Classifying movement behaviour in relation to environmental conditions using hidden Markov models

Toby A. Patterson[1,2]*, Marinelle Basson[1], Mark V. Bravington[3] and John S. Gunn[1]

[1]*CSIRO Marine & Atmospheric Research, GPO Box 1538, Hobart, Tas. 7001, Australia;* [2]*Antarctic Wildlife Research Unit, School of Zoology, University of Tasmania, PO Box 252-05 Hobart, Tas. 7001, Australia; and* [3]*CSIRO Mathematics and Information Science, GPO Box 1538, Hobart, Tas. 7001, Australia*

Summary

1. Linking the movement and behaviour of animals to their environment is a central problem in ecology. Through the use of electronic tagging and tracking (ETT), collection of *in situ* data from free-roaming animals is now commonplace, yet statistical approaches enabling direct relation of movement observations to environmental conditions are still in development.

2. In this study, we examine the hidden Markov model (HMM) for behavioural analysis of tracking data. HMMs allow for prediction of latent behavioural states while directly accounting for the serial dependence prevalent in ETT data. Updating the probability of behavioural switches with tag or remote-sensing data provides a statistical method that links environmental data to behaviour in a direct and integrated manner.

3. It is important to assess the reliability of state categorization over the range of time-series lengths typically collected from field instruments and when movement behaviours are similar between movement states. Simulation with varying lengths of times series data and contrast between average movements within each state was used to test the HMMs ability to estimate movement parameters.

4. To demonstrate the methods in a realistic setting, the HMMs were used to categorize resident and migratory phases and the relationship between movement behaviour and ocean temperature using electronic tagging data from southern bluefin tuna (*Thunnus maccoyii*). Diagnostic tools to evaluate the suitability of different models and inferential methods for investigating differences in behaviour between individuals are also demonstrated.

Key-words: archival and satellite tags, behavioural analysis, bluefin tuna, hidden Markov models.

Introduction

Understanding how individual animals alter their movements in relation to habitat is an important problem in foraging ecology (e.g. Bestley *et al*. 2008), habitat selection studies (e.g. Aarts *et al*. 2008) and spatial population ecology (Bowler & Benton 2005). An important starting point for all these branches of ecology is that animals should maximize their time in productive areas (Charnov 1976). However, foraging itself is rarely observed so changes in individuals' movement patterns are widely interpreted as indicating switching between different underlying behavioural states (Nathan *et al*. 2008). Thus, episodes of slow and variable movement are interpreted as the visible consequence of intensive foraging (Barraquand & Benhamou 2008). Faster, directed movements are taken to indicate either inter-patch movements or searching behaviour and are often viewed as responses to

factors like habitat change or food availability (see Bowler & Benton 2005 for a review).

Understanding how habitat alters behaviour is important because collective individual behaviour, by definition, has population-level consequences; influencing spatial distribution (Morales & Ellner 2002); causing aggregations of predators in areas of high prey density (i.e. prey taxis Kareiva & Odell 1987); and in influencing the spread of invasive species and diseases (Haydon 2008). Classifying important regions for conservation on the basis of the relationship between individuals' space usage and the environment is also becoming more common (e.g. Hyrenbach *et al*. 2006; Shillinger *et al*. 2008).

Relating individual movement and behaviour to habitat requires adequate data. Bio-logging technology (Hooker *et al*. 2007), via electronic tagging and telemetry (ETT), enables collection of high-resolution data of movement and environment. However, a weak link in understanding the relationship between habitat and behaviour is that

*Correspondence author. E-mail: toby.patterson@csiro.au

ETT devices typically measure basic physical data – e.g. pressure, temperature and salinity – and not quantities which directly influence the foraging behaviour of animals. Putative behaviours identified from movement data must instead be related to physical proxies (Hindell *et al.* 2003; Teo, Boustany & Block 2007; Gremillet *et al.* 2008). Indirect data of foraging success such as stomach temperature (Austin *et al.* 2006; Bestley *et al.* 2008) or body condition such as buoyancy (Biuw *et al.* 2003) have been collected. However, direct measures of say, feeding success or spawning behaviour are rare and *in situ* data of prey availability is even less common.

Therefore, both behaviour and its probable drivers are hidden to the observer with most types of ETT technology and uncovering hidden behaviour becomes a statistical problem. Many recent studies have sought to uncover hidden behaviour solely from the geometry of the movement path. Amongst these are first-passage time methods (FPT; Fauchald & Tveraa 2003), Lévy flight analysis (Sims *et al.* 2008) and state-space models (SSM; Jonsen, Myers & James 2007). However, these recent approaches have some limitations. Lévy methods are limited by their lack of predictive capacity, making them unsuited to behavioural categorization. FPT and SSM studies to date have mostly considered the movement data in isolation, without reference to habitat data, to categorize behaviour (also see Barraquand & Benhamou 2008 for a discussion of further issues with FPT). Some analyses have used pre-calculated sinuosity indices to reduce the track to bouts of behaviour which are taken as input into a statistical model in order to infer relationships to habitat variables (e.g. Bailleul *et al.* 2007). In general, methods which first classify behaviour before relating these to habitat in a separate analysis faces two problems; first, uncertainty in the behavioural classification is either not calculated (e.g. FPT) or unused (e.g. SSM analyses by Jonsen *et al.* 2007; Shillinger *et al.* 2008); second, habitat data is not able to influence the behavioural categorization despite our expectation that behaviour should be heavily influenced by the environment. Either issue could lead to spuriously precise predictions of the behavioural response to habitat or worse, incorrect inferences about the relationship between habitat and behaviour.

In this paper, we investigate hidden Markov models (HMM; MacDonald & Zucchini 1997) as an alternative method for behavioural analysis of ETT data capable of categorizing movement modes and linking these to physical descriptors of habitat. Our approach is consistent with the fact that neither behaviour nor its drivers are ever directly observed. HMMs are state-space models which assume that the observed distribution of the observations is conditional on a finite number of unobservable or hidden discrete states (Cappe, Moulines & Ryden 2005). While SSM are increasingly being applied to ETT data as they are a natural way to handle location error and behavioural switching (Patterson *et al.* 2008), in this paper, we concentrate on their use in behavioural classification only. We show how HMMs provide a statistically rigorous framework for incorporating covariates, for allowing for the autocorrelation commonly encountered in ETT data, and for making inferences about behavioural states.

Materials and Methods

To introduce HMM, we first develop a simple simulation and then provide details of the HMM machinery applied to the simulated data.

Consider an idealized foraging tuna. Assume that our tuna can be in two possible behavioural states, searching (*S*) and foraging (*F*). At each time step of the simulation a distance y_t is moved. These distances are random variables generated from separate probability functions pertaining to each state. Movement distances were drawn from two separate exponential distributions $P(y_t | \text{state} = j) = \lambda_j e^{-\lambda_j y_t}$ although other distributions could be used. The average distance moved is given by $1/\lambda_j$ so the parameters λ_j were chosen to mimic the situation where a larger distance is expected for searching moves and a smaller distance with foraging moves ($1/\lambda_1 > 1/\lambda_2$). This reflects decreased likelihood of large movements when foraging is successful on a particular patch. Switching between behaviours is a Markov process governed by a 2×2 transition matrix τ,

$$
\begin{aligned}
\tau(X, \theta) &= \begin{pmatrix} F \to F & F \to S \\ S \to F & S \to S \end{pmatrix} \\
&= \begin{pmatrix} P(F|F) & 1 - P(F|F) \\ 1 - P(S|S) & P(S|S) \end{pmatrix} \\
&= P(s_t = j | s_{t-1} = k, X, \theta)
\end{aligned}
\tag{1}
$$

Here, the elements of the transition matrix are $\tau_{k,j} = P(s_t = j | s_{t-1} = k, X, \theta)$, $s_t = j$ denotes the state (*F* or *S*) at time *t*. The variable *X* could be any covariate of interest, but throughout this paper we set this to be a temperature anomaly calculated as daily mean temperature from the tag record subtracted from the global mean across all samples. Hereafter, we denote this anomaly as $\langle X \rangle$ to differentiate it from a generic covariate *X*.

