

momentuHMM: R package for analysis of
telemetry data using generalized multivariate
hidden Markov models of animal movement

Brett T. McClintock¹ and Théo Michelot²

¹Marine Mammal Laboratory
Alaska Fisheries Science Center
NOAA National Marine Fisheries Service
Seattle, U.S.A.

Email: brett.mcclintock@noaa.gov

²School of Mathematics and Statistics
University of Sheffield
Sheffield, U.K.

RUNNING HEAD: R package **momentuHMM**

June 16, 2017

Summary

1. Discrete-time hidden Markov models (HMMs) have become an immensely popular tool for inferring latent animal behaviors from telemetry data, largely because they are relatively fast and easy to implement when data streams are observed without error and at regular time intervals. While movement HMMs typically rely solely on location data, auxiliary biotelemetry and environmental data are powerful and readily-available resources for incorporating much more behavioral realism and inferring ecological relationships that would otherwise be difficult or impossible to infer from location data alone. However, there is a paucity of generalized user-friendly software available for implementing (multivariate) HMMs of animal movement. Furthermore, location measurement error, temporal irregularity, and other forms of missing data are often pervasive in telemetry studies (particularly in marine systems).

2. Here we introduce an open-source R package, `momentuHMM`, that addresses many of the deficiencies in existing software. Features for multivariate HMMs in `momentuHMM` include: 1) tools for data pre-processing and visualization; 2) user-specified probability distributions for an unlimited number of data streams, such as those based on location (e.g., step length, turning angle) and auxiliary biotelemetry data (e.g., from pressure, conductivity, heart rate, or motion sensors); 3) user-specified design matrices and constraints for covariate modelling of state transition probability and probability distribution parameters using linear model formulas familiar to most R users; 4) multiple imputation methods that account for observation error attributable to measurement error and temporally-irregular or missing data; 5) seamless integration of spatio-temporal environmental covariate data (e.g., wind direction, forest cover, sea ice concentration) using rasters; 6) incorporation of “activity center” effects on parameters (e.g., areas associated with attractive or repulsive forces); 7) circular-circular regression models for angular probability distributions; 8) cosinor and spline regression formulas for cyclical (e.g., daily, seasonal) and other complicated patterns; and 9) data simulation capabilities for power analyses and assessing model performance, including simulation of location data subject to temporal irregularity and/or measurement error.

3. After providing a brief introduction to (multivariate) HMMs for telemetry data, we demonstrate some of the capabilities of `momentuHMM` using real-world examples. This brief tutorial includes workflows for data formatting, model specification, model fitting, and diagnostics.

4. While many of the features of `momentuHMM` were motivated by animal movement data, the package can be used for analyzing any type of data that is amenable to (multivariate) HMMs. Practitioners interested in additional features for `momentuHMM` are encouraged to contact the authors.

Key-words biologging, biotelemetry, `crawl`, `moveHMM`, state-space model, state-switching

1 Introduction

Discrete-time hidden Markov models (HMMs) have become immensely popular for the analysis of animal telemetry data (e.g. Morales *et al.* 2004; Jonsen *et al.* 2005; Langrock *et al.* 2012; McClintock *et al.* 2012). In short, an HMM is a time series model composed of a (possibly multivariate) observation process $(\mathbf{Z}_1, \dots, \mathbf{Z}_T)$, in which each data stream is generated by N state-dependent probability distributions, and where the unobservable (hidden) state sequence $(S_t \in \{1, \dots, N\}, t = 1, \dots, T)$ is assumed to be a Markov chain. The state sequence of the Markov chain is governed by (typically first-order) state transition probabilities, $\gamma_{ij}^{(t)} = \Pr(S_{t+1} = j \mid S_t = i)$ for $i, j = 1, \dots, N$, and an initial distribution $\boldsymbol{\delta}^{(0)}$. The likelihood of an HMM can be succinctly expressed using the forward algorithm:

$$\mathcal{L} = \boldsymbol{\delta}^{(0)} \boldsymbol{\Gamma}^{(1)} \mathbf{P}(\mathbf{z}_1) \boldsymbol{\Gamma}^{(2)} \mathbf{P}(\mathbf{z}_2) \boldsymbol{\Gamma}^{(3)} \dots \boldsymbol{\Gamma}^{(T-1)} \mathbf{P}(\mathbf{z}_{T-1}) \boldsymbol{\Gamma}^{(T)} \mathbf{P}(\mathbf{z}_T) \mathbf{1}^N \quad (1)$$

where $\boldsymbol{\Gamma}^{(t)} = \left(\gamma_{ij}^{(t)} \right)$ is the $N \times N$ transition probability matrix, $\mathbf{P}(\mathbf{z}_t) = \text{diag}(p_1(\mathbf{z}_t), \dots, p_N(\mathbf{z}_t))$, $p_s(\mathbf{z}_t)$ is the conditional probability density of \mathbf{Z}_t given $S_t = s$, and $\mathbf{1}^N$ is a N -vector of ones (for a thorough introduction to HMMs see Zucchini *et al.* 2016).

One of the most common discrete-time animal movement HMMs for telemetry location data is composed of two data streams, step length and turning angle (or bearing), which are calculated for each of the T time steps from the temporally-regular observations of an animal’s position, (x_t, y_t) , for $t = 1, \dots, T + 1$ (e.g. Morales *et al.* 2004; Langrock *et al.* 2012; McClintock *et al.* 2012). Step length (l_t) is typically calculated as the Euclidean distance between the locations (x_t, y_t) and (x_{t+1}, y_{t+1}) , while turning angle (ϕ_t) is calculated as the change in bearing ($b_t = \text{atan2}(y_{t+1} - y_t, x_{t+1} - x_t)$) between the intervals $[t - 1, t]$ and $[t, t + 1]$ (e.g. $\phi_t = 0$ if $b_{t-1} = b_t$). For this HMM composed of 2 data streams, $\mathbf{z}_t = (l_t, \phi_t)$, and, conditional on the latent state S_t , independent prob-

ability distributions are typically assumed for each stream, e.g., $p_s(\mathbf{z}_t) = p_s(l_t)p_s(\phi_t)$. Some common probability distributions for the step length data stream are the gamma or Weibull distributions, while the wrapped Cauchy or von Mises distributions are often employed for turning angle or bearing.

