment that is approximately stationary (i.e., parameters should be constant over the window). The 26-minute window was a reasonable compromise, and the 13-minute overlap allows for us to resolve more abrupt changes in the parameter values. Note that some types of models with low resolution data for which behavior can be directly inferred from the state may not need to be time windowed and state augmentation can proceed directly without recourse to the multiple iterated filtering.

Parameter identifiability is a concern for many ecological and movement models, especially as complexity is increased to account for more features of behavior, and additional spatial dimensions are incorporated. One possibility is to use prior information on the parameter values, as is done with hierarchical Bayesian state-space models of animal behavior (Morales et al. 2004). Within the context of our particle filtering approach, priors could be built into the resampling step (thereby changing the weights), or alternatively hybrid particle-MCMC approaches (e.g., Andrieu et al. 2010) could be considered.

In summary, recent review articles have identified state-space models (Patterson et al. 2008) and hierarchical Bayesian approaches (Schick et al. 2008) as two important directions for extracting ecologically meaningful information from animal tag data. Here we have used a particle filter approach for an augmented state-space model that is designed for the type of large volume, high resolution motion time series recorded by archival animal tags. The approach allows for estima-

tion of the parameters in any movement model, from which behavior can be then inferred. The approach of this study offers a promising direction for more fully exploiting the behavioral information in these rich data sets.

ACKNOWLEDGMENTS

These data were collected as part of the North Pacific Research Board Project F2636. We are indebted to Pamela Lestenkof for the collection of these data, and for her motivating interest in understanding the at-sea behavior of the northern fur seal. We are grateful to Jordan Rosenfeld and to two anonymous reviewers for their constructive comments on an earlier manuscript draft, and to Andrew Trites for his support throughout this project. M. Dowd was supported by an NSERC Discovery Grant. R. Joy was supported by funding to the North Pacific Universities Marine Mammal Research Consortium from the North Pacific Marine Science Foundation.

LITERATURE CITED

- Andrieu, C., A. Doucet, and R. Holenstein. 2010. Particle Markov chain Monte Carlo. *Journal of the Royal Statistical Society B* 72(3):269–342.
- Codling, E. A., M. J. Plank, and S. Benhamou. 2008. Random walk models in biology. *Journal of the Royal Society Interface* 5:813–834.
- Doucet, A., and V. B. Tadić. 2003. Parameter estimation in general state space models using particle methods. *Annals of the Institute of Statistical Mathematics* 55(2):406–422.
- Gurarie, E., R. D. Andrews, and K. L. Laidre. 2009. A novel method for identifying behavioural changes in animal movement data. *Ecology Letters* 12:395–408.
- Ionides, E. L., C. Breto, and A. A. King. 2006. Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences USA* 103:18438–18443.
- Johnson, D. S., J. M. London, M.-A. Lea, and J. W. Durban. 2008. Continuous-time correlated random walk model for animal telemetry data. *Ecology* 89:1208–1215.
- Jonsen, I. D., R. A. Myers, and J. M. Flemming. 2005. Robust state-space modeling of animal movement data. *Ecology* 86:2874–2880.
- Jonsen, I. D., R. A. Myers, and M. C. James. 2007. Identifying leatherback turtle foraging behaviour from satellite-telemetry using a switching state-space model. *Marine Ecology Progress Series* 337:255–264.
- Kitagawa, G. 1998. A self-organising state-space model. *Journal of the American Statistical Association* 93:1203–1215.
- Morales, J. M., D. T. Haydon, J. Frair, K. E. Holsinger, and J. M. Fryxell. 2004. Extracting more out of relocation data: building movement models as mixtures of random walks. *Ecology* 85:2436–2445.
- Patterson, T. A., M. Basson, M. V. Bravington, and J. S. Gunn. 2009. Classifying movement behaviour in relation to environmental conditions using hidden Markov models. *Journal of Animal Ecology* 78:1113–1123.
- Patterson, T. A., L. Thomas, C. Wilcox, O. Ovaskainen, and J. Matthiopoulos. 2008. State-space models of individual animal movement. *Trends in Ecology and Evolution* 23:87–94.
- Polansky, L., G. Wittemyer, P. C. Cross, C. J. Tambling, and W. M. Getz. 2010. From moonlight to movement and synchronized randomness: Fourier and wavelet analyses of animal location time series data. *Ecology* 91:1506–1518.
- Priestley, M. B. 2004. *Spectral analysis and time series*. Academic Press, London, UK.
- Ristic, B., S. Arulampalam, and N. Gordon. 2004. *Beyond the Kalman filter: particle filters for tracking applications*. Artech House, Boston, Massachusetts, USA.
- Ropert-Coudert, Y., and R. P. Wilson. 2005. Trends and perspectives in animal-attached remote sensing. *Frontiers in Ecology and Environment* 3:437–444.
- Schick, R. S., S. R. Loarie, F. Colchero, B. D. Best, A. Boustany, D. A. Conde, P. N. Halpin, L. N. Joppa, C. M. McLennan, and J. S. Clark. 2008. Understanding movement data and movement processes: current and emerging directions. *Ecology Letters* 11:1338–1350.
- Wittemyer, G., L. Polansky, I. Douglas-Hamilton, and W. M. Getz. 2008. Disentangling the effects of forage, social rank, and risk on movement autocorrelation of elephants using Fourier and wavelet analyses. *Proceedings of the National Academy of Sciences USA* 105:19108–19113.

APPENDIX A

The approach used to specify the observation error and system noise variance (*Ecological Archives* E092-049-A1).

APPENDIX B

A primer on particle filtering (*Ecological Archives* E092-049-A2).

SUPPLEMENT

R code for implementing the multiple iterative filtering methodology (based on an idealized example) (*Ecological Archives* E092-049-S1).