Equation 1 states that probability of switching from one state to another is a function of a covariate *X* and parameters θ. In our model, this is given by $P(s_t = j | s_{t-1} = j) = \text{logit}^{-1}(\alpha_j + \beta_j X)$ and $P(s_t = j | s_{t-1} = k) = 1 - P(s_t = j | s_{t-1} = j)$. The inverse logit transform $e^x / (1 + e^x)$ maps $x \in \{-\infty, +\infty\} \to [0, 1]$ and thereby constrains the transition matrix entries to be probabilities. Temperatures used in simulations were taken from real data from electronic tags implanted in southern bluefin tuna (SBT) (further details follow).

ESTIMATING STATE PROBABILITIES AND LIKELIHOOD CALCULATIONS

The problem we address is estimating the parameters governing the hidden Markov process described above, knowing only the distance moved, y_t and the temperature anomaly $\langle X \rangle_t$. Maximum likelihood estimation was used to estimate the parameters $\theta = (\lambda_1, \lambda_2, \alpha_1, \alpha_2, \beta_1, \beta_2)$. From these and the data we can estimate the probability of the behavioural state.

Foraging and searching behaviours occur in bouts that persist through time. Therefore, we assume that the data are not temporally independent (although see the Supporting Information, Appendix B for alternative models). Therefore, the likelihood is formed from the probabilities of the observed data conditional on the previous state. This means that if we really did know the hidden state we could treat all data pertaining to this state as independent draws from a (in this

case, exponential) probability distribution. As we in fact, do not know the true state we must estimate it from the data.

Let $Y_{1:T}$ be a set of movement observations (in this case distances moved) over $t = 1,...,T$ from two hidden states $s_t \in \{1, 2\}$ which are governed by a Markov process with the transition matrix τ. Let $f(y_{t+1}|Y_t)$ refer to the likelihood of observing y_{t+1} conditional on Y_t, the set of observations of y from the first to the t^{th} time step. The overall likelihood is given by

$$L(Y, \theta) = \prod_{t=0}^{T-1} f(y_{t+1}|Y_{1:t}, \theta) \qquad (2)$$

We show in Supporting Information, Appendix A that, given the first term in the likelihood $f(y_1|Y_0, \theta)$ a recursive scheme can be used to calculate each subsequent term. Using eqn 2, the parameters $\hat{\theta}$ can be estimated by numerical minimization of $-\log[L(Y, \theta)]$. This was performed using quasi-Newton minimization in ʀ (R Development Core Team 2008).

Computation of $\hat{\theta}$ requires only that we compute the distribution of a state at time t, given all previous observations. To compute the full posterior state probability at t, i.e. incorporating observations from t onwards (Wikle & Berliner 2007), we can apply the smoothing algorithm described in Supporting Information, Appendix A.2. States are categorized by choosing the state with the highest posterior (smoothed) probability.

Diagnostics

For a HMM no strict analog to a residual exists as the value of a residual depends on the state (MacDonald & Zucchini 1997), which is unobservable. Pseudo-residuals (described in Supporting Information, Appendix B) were used to assess the goodness-of-fit graphically.

We also used diagnostics to consider two simpler candidate models, both special cases of the two-state HMM. The first hypothesizes that the data are generated from a single behavioural state. Given that an assumption of the HMM is conditional independence of the observations given state, a one-state model is simply the exponential density. This model can be seen to be a special case of the two-state HMM by setting the transition matrix to be one in the first column and zero in the second. The second model considered was a 'memoryless' model (aka. a mixture model – McLachlan & Peel 2000). In this case, the probability of states are a function of temperature but process has no memory of its previous state and each observation is

considered independent. This is a special case of the HMM where the rows of τ are set to be identical (e.g. $F \rightarrow F = S \rightarrow F$). QQ-plots, trends of pseudo-residuals through time and pseudo-residual autocorrelation functions were used to assess the goodness of fits of each model.

STATIONARY DISTRIBUTIONS WITH RESPECT TO COVARIATES

It is not intuitive to see how particular α and β values influence the transition probabilities. One way to summarize their implications for the predicted behaviour is to examine the equilibrium state (stationary distribution) of the estimated Markov process which gives the marginal probability of a state assuming the covariate is fixed at a given value. This requires calculating the dominant eigenvector of the transition matrix τ at some temperature anomaly of interest, e.g. $\langle X \rangle = 2°\text{C}$. Mathematically, we seek the vector p^* such that $p^* = \tau(\langle X \rangle, \hat{\theta})p^*$ where $\sum p^* = 1$. This summarizes the marginal behaviour of the model given a value of the covariate X and estimated parameters $\hat{\theta}$. The delta method (Oehlert, 1992) was used to calculate a confidence interval on this stationary distribution:

$$Var(X) \approx \tau'(X, \hat{\theta})^{\top} Var(\hat{\theta})\tau'(X, \hat{\theta}) \qquad (3)$$

The derivatives $\tau'(X, \hat{\theta})$ were calculated numerically and $Var(\hat{\theta})$ was calculated from the inverse Hessian matrix of the parameters.

Figure 1 shows that for $\alpha = (3,3)$ and $\beta = (-0.4, 1.3)$, the values used in the simulation (see below), the transition probabilities are asymmetrical with respect to the zero temperature anomaly with high probability of switching from state 2 to state 1 at negative temperature anomalies. However, the stationary probabilities of being in state given an anomaly shows a symmetrical preference for state 1 at negative anomalies and state 2 at positive anomalies. This highlights the difference between the probability of a state transition given covariates and the stationary distribution of a state.

SIMULATIONS

It is important to determine the necessary time-series length required for reliable inference of animal behaviour. To investigate this we simulated 100 datasets of 2000 time steps and chose subsets of the data of varying lengths from 100, 200, 300, 400, 700, 1000 and the full time-series of 2000 time steps. Also, if the average move distance in a
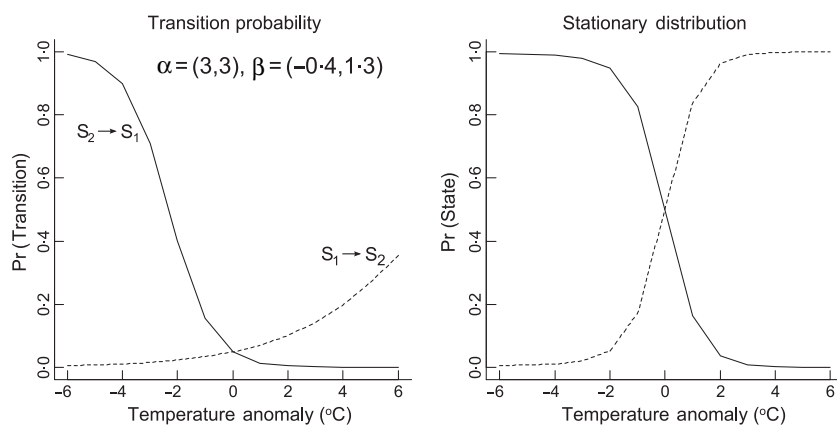
Fig. 1. Left hand panel: Probability of transition from state 2 to state 1 (solid line) or state 1 to 2 (dashed line) as a function of temperature anomaly (SST) given transition matrix parameters α_i and β_i shown. Right hand panel: Probability of being in state 1 (solid line) or state 2 (dashed line) at a given temperature anomaly and the same values of α_i and β_i.

feeding state is similar to average move distance in a searching state accurate estimation may be difficult. Therefore, we also explored the implications of states where the distances moved are similar (see Table 2). One hundred realizations were used in these simulations as we were primarily interested in the average behaviour of estimates.

Results

SIMULATION RESULTS

The true parameter values used in the simulation and the estimated values from a time series of length 2000 are given in Table 1. Parameters are generally well estimated, particularly $1/\lambda$, although we note that the clear difference between the true λ values implies strong contrast in the (simulated) data on distance moved. We explore this further below.