While HMMs for animal movement based solely on location data are somewhat limited in the number and type of biologically-meaningful movement behavior states they are able to accurately identify, advances in biologging technology are now allowing the collection of valuable auxiliary biotelemetry data (e.g., dive activity, accelerometer, heart rate, stomach temperature), which, when combined with location data, allow for multivariate HMMs that can incorporate much more behavioral realism and facilitate inferences about complex ecological relationships that would otherwise be difficult or impossible to infer from location data alone (e.g. McClintock *et al.* 2013; DeRuiter *et al.* 2017; McClintock *et al.* 2017). Multivariate HMMs that utilize both location and auxiliary biotelemetry data can facilitate the identification of additional states that go beyond the $N = 2$ state approaches that are most frequently used by practitioners. For example, the most widely used 2-state HMMs for animal movement include “encamped” (or “foraging”) and “exploratory” (or “transit”) states characterized by area-restricted-search-type movements (shorter step lengths with little to no directional persistence) and migratory-type movements (longer step lengths with high directional persistence), respectively (Morales *et al.* 2004; Jonsen *et al.* 2005). However, very different behaviors can exhibit similar horizontal trajectories. For example, for herbivores such as North American elk (Morales *et al.* 2004) or central-place foragers such as harbour seals (McClintock *et al.* 2013), the horizontal trajectories of “resting” and “foraging” movements can be very difficult to distinguish. Standard 2-state HMMs based solely on horizontal trajectory will tend to lump these behaviors together, and this could have unintended consequences if, for example, one intends to use the estimated state sequences to identify foraging habitat. In order to tweeze apart distinct behaviors with similar horizontal trajectories, additional states can be informed by auxiliary information (such as mandible accelerometer or dive data), incorporated as additional data stream(s) in a multivariate HMM.

When data streams are observed without error and at regular time intervals, a major advantage of HMMs is the relatively fast and efficient maximization of the likelihood using the forward algorithm (Eq. 1). However, location measurement error is rarely

non-existent in animal-borne telemetry studies and depends on both the device and the system under study. For example, GPS errors are typically less than 50m, but Argos errors can exceed 10km (e.g. Costa *et al.* 2010). An extreme case of missing data can arise when location data are obtained with little or no temporal regularity, as in many marine mammal telemetry studies (e.g. Jonsen *et al.* 2005), such that few (if any) observations align with the regular time steps required by discrete-time HMMs. When explicitly accounting for uncertainty attributable to location measurement error, temporally-irregular observations, or other forms of missing data, one must typically fit (multivariate) HMMs using computationally-intensive (and often time-consuming) model fitting techniques such as Markov chain Monte Carlo (Jonsen *et al.* 2005; McClintock *et al.* 2012). However, complex analyses requiring novel statistical methods and custom model-fitting algorithms are not practical for many practitioners.

While statisticians have been applying HMMs to telemetry data for decades, R (R Core Team 2016) packages such as `bsam` (Jonsen *et al.* 2005), `moveHMM` (Michelot *et al.* 2016), and `swim` (Whoriskey *et al.* 2017) have recently helped make these models of animal movement behavior more accessible to the practitioners that are actually conducting telemetry studies. These advances represent important steps toward making HMMs of animal movement more accessible, but the models that can currently be implemented using existing software remain limited in many key respects. For example, existing HMM software for animal movement is limited to two data streams based solely on location data (e.g. step length and turning angle), and while `moveHMM` allows for a user-specified number of latent behavioral states (`bsam` and `swim` are limited to $N = 2$ states), it is typically difficult to identify ≥ 2 biologically-meaningful behavior states from only 2 data streams (e.g. Morales *et al.* 2004; Beyer *et al.* 2013; McClintock *et al.* 2014). Both `moveHMM` and `swim` are designed for temporally-regular (or linearly-interpolated) location data with negligible measurement error, but the realities of animal-borne telemetry often yield temporally-irregular location data subject to error (particularly in aquatic environments). Other notable deficiencies of existing software include limited abilities to incorporate spatio-temporal environmental or individual covariates on parameters, biased (or directed) movements in response to attractive or repulsive forces (e.g. McClintock *et al.* 2012; Langrock *et al.* 2014), cyclical (e.g. daily, seasonal) and other more complicated behavioral patterns, or constraints on parameters.

To address these deficiencies in existing software, we introduce a new user-friendly R package, **momentuHMM** (Maximum likelihood analysis Of animal MovemENT behavior Using multivariate Hidden Markov Models), intended for practitioners wishing to implement more flexible and realistic (multivariate) HMM analyses of animal movement while accounting for common challenges associated with telemetry data. Features for multivariate HMM analyses in **momentuHMM** include: 1) tools for data pre-processing and visualization; 2) user-specified probability distributions for an unlimited number of data streams and latent behavior states; 3) user-specified design matrices and constraints for covariate modelling of state transition probability and probability distribution parameters using linear model formulas familiar to most R users; 4) multiple imputation methods that account for observation error attributable to measurement error and temporally-irregular or missing data (Hooten *et al.* 2017; McClintock 2017); 5) seamless integration of spatio-temporal environmental covariate data (e.g., wind direction, forest cover, sea ice concentration) using the **raster** package (Hijmans 2016); 6) incorporation of “activity center” effects (e.g., areas associated with attractive or repulsive forces); 7) circular-circular regression models for angular probability distributions (Duchesne *et al.* 2015); 8) cosinor (e.g. Cornelissen 2014) and spline regression formulas for cyclical and other complicated behavioral patterns; and 9) data simulation capabilities for power analyses and assessing model performance, including simulation of location data subject to temporal irregularity and/or measurement error.