Parameter estimates (Fig. 2) stabilized to values close to the true values at time-series lengths of 300 or more time steps. We explored several other examples with different α, β and λ values, and results for these examples were very similar to those shown in Fig. 2. Generally, estimates of λ remained accurate with less than 300 data points if good starting values were used in the MLE process.

A direct comparison of results for a case with $\lambda = (20, 0.5)$, and a case with $\lambda = (2, 0.5)$ is shown in Fig. 2. True values of α and β were the same for both examples. In both examples, estimates again converged to values close to the true values at time-series lengths of about 300 time steps (Fig. 2). Even for long time series, the mean squared errors of some parameter estimates were higher for the case where the λ values in the two states were more similar. This is not surprising, but illustrates the potential value of simulation trials as part of the process of analysing actual data.

Estimates from the HMM of the most likely state at each time step are of particular interest. Knowing the 'true' state in the simulated data allows model estimates of state to be summarized into a percentage of incorrect state identifications. Table 2 shows results of the mean number of misidentified states for the two simulation examples discussed above. As expected, when λ values were more similar, the percentage misidentified states was higher than when the two λ were further apart. The actual percentage misidentified states was relatively low in these examples. Other cases, for example where the stationary probabilities were more similar for the two states over the range of SST anomalies, lead to higher levels of misidentified states, as shown for Case 2 (Table 2).

Table 1. Summary results of parameter estimates from 100 simulations, each with a time-series of 2000 time steps

|	Truth	Median estimated value	MSE	Relative error	
$1/\lambda_1$	0·05	0·0497	2×10^{-6}	−0·56	$\lambda_1 = 20$
$1/\lambda_2$	2	1·994	0·005	−0·096	$\lambda_2 = 0.5$
α_1	3	2·988	0·004	−0·8	
α_2	3	3·013	0·069	0·85	
β_1	−0·4	−0·417	0·021	6·46	
β_2	1·3	1·324	0·044	3·46	

APPLICATION TO REAL DATA

Background

Having evaluated the model with simulated data, we applied it to data from three juvenile SBT (*Thunnus maccoyyi*, Castelnau, 1872) carrying archival tags. Our test data set is kept small purely for brevity but each track contained roughly data over 1 year (Table 3). We do not aim to provide a thorough analysis of SBT movement and behaviour here but simply to demonstrate the potential of the technique.

Again, we examined horizontal movements in relation to the daily average ambient temperature anomaly, aiming to detect switches in movement behaviour. Juvenile SBT undertake large migrations from the Great Australian Bight (GAB), south of the Australian mainland, into the Indian Ocean (IO) every year (Gunn & Block 2001) and return to the GAB the next summer. The reasons for these migrations are unknown (Gunn & Block 2001; Bestley *et al.* 2008). Movements of SBT appear to indicate resident periods punctuated with episodes of rapid directional movement (Bestley *et al.* 2008), a process possibly driven by oceanographic forcing. Therefore, the following analysis assumes that temperature is a proxy for favourable environmental conditions in some unknown way. While tunas do respond to ocean temperatures (Laurs, Yuen & Johnson 1977; Bertignac, Lehodey & Hampton 1998; Brill *et al.* 1999, 2002), the scenario considered here is a simplification of the ecophysiological processes driving tuna movement and behaviour.

Following the simulation, we assume that the SBT occupy one of two possible states, although here we do not infer foraging behaviour. This is because Bestley *et al.* (2008) examined SBT-feeding behaviour and found that residence periods can sometimes coincide with periods of fasting. Therefore, we simply aim to detect changes in the movement behaviour and label these as either 'resident' or 'migratory'. While this categorization is simple, two modes of movement have been noted in bluefin tuna (Newlands, Lutcavage & Pitcher 2004; Teo *et al.* 2007). We consider data from three fish implanted with archival tags that measured depth, ambient light and internal and external temperature (Wildlife Computers Mk7 model archival tags; Wildlife Computers, Redmond, WA, USA) every 4 min. The fish were tagged as part of CSIRO Marine Research tagging operations in the GAB in the austral summer of 1998 [see Gunn & Block (2001) and Bestley *et al.* (2008) for details]. Position data here are a subset of those used by Bestley *et al.* (2008) and were estimated by combining light-based methods to estimate longitude (Hill & Braun 2001) and matching the onboard temperature sensor measurements to sea surface temperature (Teo *et al.* 2004). Estimates of position are calculated daily and the average daily temperate anomaly is calculated from the 4 min temperature record for each 24 h period. Therefore, the HMM input data consists of $\{t_i, y_i, \langle X \rangle_i\}$ where t_i is the day of the track, y_i is the distance between subsequent position estimates and $\langle X \rangle_i$ is the temperature anomaly.
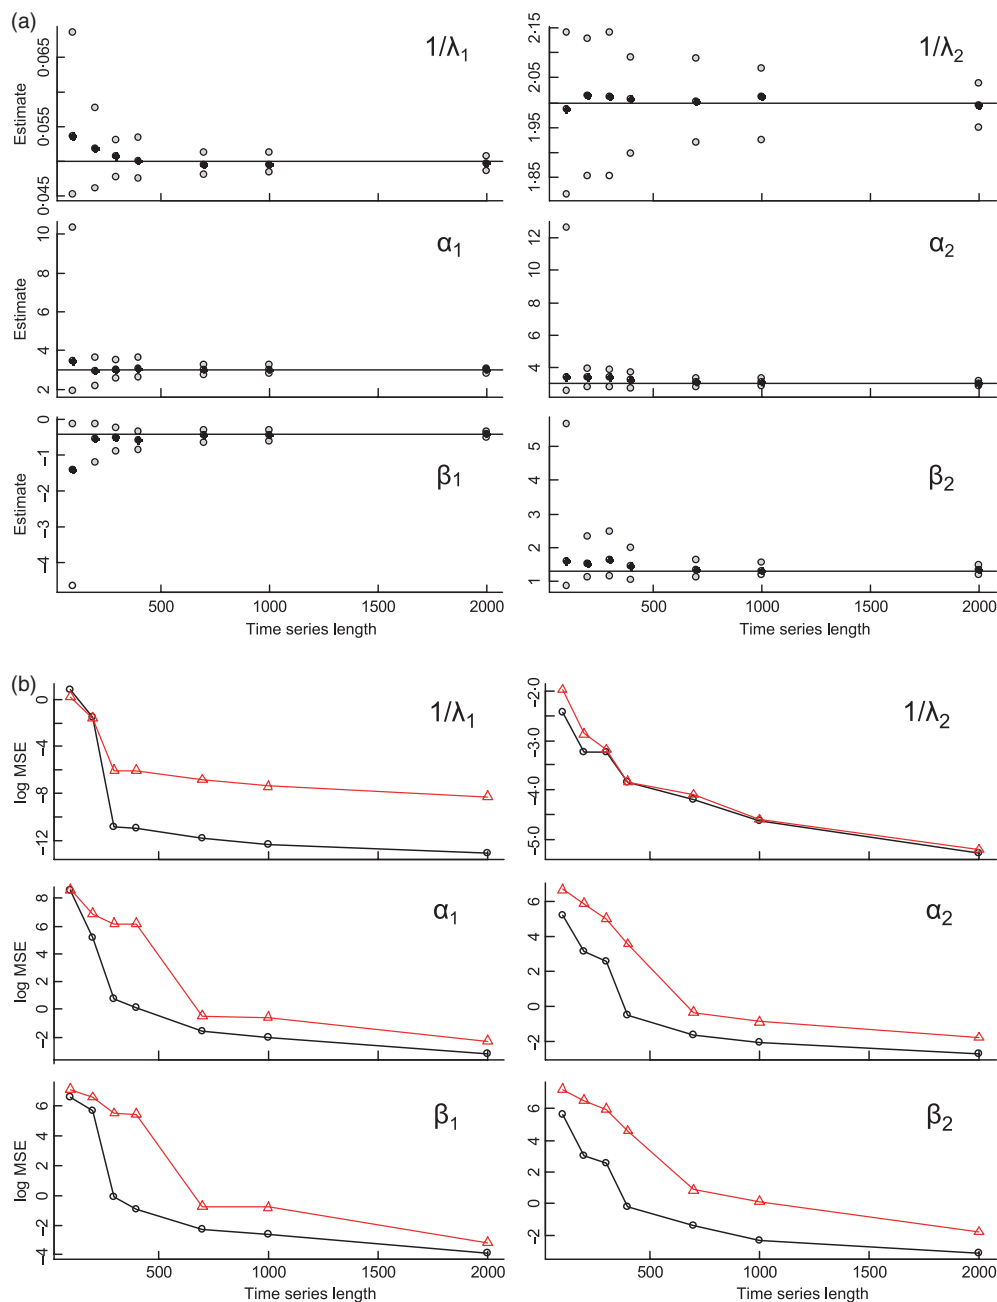
Fig. 2. (a) Parameter estimates for 100 datasets of different time-series lengths (*x*-axis). The horizontal line indicates the true value, solid dots are the median estimates and the grey circles show the lower and upper quantiles of the 100 estimates. Note that the $1/\lambda$ (average distance within state) is shown in the first two panels. (b) Log mean squared error (log MSE) for each of the six estimated parameters as a function of the time-series length of data, and for two examples: black line, open circles has true $1/\lambda_i = (0 \cdot 05, 2)$ and red line, open triangles $1/\lambda_i = (0 \cdot 5, 2)$. True $\alpha_i = (3, 3)$ and $\beta = (-0 \cdot 4, 1 \cdot 3)$ in both examples.

Testing for differences between individuals

We applied standard likelihood ratio methods (Buse 1982) to test for differences between SBTs behaviour with respect to temperature. This is done by comparing the log likelihoods from models fitted to individual data sets, $L_{separate}$, to the log likelihood of a global model fitted to all animals simultaneously, L_{joint}. Then $2\Delta = L_{joint} - L_{separate}$ is asymptotically distributed as $\chi^2_{n(\theta)}$, where $n(\theta)$ is the number of parameters in the global model. Hence individuals are significantly different at p if $2\Delta \geq \chi^2_{n(\theta)}(p)$.

MODEL RESULTS

The diagnostic QQ-plots (Fig. 3) of pseudo residuals (Supporting Information, Appendix B) showed that the models fit reasonably well, although in some states the QQ-plots departed from the 1 : 1 relationship. This might indicate the

Table 2. Mean percentage of mis-identified states for four simulations: two sets of λ values crossed with two sets of α and β values

Scenario	Different movement $1/\lambda = (0{\cdot}05, 2)$	Similar movement $1/\lambda = (0{\cdot}5, 2)$
Case 1		
$\alpha = (3,3), \beta = (-0{\cdot}4, 1{\cdot}3)$		
$N = 300$	≤ 1	6
$N = 2000$	≤ 1	4
Case 2		
$\alpha = (0,0), \beta = (0{\cdot}8, 0{\cdot}8)$		
$N = 300$	$3{\cdot}50$	20
$N = 2000$	$3{\cdot}00$	16

Results are given for time-series lengths of 300 or 2000 time-steps. Case 1 is the same as that shown in Table 1 and Figs 1 and 2. Case 2 has α and β values that imply stationary state probabilities of $0{\cdot}5$ at all anomaly values.

Table 3. Parameter estimates and standard errors (SE) estimated from three southern bluefin tuna

SBT ID (Dataset length)	State	$\hat{\alpha}$	SE	$\hat{\beta}$	SE	$\hat{\lambda}^{-1}$	SE
Individual models							
99267	Resident	3·06	0·959	0·413	0·516	38·3	0·00256
(354 d)	Migratory	2·95	1·19	−1·26	0·803	88·1	0·00117
98007	Resident	2·29	0·492	1·87	2·06	48·8	0·00397
(385 d)	Migratory	1·31	1·31	−0·483	5·56	82·8	0·00219
99629	Resident	−2·86	5·5	3·21	2·79	43·1	0·00231
(263 d)	Migratory	3·21	4·14	−0·722	1·09	68·9	0·00121
Global model							
All	Resident	2·622	0·638	3·263	1·54	43·9	0·000719
	Migratory	0·48	0·356	−2·0395	2·18	79·37	0·00147

The parameters α and β are used in forming the transition matrix and the λ are the rate parameter of an exponential distribution used to model movements. The table shows estimated parameters from fitting to the individual data sets separately (labelled 'Individual Models') and when using all data simultaneously to estimate the parameters (labelled 'Global Model'). Also shown in parentheses in the left hand column is the length in days of each data series.

need for another state to be incorporated into the models, or that the exponential model of movements was at times inadequate. Diagnostics for the one-state model and the memoryless two-state model indicated sizable trends and significant autocorrelation in the residuals (see Supporting Information, Appendix B). The diagnostics for the two-state HMM were much improved with no apparent trend in the residuals and much reduced auto correlation. Hence, we consider the two-state HMM to be the best of the models considered and that simpler models were unable to model the data adequately.

The HMM predicted quite different average movement distances in each state. However, estimates of these $(1/\lambda)$ were quite similar between the three fish considered. In the resident state, SBT moved between 38 and 49 nm day^{-1} (nm = nautical mile, 1 nm = 1·852 km) and in the migratory state from 69 to 88 nm day^{-1} (Table 3). The time spent in each state varied between the tuna (Fig. 4); 59·1% (99267),

66% (98007) and 35% (99629). By using the posterior probability of being in a state as an indicator that a tuna is in the resident or searching state, we mapped the positions where states occurred (Fig. 4). For example, the HMM inferred that SBT 97627 was in resident mode in the GAB and also around the boundary of the Northern Subtropical Front (Kostianoy *et al.* 2004). This tuna was also categorized as resident in the area to the southwest of Western Australia.

A likelihood-ratio test comparing a global model (shared parameters between SBT compared to individual parameters for each fish, Table 3) to the individual model fits found no significant differences between individuals ($\chi^2_{n(\theta)=6} = 9 \cdot 791$, $p = 0 \cdot 866$). Under the assumption of shared parameters between all fish, the SBT were expected to move 79 km day^{-1} in whilst in a migratory mode and 43·9 km day^{-1} in a resident mode. Additionally, pairwise comparisons between tuna also found no significant differences. This result was qualitatively supported by the stationary distribution of the transition matrix with respect to the mean temperature anomaly. These predict similar responses from each fish with SBTs 98007 and 99267 predicted to flip state about the mean ambient temperature anomaly (i.e. $\langle X \rangle = 0$) (Fig. 5). SBT 99629 varied slightly, with the inflection point in the ogive occurring at a temperature anomaly around −3 °C. Note that the confidence intervals on these fits are wide, indicating considerable uncertainty in these results. This is likely to reflect the simplicity of the model and may indicate that other covariates are required to better capture the relationship between movement and temperature.

Discussion

STATISTICALLY LINKING MOVEMENT AND HABITAT DESCRIPTORS

The relationship between movement and habitat is shaped by factors, such as foraging success or physiological constraints placed by the environment. As these factors mediate animals' ability to survive and reproduce, determining the relationship between movement and environment is a critical part of understanding population processes (Alerstam, Hedenstrom & Akesson 2003; Bowler & Benton 2005) and spatial pattern (Morales & Ellner 2002; Del Mar & Penteriani 2008). This paper has demonstrated a statistical method for categorizing behavioural modes in individual movement data. Importantly, the model we describe relates the behaviour directly to covariate data – the relationship between partially observed or hidden behaviours and descriptors of habitat is estimated within the model. This avoids the need for *post hoc* analysis (Shillinger *et al.* 2008) to examine the relationship with covariates. Thus, direct estimation of the statistical relationship between movement and habitat is achievable.