In the following tutorial, we demonstrate some of the capabilities of **momentuHMM** using real-world examples, including an example of periodic cycles in African elephant movement, a 3-state (“resting”, “foraging”, “transit”) northern fur seal example incorporating auxiliary dive activity data (McClintock *et al.* 2014), a loggerhead turtle example relating “foraging” and “transit” movements to ocean surface currents, and a 5-state grey seal example incorporating biased movements toward haul-out and foraging locations (McClintock *et al.* 2012). This brief tutorial includes workflows for data formatting, model specification, model fitting, and diagnostics. While many of the features of **momentuHMM** were motivated by animal movement data, the package can be used for analyzing any type of data that is amenable to (multivariate) HMMs. Additional information, including help files, data, examples, and package usage is available by downloading the **momentuHMM** package from CRAN (cran.r-project.org) or Github (github.com/bmccclintock/momentuHMM). This article describes **momentuHMM** version

Table 1. Workhorse functions for the R package `momentuHMM`.

Function	Description
<code>crawlMerge</code>	Merge <code>crawlWrap</code> output with additional data streams or covariates
<code>crawlWrap</code>	Fit <code>crawl</code> models and predict temporally-regular locations
<code>fitHMM</code>	Fit a (multivariate) HMM to the data
<code>MIfitHMM</code>	Fit (multivariate) HMMs to multiple imputation data
<code>MIpool</code>	Pool <code>momentuHMM</code> model results across multiple imputations
<code>plot.crwData</code>	Plot <code>crawlWrap</code> output
<code>plot.miSum</code>	Plot summaries of multiple imputation <code>momentuHMM</code> models
<code>plot.momentuHMM</code>	Plot summaries of <code>momentuHMM</code> models
<code>plot.momentuHMMData</code>	Plot summaries of selected data streams and covariates
<code>plotPR</code>	Plot time series, qq-plots and sample ACFs of pseudo-residuals
<code>plotSat</code>	Plot locations on satellite image
<code>plotSpatialCov</code>	Plot locations on raster image
<code>plotStates</code>	Plot the (Viterbi-)decoded states and state probabilities
<code>prepData</code>	Pre-process data streams and covariates
<code>pseudoRes</code>	Calculate pseudo-residuals for <code>momentuHMM</code> models
<code>simData</code>	Simulate movement data using a (multivariate) HMM
<code>stateProbs</code>	State probabilities for each time step
<code>viterbi</code>	Most likely state sequence (using the Viterbi algorithm)

1.0.0.

2 `momentuHMM` overview

Before delving into some of the finer details, we will first provide an overview of the main features and functions of the package. While space is limited in this tutorial, further details on implementation can be found in the package’s documentation and vignette. The workhorse functions of `momentuHMM` are listed in Table 1. Usage of several of these functions (e.g. `fitHMM`, `prepData`, `simData`) is deliberately very similar to equivalent functions in `moveHMM` (Michelot *et al.* 2016), but the `momentuHMM` arguments for these functions have been generalized and expanded to accommodate a more flexible framework for data pre-processing, model specification, parameterization, and simulation. R users already familiar with `moveHMM` will therefore likely find it easy to immediately begin using `momentuHMM`.

One of the key features of `momentuHMM` is the ability to specify an unlimited number of HMM data streams from a broad range of commonly used probability distribu-

tions. Any of the parameters of the probability distributions used for the observed data can be modelled as a function of environmental and individual covariates using link functions (Table 2). For any given “natural scale” (or “real scale”) probability distribution parameter θ , all of the link functions (g) in `momentuHMM` are of the general form $g(\boldsymbol{\theta}) = \mathbf{X}_\theta \boldsymbol{\beta}_\theta$, where \mathbf{X}_θ is the $T \times k$ design matrix (composed of k covariates) and $\boldsymbol{\beta}_\theta$ is the corresponding k -vector of “working scale” (or “beta scale”) parameters for θ . For example, suppose step length is assumed to have a gamma distribution, $l_t \mid S_t = s \sim \text{gamma}(\mu_s, \sigma_s)$. In `momentuHMM`, the natural scale parameters for the gamma distribution are the (state-dependent) step length mean ($\mu_s > 0$) and standard deviation ($\sigma_s > 0$). Because both of these parameters must be positive, the log link function is a natural choice for modelling these parameters as a function of covariates, e.g., $\log(\boldsymbol{\mu}) = \mathbf{X}_\mu \boldsymbol{\beta}_\mu$.

The state transition probabilities can also be modelled as functions of covariates, using a multinomial logit link, as described e.g. by Michelot *et al.* (2016).

2.1 Data preparation and visualization

For temporally-regular location data with negligible measurement error, the `prepData` function is used to create a `momentuHMMData` object that can be used for data visualization and further analysis. The arguments for `prepData` include:

- **data** A data frame with $T + 1$ rows including optionally a field ‘ID’ (identifiers for different individuals), coordinates from which step length (‘step’) and turning angle (‘angle’) data streams are to be calculated, any additional data streams, and any covariates identified in the `covNames` and `angleCovs` arguments. Alternatively, **data** can be a `crwData` object returned by `crawlWrap`.
- **type** Coordinate type; ‘UTM’ if easting-northing or ‘LL’ if longitude-latitude.
- **coordNames** Names of the two coordinate columns in **data**. If `coordNames=NULL` then step lengths, turning angles, and any location-based covariates (i.e., those specified by `spatialCovs`, `centers`, and `angleCovs`) are not calculated.
- **covNames** Character vector indicating the names of any covariates in **data**. Any variables in **data** (other than “ID”) that are not identified in `covNames` or `angleCovs` are assumed to be data streams.

Table 2. Data stream (z) probability distributions, natural parameters, and default link functions for covariate modelling. Probability distributions with positive support can be zero-inflated (with additional zero-mass parameters), while the beta distribution can be zero- and one-inflated (with additional one-mass parameters). If user-specified bounds are provided, then custom link functions are used instead of the defaults (see package documentation for further details). If circular-circular regression is specified for the mean of angular distributions (“vm” and “wrpcauchy”), then the link function described in Duchesne et al. (2015) is used. Users seeking additional probability distributions are encouraged to contact the authors.