The description of movement we have used is simple. However, it is biologically meaningful and provides a tractable way to deal with a highly complex problem. Quantitative investigation of individual movement requires a reasonable (albeit often simple) model of the interaction between
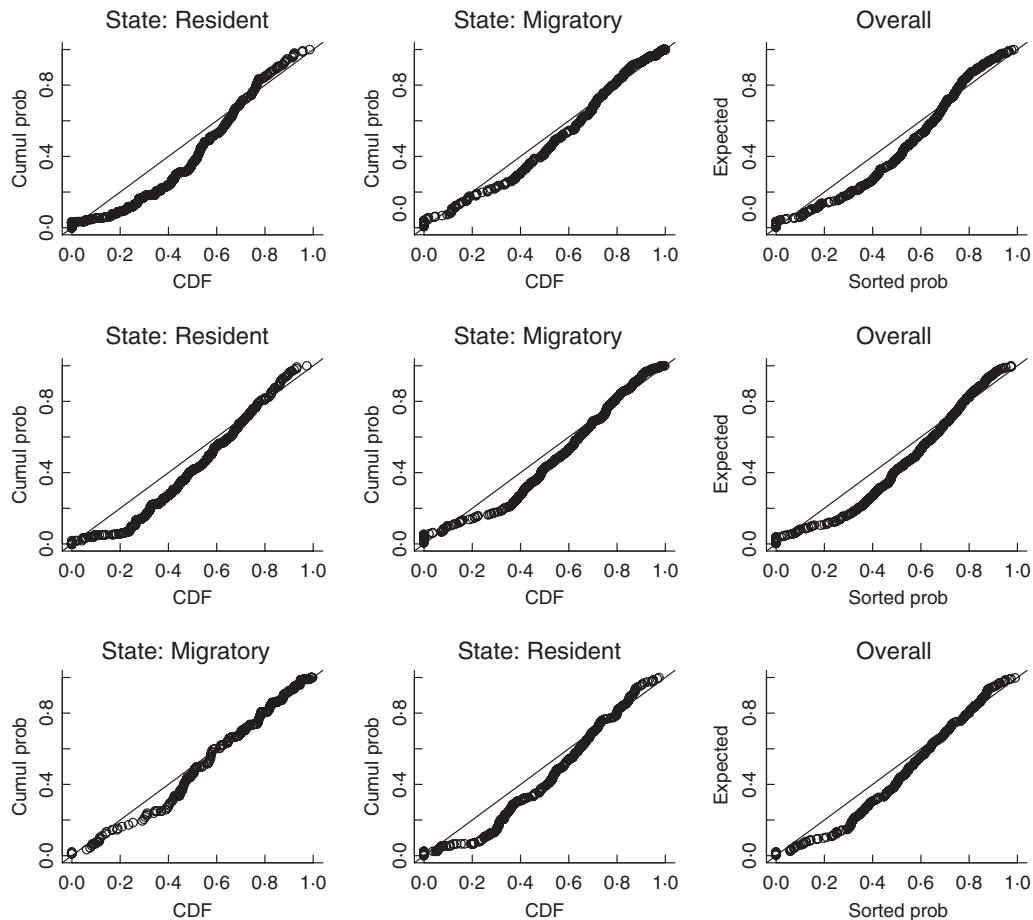
Fig. 3. Diagnostic plots for hidden Markov models fitted to the three southern bluefin tuna; (top) 99629, (middle) 98007, (bottom) 99267. QQ-plots for state 1 (resident), state 2 (migratory) and overall model fits are given.

behaviour and habitat which may be estimated from data (Patterson *et al*. 2008). While a specific hypothesis will often be implied by the model, HMMs are useful for dealing with unmodelled or unmeasured processes – which are all pervasive in movement analysis. For example, the factors driving the ocean-basin scale, cyclic migrations shown by juvenile-SBT are largely unknown (Gunn & Block 2001). Similarly, which factors influence residency behaviour is poorly understood, although for non-breeding animals, growth through productive foraging is expected to be the primary impetus (Bestley *et al*. 2008).

Extensions to our model could include improved descriptions of the ocean state, individual data from the SBT (e.g. size or year class) or productivity data (e.g. Gremillet *et al*. 2008). In addition, SBT are unusual in that direct indices of feeding behaviour can be gained from archival tag data (Gunn, Hartog & Rough 2001). A logical next step would be to apply these models to a larger data set to develop more precise inferences about juvenile SBT movement phases to compare against empirical measures of feeding activity. However, for most species, direct observations of feeding are unavailable and models such as ours are required to detect and categorize behaviour. Moreover, a suite of covariates at multiple spatial and temporal scales are likely to influence

prey distributions and therefore most likely movement behaviour. Untangling these multiple interactions is challenging. A strength of the approach we describe is the capacity for quantitative inference linking movement and observable habitat variables, even though they may only indirectly influence behaviour via a set of unobserved processes. This may aid in better understanding the direct biological drivers of movement for many species, especially in marine systems where observations of the biological components of the system (e.g. forage distributions) are particularly sparse.

SIMULATION TESTING

By conducting simulation tests, we demonstrated how the length of the data series affects parameter estimation. Generally, the parameters were well estimated for data series longer than ~300 time steps. However, this varied with the contrast in the average distance moved within state. This gives a practical insight into the duration of time series required for effective inference and therefore which ETT data sets and methods are amenable to HMM (and possibly other SSM) analysis. Also, our results indicate that useful results can be obtained for shorter data sets but with greater uncertainty.
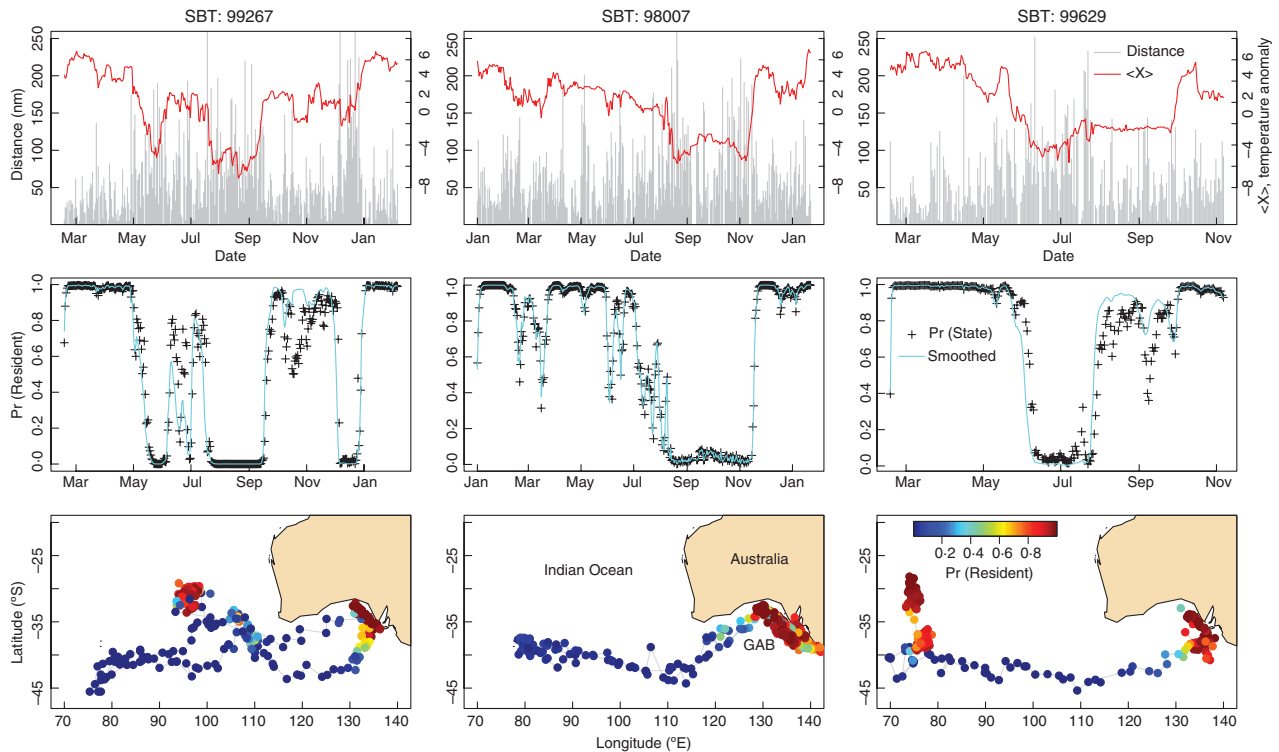
Fig. 4. Input data and state predictions from the hidden Markov model. Each column shows the result for a particular southern bluefin tuna. Left column: the line shows the smoothed probability of state. Top row: daily distance moved and mean temperature anomaly, $<X>$ (right axis). Middle row: categorization of most likely state for three southern bluefin tuna along their tracked movements. Bottom row: State categorizations mapped onto movement paths. The colour of the points mapped on the locations shows smoothed P(Resident) i.e. Red indicates higher probability of the resident state relative to blue.

This is not surprising as detecting state switches is more difficult if each state 'looks' similar.

There was also an interaction between the average distance moved ($1/\lambda$) and the transition matrix parameters (α_j and β_j) and estimation was less accurate when the probability of switching was uniform across the range of the covariate. Understanding interactions between the implications of the model and parameter estimability can be complex, even for simple models such as this. Simulation testing of models assists in understanding how models behave and thus influence biological inferences. To our knowledge, this analysis of the interaction between the data series length and the reliability of estimation has not been considered elsewhere. As the modelling approaches to tracking and bio-logging data increase in complexity, we suggest that simulation become a routine analysis tool.

DIAGNOSTICS AND INFERENCE

A key part of the material we present here are the diagnostic techniques (Supporting Information, Appendix B). Diagnostics are an integral part of statistical modelling, but are lacking in many analyses of animal movement with SSMs. Morales *et al.* (2004) used the Deviance Information Criterion to compare between models. This and other information theoretic approaches are certainly useful, but still rely on the likelihood coming from a well-fitting model, making the use

of residual plots etc. very useful. Those produced here suggested reasonable fit to the data but possibly indicates the need for more states or an error distribution other than the exponential. Nonetheless, a two-state HMM was preferred over simpler models. While a state switching model has greater utility than a one-state model, and thus may be preferred *a priori*, it is worth determining if a more parsimonious single-state model could explain the data equally well. That the alternative models fitted poorly and failed to model both long and short term autocorrelation informs us of persistence in the observed movement.

The inference methods we have presented are limited and multiple comparisons are best avoided for large numbers of animals due to increased type-I error rate. Hierarchical/mixed-modelling (Jonsen, Myers & Flemming 2003) frameworks are a more powerful approach to individual heterogeneity. However, the likelihood ratio test could be used to compare between groups of animals with shared parameters. Deciding which animals to group together may not be straightforward – but sound biological hypotheses often entail groupings (e.g. animals from the same population/age/sex) where likelihood ratio methods may be appropriate. Recent Bayesian methods have demonstrated automatic groupings of individuals with mark–recapture data (Dorazio *et al.* 2008). Information theoretic criterion such as Akiake's information criterion (Burnham & Andersen 2003) may also be useful for covariate selection.
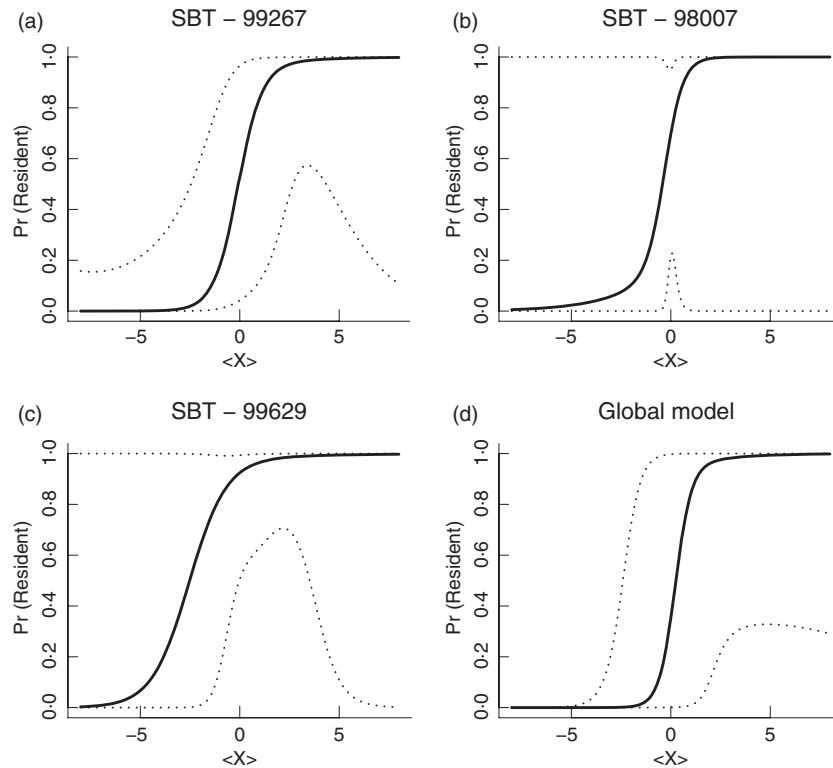
Fig. 5. Stationary distributions of the state transition matrix formed from the estimates of α and β (see Table 3) at varying temperature anomaly values ranging from −8 to 8 °C. Black lines are the expected probability and the dashed lines are 95% confidence intervals. (a–c) Results using data from individual tuna. (d) Results from shared parameters estimated from all three data sets simultaneously.

COMPARISONS TO RELATED APPROACHES

The treatment of temporal dependence in movement data marks a major distinction between SSM and other movement analysis approaches (Patterson *et al.* 2008). For example Lévy analyses (e.g. Viswanathan *et al.* 1996) mostly ignore short-term temporal correlation potentially leading to spuriously precise estimates of parameters (see also Edwards *et al.* 2007 for a discussion of the statistical issues with Levy approaches). The state switching behaviour of HMMs is not necessarily first order Markovian (MacDonald & Zucchini 1997), since using covariates to mediate switching between states, dependence relates time in a behavioural mode to other factors besides the previous state. This is important when modelling movements where states are expected to persist through time. Several SSMs have estimated a transition matrix which is independent of covariates (Jonsen, Mills-Flemming & Myers 2005; Bailey *et al.* 2008). Determining the resulting differences between these models is an important step required for further work.

A novel aspect of our paper is the characterization of behavioural switching with respect to covariates via the stationary distribution (Fig. 5). The simulation study found that the nature of the stationary distribution with respect to the covariate (temperature anomaly) was linked to parameter estimability. We note that the stationary distribution should not be interpreted too literally. It is not clear what biological interpretation should not be attached

to the quantity as animals' behaviour may be unlikely to be in equilibrium. Nevertheless, the technique can be convenient for summarizing covariate effects on estimated switching probabilities.

Several studies have used statistical approaches similar to the methods outlined here (Jonsen *et al.* 2003; Franke, Caelli & Hudson 2004; Morales *et al.* 2004; Roberts *et al.* 2004; Royer, Fromentin & Gaspar 2005) The closest to our methods are from the examples of Franke *et al.* (2004) using HMMs to examine the movements of Caribou and analysis of elk by Morales *et al.* (2004).

There are several technical differences between our approach and these studies. Franke *et al.* used the expectation-maximization algorithm for parameter estimation (MacDonald & Zucchini 1997). Roberts *et al.* used variational inference to categorize areas of high positional entropy in the track. Morales *et al.* (2004) and Jonsen *et al.* (2003) used Markov Chain Monte Carlo (MCMC). A major advantage of the direct likelihood maximization approach we used is speed – models were fit within a few seconds. Additionally, our approach is not simply faster, it is also reliable. While MCMC is without doubt an extremely powerful tool in expert hands, it is hard to detect problems (e.g. Carlin & Louis 1997, p177). Computation time can be a serious impediment to fitting a variety of models (for model selection), for verifying performance (through simulation testing) or when analysing very large data sets. For these reasons, it is of great practical value to have quick and reliable estimation

algorithms such as quasi-Newton maximization (used here), particularly for users without extensive statistical expertise.