Distribution	Support	Parameters	Link function
Beta (“beta”)	$z_t \in (0, 1)$	shape1 > 0 shape2 > 0 zero-mass $\in (0, 1)$ one-mass $\in (0, 1)$	log log logit logit
Exponential (“exp”)	$z_t > 0$	rate > 0 zero-mass $\in (0, 1)$	log logit
Gamma (“gamma”)	$z_t > 0$	mean > 0 sd > 0 zero-mass $\in (0, 1)$	log log logit
Log normal (“lnorm”)	$z_t > 0$	location $\in \mathbb{R}$ scale > 0 zero-mass $\in (0, 1)$	identity log logit
Poisson (“pois”)	$z_t \in \{0, 1, \dots\}$	lambda > 0	log
Von Mises (“vm”)	$z_t \in (-\pi, \pi]$	mean $\in (-\pi, \pi]$ concentration > 0	tan(mean/2) log
Weibull (“weibull”)	$z_t > 0$	shape > 0 scale > 0 zero-mass $\in (0, 1)$	log log logit
Wrapped Cauchy (“wrpcauchy”)	$z_t \in (-\pi, \pi]$	mean $\in (-\pi, \pi]$ concentration $\in (0, 1)$	tan(mean/2) logit

- **spatialCovs** List of Raster-class objects (Hijmans 2016) containing spatio-temporally referenced covariates. Covariates specified by **spatialCovs** are extracted from the raster layer(s) based on the location data. Raster stacks may also be included, in which case the appropriate z values (e.g. time, date) must also be included in **data**.
- **centers** 2-column matrix providing the coordinates for any activity centers (e.g., potential centers of attraction or repulsion) from which distance and angle covariates will be calculated based on the location data and returned in the **momentuHMMData** object.
- **angleCovs** Character vector indicating the names of any circular-circular regression angular covariates in **data** or **spatialCovs** that need conversion from standard direction (in radians relative to the x-axis) to turning angle (relative to previous movement direction).

Summary plots of the **momentuHMMData** object returned by **prepData** can be created for any data stream or covariate using the generic **plot** function.

If location data are temporally-irregular or subject to measurement error, then they are not suitable for **prepData**. In this case, **momentuHMM** can be used to perform a 2-stage multiple imputation approach (McClintock 2017). We discuss this pragmatic approach to incorporating uncertainty attributable to observation error and temporal irregularity into multivariate HMM analyses in section 2.3.

2.2 HMM specification and fitting

Once a **momentuHMMData** object has been created using **prepData**, then the data are ready to be passed to the generalized multivariate HMM-fitting function **fitHMM**. There are many different options for specifying HMMs using **fitHMM**, so here we will only focus on several of the most important and useful features (further details of all **fitHMM** arguments are in the package documentation). The bare essentials of **fitHMM** include the arguments:

- **data** A **momentuHMMData** object
- **nbStates** Number of latent states (N)

- `dist` A named list indicating the probability distributions of the data streams.
- `estAngleMean` An optional named list indicating whether or not to estimate the angle mean for data streams with angular distributions (e.g. turning angle). If not estimated (the default), the angle mean is fixed to 0.
- `formula` Regression formula for the transition probability covariates
- `stationary` Logical indicating whether or not the initial distribution is considered equal to the stationary distribution (must be `FALSE` if `formula` includes covariates)
- `Par0` A named list containing vectors of starting values for the state-dependent probability distribution parameters of each data stream

These seven arguments are all that are needed in order to fit the HMMs currently supported in `moveHMM` (Michelot *et al.* 2016). For example, here is how the analysis of 15 “wild haggis” tracks described in Michelot *et al.* (2016) would be implemented using `momentuHMM`:

```
library(momentuHMM)
### Load raw data
rawHaggis<-read.csv("rawHaggises.csv")
### Process data
processedHaggis<-prepData(data=rawHaggis,covNames=c("slope","temp"))

### Fit HMM
# initial step distribution natural scale parameters
stepPar0 <- c(1,5,0.5,3) # (mu_1,mu_2,sd_1,sd_2)
# initial angle distribution natural scale parameters
anglePar0 <- c(0,0,1,8) # (mean_1,mean_2,concentration_1,concentration_2)
fitHaggis <- fitHMM(data = processedHaggis, nbStates = 2,
                    dist = list(step = "gamma", angle = "vm"),
                    Par0 = list(step = stepPar0, angle = anglePar0),
                    formula = ~ slope + I(slope^2),
                    estAngleMean = list(angle=TRUE))
```

Note that many of the arguments in `fitHMM` are lists, with each element of the list corresponding to a data stream. The list names provided in `dist`, `Par0`, and `estAngleMean` (e.g. ‘step’ and ‘angle’) must therefore have a corresponding column in `data` with the same name. Additional data streams can be added to the model by

simply adding the additional elements to these list arguments (see example in section 3.2).

As seen above, the `formula` argument can include many of the functions and operators commonly used to construct terms in R linear model formulas (e.g. `a*b`, `a:b`, `cos(a)`). Unique to `momentuHMM`, the `formula` argument can also be used to specify transition probability matrix models that incorporate cyclical patterns (using the `cosinor` special function; see example in section 3.1), splines for explaining other more complicated patterns (e.g., `bs` and `ns` functions in the R base package `splines`), and factor variables (e.g., `formula=~ID` for individual-level effects). By default the `formula` argument applies to all state transition probabilities, but the special functions `state` and `betaCol` allow for state- and parameter-specific formulas to be specified (see example in section 3.4). Specific state transition probabilities can also be fixed to zero (or any other value) using the `fixPar` argument, which can be useful for incorporating more behavioral realism. For example, `fixPar` can be used to prohibit or enforce switching from one particular state to another (possibly as a function of spatio-temporal covariates).

Similar to the `formula` argument for state transition probability modelling, it is through the `DM` argument of `fitHMM` that models are specified for the state-dependent probability distribution parameters for each data stream. `DM` is a list argument containing an element for each data stream, but each element itself is also a list specifying the design matrix formulas for each parameter. For example, the following fits the exact same wild haggis model as above, but employs a user-specified (intercept-only) design matrix for the step length data stream:

```
stepDM <- list(mean = ~1, sd = ~1)

### Fit HMM using user-specified DM
fitHaggisDM <- fitHMM(data = processedHaggis, nbStates = 2,
  dist = list(step = "gamma", angle = "vm"),
  DM = list(step = stepDM),
  Par0 = list(step = log(stepPar0), angle = anglePar0),
  formula = ~ slope + I(slope^2),
  estAngleMean = list(angle=TRUE))
```

Note that when `DM` is specified for a data stream, the initial parameter values (`Par0`) for that data stream now correspond to columns of the resulting design matrix and must

be on the working scale instead of the natural scale. In this case, because the log link is used for the natural parameters of the gamma distribution, `Par0$step` was specified on the log scale. The functions `getPar0` and `getParDM` are designed to assist users in the specification of design matrices and corresponding initial values on the working scale for any given model (see package documentation for further details). DM formulas are just as flexible as the `formula` argument and, in addition to common linear model formula functions and operators, can also include cyclical cosinor models, splines, factor variables, and state-specific probability distribution parameter formulas (see examples in sections 3.3 and 3.4). As with the state transition probabilities, working parameters for probability distributions can also be fixed to user-specified values using the `fixPar` argument.

Specification of design matrices using DM is not limited to formulas. Alternatively, “pseudo-design” matrices can be specified, using an R matrix with rows corresponding to the natural parameters and columns corresponding to the working parameters. The elements in the matrix may be numeric or character strings containing model formula terms (see examples in sections 3.2 and 3.4). Using a pseudo-design matrix for step length, the following is yet another way to implement the exact same wild haggis model:

```
stepDMp <- matrix(c(1,0,0,0,
                  0,1,0,0,
                  0,0,1,0,
                  0,0,0,1),4,4,byrow=TRUE)
rownames(stepDMp) <- c("mean_1","mean_2","sd_1","sd_2")
colnames(stepDMp) <- c("mean_1:(Intercept)","mean_2:(Intercept)",
                    "sd_1:(Intercept)","sd_2:(Intercept)")

### Fit HMM using user-specified DM
fitHaggisDMp <- fitHMM(data = processedHaggis, nbStates = 2,
                    dist = list(step = "gamma", angle = "vm"),
                    DM = list(step = stepDMp),
                    Par0 = list(step = log(stepPar0), angle = anglePar0),
                    formula = ~ slope + I(slope^2),
                    estAngleMean = list(angle=TRUE))
```

(note that column and rownames for pseudo-design matrices are not required but can be useful). Pseudo-design matrices allow for the sharing of common working parameters

(such as intercept terms) among natural scale parameters, and this can be used to constrain natural scale parameters (e.g., $\mu_1 \leq \mu_2$) when used in tandem with the `cons` argument. This is particularly useful for preventing state label switching when repeatedly fitting the same HMM using multiple imputation methods (see section 2.3).

Another noteworthy `fitHMM` argument, `circularAngleMean`, is a list argument that enables users to specify circular-circular regression models for the mean (μ) parameter of angular distributions, such as the wrapped Cauchy and von Mises, instead of circular-linear models based on the tangent link function (Table 2). When `circularAngleMean` is specified as `TRUE` for any given angular data stream (e.g. turning angle), then a special link function based on Duchesne *et al.* (2015) is used:

$$\boldsymbol{\mu} = \text{atan2}(\sin(\mathbf{X}_\mu)\boldsymbol{\beta}_\mu, 1 + \cos(\mathbf{X}_\mu)\boldsymbol{\beta}_\mu), \quad (2)$$

where \mathbf{X}_μ is a $T \times k$ matrix composed of the turning angles between k angular covariates (e.g., wind direction, sea surface currents) and the bearing of movement during the previous time step, i.e., each element $x_{t,i} = \text{atan2}(\sin(r_{t,i} - b_{t-1}), \cos(r_{t,i} - b_{t-1}))$ for angular covariate $r_{t,i}$ and $i = 1, \dots, k$ (note that `prepData` calculates \mathbf{X}_μ based on the `angleCovs` argument so users need not bother). Because this link function is designed for turning angles, a turning angle of 0 is provided as the reference angle (hence the “1+” preceding the cosine term in Eq. 2). Thus as a trade-off between biased and correlated movements, the working parameters ($\boldsymbol{\beta}_\mu$) for the expected turning angle at time t weight the attractive (or repulsive) strengths of the angular covariates relative to directional persistence. When all $\boldsymbol{\beta}_\mu = 0$, the model reduces to a correlated random walk, but an increasingly biased random walk results as $\boldsymbol{\beta}_\mu$ gets larger (or smaller). Many interesting hypotheses about animal movement can be addressed using circular-circular regression on movement direction, including the effects of wind, sea surface currents (see example in section 3.3), and centers of attraction or repulsion (see example in section 3.4)

2.3 Multiple imputation

When location data are temporally-irregular or subject to measurement error, then they are not suitable for standard maximum-likelihood HMM analyses based on the forward algorithm (Eq. 1). In this case, `momentuHMM` can be used to perform the 2-stage multiple imputation approach of McClintock (2017). The basic concept is to

first employ a single-state (i.e., $N = 1$) movement model that is relatively easy to fit but can accommodate location measurement error and temporally-irregular or missing observations (e.g. Johnson *et al.* 2008). The second stage involves repeatedly fitting the desired HMM to m temporally-regular realizations of the position process drawn from the model output of the first stage. Data streams or covariates that are dependent on location (e.g., step length, turning angle, habitat type, snow depth, sea surface temperature) will of course vary among the m realizations of the position process, and the pooled inferences across the HMM analyses therefore reflect location uncertainty.

There are three primary functions (`MIfitHMM`, `MIpool`, and `crawlWrap`) for performing multiple imputation HMM analyses in `momentuHMM`, and all rely on parallel processing to speed up computations. `crawlWrap` is a wrapper function for fitting the continuous-time correlated random walk (CTCRW) model of Johnson *et al.* (2008) to one or more tracks (subject to location measurement error and/or temporal irregularity) and then predicting temporally-regular tracks of the user's choosing (e.g. 15 min, hourly, daily) based on the CTCRW model output. `crawlWrap` returns a `crwData` object that can be used to draw m realization of the position process within the `MIfitHMM` function. `MIfitHMM` is essentially a wrapper function for `fitHMM` that repeatedly fits the same user-specified HMM to m imputed data sets and stores the output from each of the m model fits. If a `crwData` object is provided, then `MIfitHMM` will first draw m imputations based on the `crwData` output and then fit the specified HMM to each imputed data set. If users wish to use a movement model other than the CTCRW to account for measurement error and temporal irregularity (e.g. Calabrese *et al.* 2016; Gurarie *et al.* 2017), or if other observation error processes (e.g. missing data) are to be accounted for in the imputation step, `MIfitHMM` can also be used for analysis of a list of m `momentuHMMData` objects that were imputed by the user. Based on the m model fits, the `MIpool` function calculates pooled estimates, standard errors, and confidence intervals for the working scale parameters, natural scale parameters (based on transformations of the pooled working parameters and mean or user-specified values for any covariates), state sequences, state probabilities, and activity budgets (i.e. the proportion of the T times step assigned to each state) using standard multiple imputation formulae (Rubin & Schenker 1986; McClintock 2017). `MIpool` can be called separately or within `MIfitHMM` (using the `poolEstimates` argument), and the function returns a `miSum` object containing the pooled output across all imputations. See sections 3.2, 3.3,