Note that the main issue here is one of practicality and potential reliability, rather than of Bayesian vs. non-Bayesian philosophy; state-space models are intrinsically Bayesian in nature because of the unobservable random states, and the issue of whether to handle the system parameters via prior distributions (full Bayes) or via MLE (empirical Bayes, as in this paper) is of secondary importance in systems with reasonable amounts of data. However, using MLE methods may also practically restrict the complexity of HMMs and some problems may only be tractable with MCMC methods. For instance random effects or hierarchical model structures may be easiest to implement using Bayesian MCMC.

A further advantage of the Bayesian state-space approach (Jonsen *et al.* 2003; Andersen *et al.* 2007; Patterson *et al.* 2008) is the ability to incorporate estimation of location error – although often location error is estimated prior to estimating movement (see Pedersen *et al.* 2008 for a likelihood based for an approach to location uncertainty with HMMs). While the exponential distribution we used may be somewhat robust to observation error, in this study we have focused on models that examine biological process rather than estimate likely position. It is likely that the current need to focus on correcting location error (Vincent *et al.* 2002; Teo *et al.* 2004; Royer & Lutcavage 2008) will lessen as technological advances are made. Wider use of the global positioning system is likely to reduce location error (Schofield *et al.* 2007). However in some cases, (e.g. tracking pelagic fish), it is unlikely that the need to deal with observation error will disappear. Nevertheless, we expect the focus to eventually shift from estimating where an animal was to what it was doing and what drove its movement.

Conclusions

Despite several recent publications detailing analogous statistical approaches, the application of HMM and related models to telemetry data is in its infancy and much further developmental work is required (Patterson *et al.* 2008). We have identified several key areas to be pursued. Some of these are straightforward such as increasing the number of covariates or allowing movement distribution parameters to also be a function of covariates. Similarly, the number of states could be increased. While each is a relatively simple extension of the simple models demonstrated they entail estimation of many more parameters.

Further development and operationalization of the HMM and allied methods will provide researchers with a robust statistical framework for the analysis of telemetry data and investigation of the processes driving the movement and behaviour. This paper has described in detail a useful model for determining the relationship between behavioural state-switching and environmental data. Building such integrative models is a necessary step in being able to elucidate the links between behaviour and the environment in free-ranging animals.

References

Aarts, G., MacKenzie, M., McConnell, B., Fedak, M. & Matthiopoulos, J. (2008) Estimating space-use and habitat preference from wildlife telemetry data. *Ecography*, **31**, 140–160.

Alerstam, T., Hedenstrom, A. & Akesson, S. (2003) Long-distance migration: evolution and determinants. *Oikos*, **103**, 247–260.

Andersen, K.H., Nielsen, A., Thygesen, U.H., Hinrichsen, H.H. & Neuenfeldt, S. (2007) Using the particle filter to geolocate Atlantic cod (*Gadus morhua*) in the Baltic Sea, with special emphasis on determining uncertainty. *Canadian Journal of Fisheries and Aquatic Sciences*, **64**, 618–627.

Austin, D., Bowen, W.D., McMillan, J.I. & Boness, D.J. (2006) Stomach temperature telemetry reveals temporal patterns of foraging success in a free-ranging marine mammal. *Journal of Animal Ecology*, **75**, 408–420.

Bailey, H., Shillinger, G., Palacios, D., Bograd, S., Spotila, J., Paladino, F. & Block, B.A. (2008) Identifying and comparing phases of movement by leatherback turtles using state-space models. *Journal of Experimental Biology and Ecology*, **356**, 128–135.

Bailleul, F., Charrassin, J.B., Monestiez, P., Roquet, F., Biuw, M. & Guinet, C. (2007) Successful foraging zones of southern elephant seals from the Kerguelen Islands in relation to oceanographic conditions. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, **362**, 2169–2181.

Barraquand, F. & Benhamou, S. (2008) Animal movements in heterogeneous landscapes: identifying profitable places and homogeneous movement bouts. *Ecology*, **89**, 3336–3348.

Bertignac, M., Lehodey, P. & Hampton, J. (1998) A spatial population dynamics simulation model of tropical tunas using a habitat index based on environmental parameters. *Fisheries Oceanography*, **7**, 326–334.

Bestley, S., Patterson, T.A., Hindell, M.A. & Gunn, J.S. (2008) Feeding ecology of wild migratory tunas revealed by archival tag records of visceral warming. *Journal of Animal Ecology*, **77**, 1223–1233.

Biuw, M., McConnell, B., Bradshaw, C.J.A., Burton, H. & Fedak, M. (2003) Blubber and buoyancy: monitoring the body condition of free-ranging seals using simple dive characteristics. *Journal of Experimental Biology*, **206**, 3405–3423.

Bowler, D.E. & Benton, T.G. (2005) Causes and consequences of animal dispersal strategies: relating individual behaviour to spatial dynamics. *Biological Reviews*, **80**, 205–225.

Brill, R.W., Block, B.A., Boggs, C.H., Bigelow, K.A., Freund, E.V. & Marcinek, D.J. (1999) Horizontal movements and depth distribution of large adult yellowfin tuna (*Thunnus albacares*) near the Hawaiian Islands, recorded using ultrasonic telemetry: implications for physiological ecology of pelagic fishes. *Marine Biology*, **133**, 395–408.

Brill, R., Lutcavage, M., Metzgerm, G., Bushnell, P., Arendt, M., Lucy, J., Watson, C. & Foley, D. (2002) Horizontal and vertical movements of juvenile bluefin tuna *Thunnus thynnus* in relation to oceanographic conditions of the western North Atlantic, determined with ultrasonic telemetry. *Fishery Bulletin*, **100**, 155–167.

Burnham, K.P. & Andersen, D. (2003) *Model Selection and Multi-Model Inference*. Springer, New York.

Buse, A. (1982) The likelihood ratio, Wald, and Lagrange multiplier tests: an expository note. *The American Statistician*, **36**, 153–157.

Cappe, O., Moulines, E. & Ryden, T. (2005) *Inference in Hidden Markov Models*. Springer, New York.

Carlin, B.P. & Louis, T.A. (1997) Bayes and empirical Bayes methods for data analysis. *Statistics and Computing*, **7**, 153–154.

Charnov, E. (1976) Optimal foraging: the marginal value theorem. *Theoretical Population Biology*, **9**, 129–136.

Del Mar, D.M. & Penteriani, V. (2008) Behavioral states help translate dispersal movements into spatial distribution patterns of floaters. *The American Naturalist*, **172**, 475–485.

Dorazio, R.M., Mukherjee, B., Zhang, L., Ghosh, M., Jelks, H.L. & Jordan, F. (2008) Modeling unobserved sources of heterogeneity in animal abundance using a Dirichlet process prior. *Biometrics*, **64**, 635–644.

Edwards, A.M., Phillips, R.A., Watkins, N.W., Freeman, M.P., Murphy, E.J., Afanasyev, V., Buldyrev, S.V., da Luz, M.G., Raposo, E.P., Stanley, H.E. & Viswanathan, G.M. (2007) Revisiting Levy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature*, **449**, 1044–1048.

Fauchald, P. & Tveraa, T. (2003) Using First-passage time in the analysis of area-restricted search and habitat selection. *Ecology*, **84**, 282–288.

Franke, A., Caelli, T. & Hudson, R.J. (2004) Analysis of movements and behaviour of caribou (*Rangifer tarandus*) using hidden Markov models. *Ecological Modelling*, **173**, 259–270.

Gremillet, D., Lewis, S., Drapeau, L., van Der Lingen, C.D., Huggett, J.A., Coetzee, J.C., Verheye, H.M., Daunt, F., Wanless, S. & Ryan, P.G. (2008) Spatial match-mismatch in the Benguela upwelling zone: should we expect chlorophyll and sea-surface temperature to predict marine predator distributions? *Journal of Applied Ecology*, **45**, 610–621.