and 3.4 for example HMM analyses that use multiple imputation to account for location measurement error and temporally irregularity.

2.4 Model visualization and diagnostics

The generic `plot` functions for `momentuHMM` models (`plot.momentuHMM` and `plot.miSum`) plot the data stream histograms along with their corresponding estimated probability distributions, the estimated natural parameters and state transition probabilities as a function of any covariates included in the model, and the tracks of all individuals (color-coded by the most likely state sequence). By default, the probability distributions are plotted based on the means of any covariate values, but user-specified covariate values for the plots can be provided using the `covs` argument. When the argument `plotCI=TRUE`, then confidence intervals for the natural parameters and state transition probabilities are also plotted. Confidence intervals are calculated from the working parameter estimates based on the delta method and finite-difference approximations of the first derivative for the transformation using the `numDeriv::grad` function (Gilbert & Varadhan 2016). For multiple imputation analyses (`plot.miSum`), all plots are based on the pooled parameter estimates and the means of any covariates (if not provided by the `covs` argument) across each imputation. Using the argument `errorEllipse`, `plot.miSum` will include estimated location error ellipses in the plots of individual tracks. The functions `plotSat`, `plotSpatialCov`, and `plotStates` (Table 1) provide further methods for visualizing model results.

Diagnostic tools include the calculation and plotting of pseudo-residuals (Zucchini *et al.* 2016) using the `pseudoRes` and `plotPR` functions, respectively. Akaike’s Information Criterion can be calculated for one or more models using the `AIC.momentuHMM` function.

2.5 Simulation

The function `simData` can be used to simulate multivariate HMM data from scratch or based on the estimated parameters of existing `momentuHMM` or `miSum` models. The `simData` arguments are very similar to those used for model specification in `fitHMM` (e.g., `dist`, `DM`) and data preparation in `prepData` (e.g., `spatialCovs`, `centers`), but they include additional arguments, `lambda` and `errorEllipse`, for simulating location data subject to temporal irregularity and measurement error, respectively. The `spatialCovs`

argument allows for rasters of spatio-temporal covariate values to be utilized in simulation models, while the `centers` argument allows activity centers to be incorporated. Thus `simData` can be used to simulate more ecologically-realistic tracks (potentially subject to observation error) that can be useful for study design, power analyses, and assessing model performance.

3 Examples

We will now demonstrate some of the capabilities of `momentuHMM` using real telemetry data. These examples are intended for demonstration purposes only, and we do not claim these example analyses represent improvements relative to previous or alternative analyses for these data sets. While only some of the key workflow elements are included here, complete R code and further details for these analyses are available on request.

3.1 African elephant

As our first example, we use an African elephant (*Loxodonta africana*) bull track described in Wall *et al.* (2014) and publicly available from the `movebank.org` data repository.

We can load the data from the URL (this requires an Internet connection):

```
URL <- paste0("https://www.datarepository.movebank.org/bitstream/handle/",
              "10255/move.373/Elliptical%20Time-Density%20Model%20%28Wall%",
              "20et%20al.%202014%29%20African%20Elephant%20Dataset%20%",
              "28Source-Save%20the%20Elephants%29.csv")
rawData <- read.csv(url(URL))
```

The data set contains two tracks; for this analysis, we only consider the first one. In addition to hourly locations, the tag also collected external temperature data. We subset the data frame to keep only the relevant rows and columns:

```
# select and rename relevant columns
rawData <- rawData[,c(11,3,4,5,6)]
colnames(rawData) <- c("ID", "time", "lon", "lat", "temp")

# only keep first track
rawData <- subset(rawData, ID==unique(ID)[1])
```

The data now has the following columns:

```
head(rawData)

##           ID           time           lon           lat temp
## 1 Salif Keita 2008-03-22 17:00:00.000 -2.160167 15.65350   38
## 2 Salif Keita 2008-03-22 18:00:00.000 -2.160075 15.65452   35
## 3 Salif Keita 2008-03-22 19:00:00.000 -2.159902 15.65451   32
## 4 Salif Keita 2008-03-22 20:00:00.000 -2.159435 15.65489   30
## 5 Salif Keita 2008-03-22 21:00:00.000 -2.158113 15.65512   29
## 6 Salif Keita 2008-03-22 22:00:00.000 -2.157848 15.65461   28
```

Location measurement error is negligible for these terrestrial GPS data, although about 1% of the hourly observations collected between 22 March 2008 and 30 September 2010 are missing. Instead of simply ignoring these missing data, we can employ `crawlWrap` to predict the missing locations based on the CTCRW model of Johnson *et al.* (2008) prior to conducting our HMM analysis.

To use `crawlWrap`, we convert times from factors to POSIX, and project the observed locations to UTM coordinates:

```
# convert times from factors to POSIX
rawData$time <- as.POSIXct(rawData$time,tz="GMT")

# project to UTM coordinates using package rgdal
library(rgdal)
llcoord <- SpatialPoints(rawData[,3:4],
                          proj4string=CRS("+proj=longlat +datum=WGS84"))
utmcoord <- spTransform(llcoord,CRS("+proj=utm +zone=30 ellps=WGS84"))

# add UTM locations to data frame
rawData$x <- attr(utmcoord,"coords")[,1]
rawData$y <- attr(utmcoord,"coords")[,2]
```

Then, we call `crawlWrap` to fit a CTCRW model and predict hourly locations:

```
# initial parameters for crawl fit
inits <- list(a = c(rawData$x[1],0,rawData$y[1],0),
              P = diag(c(5000^2, 10*3600^2, 5000^2, 10*3600^2)))

# fit crawl model
```

```
crwOut <- crawlWrap(obsData=rawData, timeStep="hour", initial.state=inits,
                    theta=c(4,-10), fixPar=c(NA,NA))
```

Here the desired time step is specified by the `timeStep` argument, and the `initial.state`, `theta`, and `fixPar` arguments are the same as for `crawl::crwMLE` (Johnson 2017). For the choice of initial parameters in `crawlWrap`, we refer the reader to the documentation of the package `crawl`, in particular `crawl::crwMLE` and `crawl::crwPredict`. We now have a complete set of temporally-regular location data.

Autocorrelation function estimates suggest there are 24-hour cycles in the step length data, and this presents an opportunity to demonstrate the use of the `cosinor` function for incorporating cyclical behavior in model parameters using `momentuHMM`. The ACF plot, shown in Figure 1, was obtained with:

```
acf(data$step[!is.na(data$step)],lag.max=300)
```

We create a `momentuHMMData` object, and the 24-hour cosinor model covariate:

```
# create momentuHMMData object from crwData object
dataHMM <- prepData(data=crwOut, covNames="temp")

# add cosinor covariate based on hour of day
dataHMM$hour <- as.integer(strftime(dataHMM$time, format = "%H", tz="GMT"))
```

As seen here, the function `prepData` can also be used for pre-processing the best predicted track data from `crawlWrap` output. The 24-hour cosinor covariate (“hour”) is simply a set of integers (0, 1, . . . , 23) indicating the hour of day for each observation.

Our aim is to fit a 2-state HMM to the elephant track that includes temperature effects on the turning angle concentration parameters and cycling temperature effects (with a 24-hour periodicity) on the step length and state transition probability parameters. A complex model, like this one, requires many parameters, such that it is challenging to choose good starting parameter values for the optimization. Here, we take an incremental approach, starting from a simpler model with no covariates. In `momentuHMM`, the function `getPar0` extracts initial parameters from a fitted (nested) HMM, given arguments for the more complex model.

For the covariate-free 2-state model, six initial parameters need to be chosen: for each state, the mean and standard deviation of the gamma distribution of step lengths,

and the concentration of the wrapped Cauchy distribution of turning angles. Looking at the histograms of the step lengths and the turning angles (e.g. output by `plot(data)`) is often useful to choose good starting parameter values.

```
# label states
stateNames <- c("encamped","exploratory")
# distributions for observation processes
dist = list(step = "gamma", angle = "wrpcauchy")

# initial parameters
Par0_m1 <- list(step=c(100,500,100,200),angle=c(0.3,0.7))

# fit model
m1 <- fitHMM(data = data, nbStates = 2, dist = dist, Par0 = Par0_m1,
             estAngleMean = list(angle=FALSE), stateNames = stateNames)
```

To ensure convergence, we could also use the argument `retryFits` to specify the number of attempts to minimize the negative log-likelihood based on random perturbations of the parameter estimates at the current minimum.

We can build on complexity, by including the temperature and time of day as covariates in the state transition probabilities. We use the function `getPar0` to extract the new starting parameter values.

```
# formula for transition probabilities
formula <- ~ temp * cosinor(hour, period = 24)

# initial parameters (obtained from nested model m1)
Par0_m2 <- getPar0(model=m1, formula=formula)

# fit model
m2 <- fitHMM(data = data, nbStates = 2, dist = dist, Par0 = Par0_m2$Par,
             beta0=Par0_m2$beta, stateNames = stateNames, formula=formula)
```

The special function `cosinor(hour, period = 24)` internally creates the `cosinor` model covariates, $\cos(2\pi \times \text{hour}/\text{period})$ and $\sin(2\pi \times \text{hour}/\text{period})$, and includes both terms (plus interactions with “temp”) in the fitted model.

Finally, we can fit the more complex model, including the effect of temperature and time of day on the parameters of the state-dependent distributions of steps and angles.

```
# formulas for parameters of state-dependent observation distributions
DM <- list(step = list(mean = ~ temp * cosinor(hour, period = 24),
                      sd = ~ temp * cosinor(hour, period = 24)),
          angle = list(concentration = ~ temp))

# initial parameters (obtained from nested model m2)
Par0_m3 <- getPar0(model=m2, formula=formula, DM=DM)

# fit model
m3 <- fitHMM(data = data, nbStates = 2, dist = dist, Par0 = Par0_m3$Par,
            beta0 = Par0_m3$beta, DM = DM, stateNames = stateNames,
            formula = formula)
```

The above model `m3` identified a state of slow undirected movement (“encamped”), and a state of faster and more directed movement (“exploratory”) (Figure 1). For a fitted model, the function `viterbi` computes the most likely state sequence:

```
# decode most likely state sequence
states <- viterbi(m3)
# derive percentage of time spent in each state
table(states)/nrow(data)
```

Here, about 74% of the steps were attributed to the “encamped” state, and 26% were attributed to the “exploratory” state.

We can use `AIC(m1,m2,m3)` to compare the three fitted models in terms of AIC; here, `m3` is overwhelmingly supported by the AIC when compared to alternative models with fewer covariates.

The model can be visualized with the generic function `plot`, which was used for the plots shown in Figure 2, and the decoded track in Figure 1.

```
plot(m3, plotCI = TRUE, covs = data.frame(hour=12))
```

Interestingly, this model suggests step lengths and directional persistence for the “encamped” state decreased as temperature increased, step lengths for both states tended to decrease in the late evening and early morning, and transition probabilities

from the “encamped” to “exploratory” state decreased as temperature increased (Figure 2).

Model fit can be assessed using the pseudo-residuals, with the functions `pseudoRes` and `plotPR`. The residual ACF plot shown in Figure 1 was produced by:

```
# compute pseudo-residuals for the steps and the angles
pr <- pseudoRes(m3)

# plot the ACF of step pseudo-residuals
acf(pr$stepRes[!is.na(pr$stepRes)], lag.max = 300)
```

Autocorrelation function plots of the pseudo-residuals indicate this model explained much of the periodicity in step length, although there does still appear to be some room for improvement.