Gunn, J.S. & Block, B.A. (2001). Tuna ecology. *Physiology and Evolution* (eds B.A. Block & E.D. Stevens), Vol. 19, pp. 167–224. Academic Press, London.

Gunn, J., Hartog, J. & Rough, K. (2001) The relationship between food intake and visceral warming in southern bluefin tuna (*Thunnus maccoyii*). In *Electronic Tagging and Tracking in Marine Fisheries* (eds J.R. Sibert & J.L. Nielsen), pp. 109–130. Kluwer Academic Publishers, Dordrecht.

Haydon, D.T. (2008) Cross-disciplinary demands of multihost pathogens. *Journal of Animal Ecology*, **77**, 1079–1081.

Hill, R.D. & Braun, M.J. (2001). Geolocation by light level. *Electronic Tagging and Tracking in Marine Fisheries Research: Methods and Technologies* (eds J.R. Sibert & J. Nielson), pp. 315–330. Kluwer Academic Press, Dordrecht.

Hindell, M.A., Bradshaw, C.J., M.D, S., Michael, K.J. & Burton, H.R. (2003) Dispersal of female southern elephant seals and their prey consumption during the austral summer: relevance to management and oceanographic zones. *Journal of Applied Ecology*, **43**, 703–715.

Hooker, S.K., Biuw, M., McConnell, B.J., Miller, P.J. & Sparling, C.E. (2007) Bio-logging science: logging and relaying physical and biological data using animal-attached tags. *Deep-Sea Research Part II-Topical Studies in Oceanography*, **54**, 177–182.

Hyrenbach, K.D., Keiper, C., Allen, S.G., Ainley, D.G. & Anderson, D.J. (2006) Use of marine sanctuaries by far-ranging predators: commuting flights to the California Current System by breeding Hawaiian albatrosses. *Fisheries Oceanography*, **15**, 95–103.

Jonsen, I.D., Myers, R.A. & Flemming, J.M. (2003) Meta-analysis of animal movement using state-space models. *Ecology*, **84**, 3055–3063.

Jonsen, I.D., Mills-Flemming, J. & Myers, R.A. (2005) Robust state-space modelling of animal movement data. *Ecology*, **86**, 2874–2880.

Jonsen, I.D., Myers, R.A. & James, M.C. (2007) Identifying leatherback turtle foraging behaviour from satellite telemetry using a switching state-space model. *Marine Ecology-Progress Series*, **337**, 255–264.

Kareiva, P. & Odell, G. (1987) Swarms of predators exhibit" preytaxis" if individual predators use area-restricted search. *The American Naturalist*, **130**, 233–270.

Kostianoy, A.G., Ginzburg, A.I., Frankignoulle, M. & Delille, B. (2004) Fronts in the Southern Indian Ocean as inferred from satellite sea surface temperature data. *Journal of Marine Systems*, **45**, 55–73.

Laurs, R.M., Yuen, H.S.H. & Johnson, J.H. (1977) Small-scale movements of albacore (*Thunnus alalunga*), in relation to ocean features as indicated by ultrasonic tracking and oceanographic sampling. *Fishery Bulletin*, **75**, 347–355.

MacDonald, I.L. & Zucchini, W. (1997) *Hidden Markov and Other Models for Discrete Valued Time Series*. Chapman & Hall, London.

McLachlan, G. & Peel, D. (2000) *Finite Mixture Models*. John Wiley & Sons, New York.

Morales, J.M. & Ellner, S.P. (2002) Scaling up animal movements in heterogeneous landscapes: the importance of behavior. *Ecology*, **83**, 2240–2247.

Morales, J.M., Haydon, D.T., Frair, J., Holsinger, K.E. & Fryxell, J.M. (2004) Extracting more out of relocation data: building movement models as mixtures of random walks. *Ecology*, **85**, 2436–2445.

Nathan, R., Getz, W.M., Revilla, E., Holyoak, M., Kadmon, R., Saltz, D. & Smouse, P.E. (2008) A movement ecology paradigm for unifying organismal movement research. *Proceedings of the National Academy of Sciences of USA*, **105**, 19052.

Newlands, N.K., Lutcavage, M.E. & Pitcher, T. (2004) Analysis of foraging movements of Atlantic bluefin tuna (*Thunnus thynnus*): individuals switch between two models of search behaviour. *Population Ecology*, **46**, 39–53.

Oehlert, G.W. (1992) A note on the delta method. *The American Statistication*, **46**, 2729.

Patterson, T.A., Thomas, L., Wilcox, C., Ovaskainen, O. & Matthiopoulos, J. (2008) State–space models of individual animal movement. *Trends in Ecology & Evolution*, **23**, 87–94.

Pedersen, M.W., Righton, D., Thygesen, U.H., Andersen, K.H. & Madsen, H. (2008) Geolocation of North Sea cod (*Gadus morhua*) using hidden Markov models and behavioural switching. *Canadian Journal of Fisheries and Aquatic Sciences*, **65**, 2367–2377.

R Development Core Team (2008) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Roberts, S., Guilford, T., Rezek, I. & Biro, D. (2004) Positional entropy during pidgeon homing I: application of Bayesian latent state modelling. *Journal of Theoretical Biology*, **227**, 39–50.

Royer, F. & Lutcavage, M. (2008) Filtering and interpreting location errors in satellite telemetry of marine animals. *Journal of Experimental Marine Biology and Ecology*, **359**, 1–10.

Royer, F., Fromentin, J.M. & Gaspar, P. (2005) A state-space model to derive bluefin tuna movement and habitat from archival tags. *Oikos*, **109**, 473–484.

Schofield, G., Bishop, C.M., MacLean, G., Brown, P., Baker, M., Katselidis, K.A., Dimopoulos, P., Pantis, J.D. & Hays, G.C. (2007) Novel GPS tracking of sea turtles as a tool for conservation management. *Journal of Experimental Marine Biology and Ecology*, **347**, 58–68.

Shillinger, G.L., Palacios, D.M., Bailey, H., Bograd, S.J., Swithenbank, A.M., Gaspar, P., Wallace, B.P., Spotila, J.R., Paladino, F.V., Piedra, R. *et al.* (2008) Persistent leatherback turtle migrations present opportunities for conservation. *PLoS Biology*, **6**, e171.

Sims, D.W., Southall, E.J., Humphries, N.E., Hays, G.C., Bradshaw, C.J., Pitchford, J.W., James, A., Ahmed, M.Z., Brierley, A.S. & Hindell, M.A. & others (2008) Scaling laws of marine predator search behaviour. *Nature*, **451**, 1098–1102.

Teo, S.L.H., Boustany, A., Blackwell, S., Walli, A., Weng, K.C. & Block, B.A. (2004) Validation of geolocation estimates based on light level and sea surface temperature from electronic tags. *Marine Ecology Progress Series*, **283**, 81–98.

Teo, S.L.H., Boustany, A.M. & Block, B.A. (2007) Oceanographic preferences of Atlantic bluefin tuna, *Thunnus thynnus*, on their Gulf of Mexico breeding grounds. *Marine Biology*, **152**, 1105–1119.

Vincent, C., McConnell, B., Ridoux, V. & Fedak, M. (2002) 'Assessment of Argos location accuracy from satellite tags deployed on captive gray seals. *Marine Mammal Science*, **18**, 156–166.

Viswanathan, G.M., Afanasyev, V., Buldyrev, S.V., Murphy, E.J., Prince, P.A. & Stanley, H.E. (1996) Levy flight search patterns of wandering albatrosses. *Nature*, **381**, 413–415.

Wikle, C.K. & Berliner, L.M. (2007) A Bayesian tutorial for data assimilation. *Physica D-Nonlinear Phenomena*, **230**, 1–16.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix S1. Patterson *et al.*, classifying movement behaviour in relation to environmental conditions using hidden Markov models.

Please note: Wiley-Blackwell are not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.