3.2 Northern fur seal

In our second example, we use the northern fur seal (*Callorhinus ursinus*) example from McClintock *et al.* (2014) to demonstrate the use of additional data streams for distinguishing behaviors with similar horizontal trajectories in a multivariate HMM. The data consist of 241 temporally-irregular Fastloc GPS locations obtained during a foraging trip of a nursing female near the Pribilof Islands of Alaska, USA, from 10-17 October 2007. The tag included time-depth recording capabilities, and the dive activity data were summarized as the number of foraging dives over $T = 228$ temporally-regular 1 hr time steps. To fit the $Z = 3$ state (1=“resting”, 2=“foraging”, 3=“transit”) of McClintock *et al.* (2014) using `momentuHMM`, we first used `crawlWrap` to predict temporally-regular locations at 1 hr time steps assuming a bivariate normal measurement error model and merged the results with the foraging dive data using the `crawlMerge` function. Then multiple imputation was used to account for location measurement error by repeatedly fitting the HMM to `nSims` realizations of the position process using `MifitHMM`:

```
nbStates <- 3
stateNames <- c("resting", "foraging", "transit")
dist <- list(step = "gamma", angle = "wrpcauchy", dive = "pois")
Par0 <- getParDM(nbStates = nbStates, dist = dist,
                 Par = Par, DM = DM, cons = cons,
                 estAngleMean = list(angle = FALSE))
```

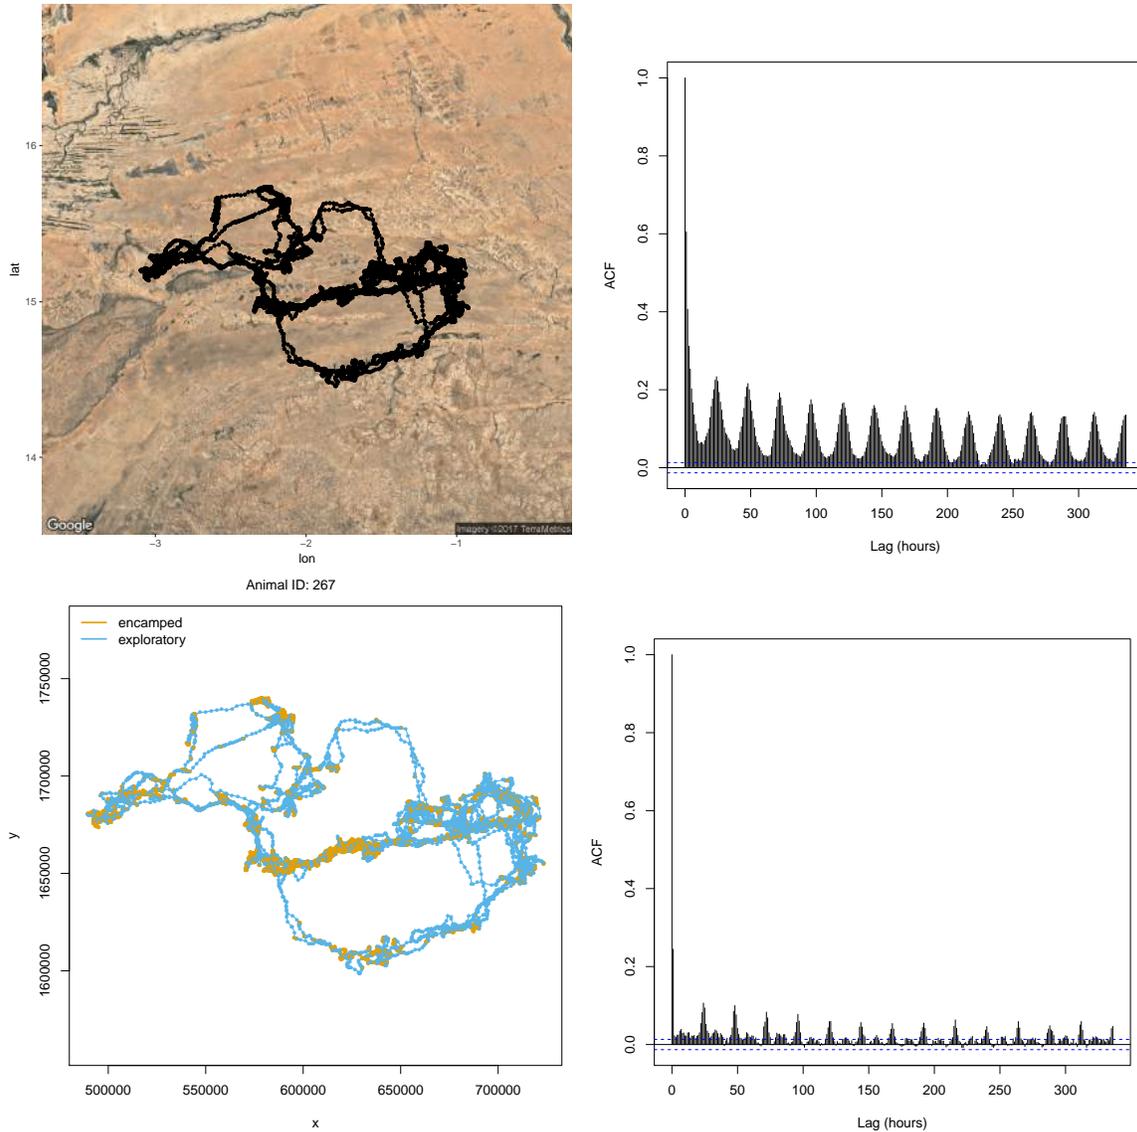


Figure 1. Plot of the elephant track produced using the ‘plotSat’ function (top-left panel), autocorrelation function (ACF) plot of the corresponding step length data (top-right panel), plot of the Viterbi-decoded state sequence for the 2-state (“encamped” and “exploratory”) model generated using the generic ‘plot’ function (bottom-left panel), and the step length pseudo-residual ACF plot for this model using the ‘plotPR’ function (bottom-right panel).

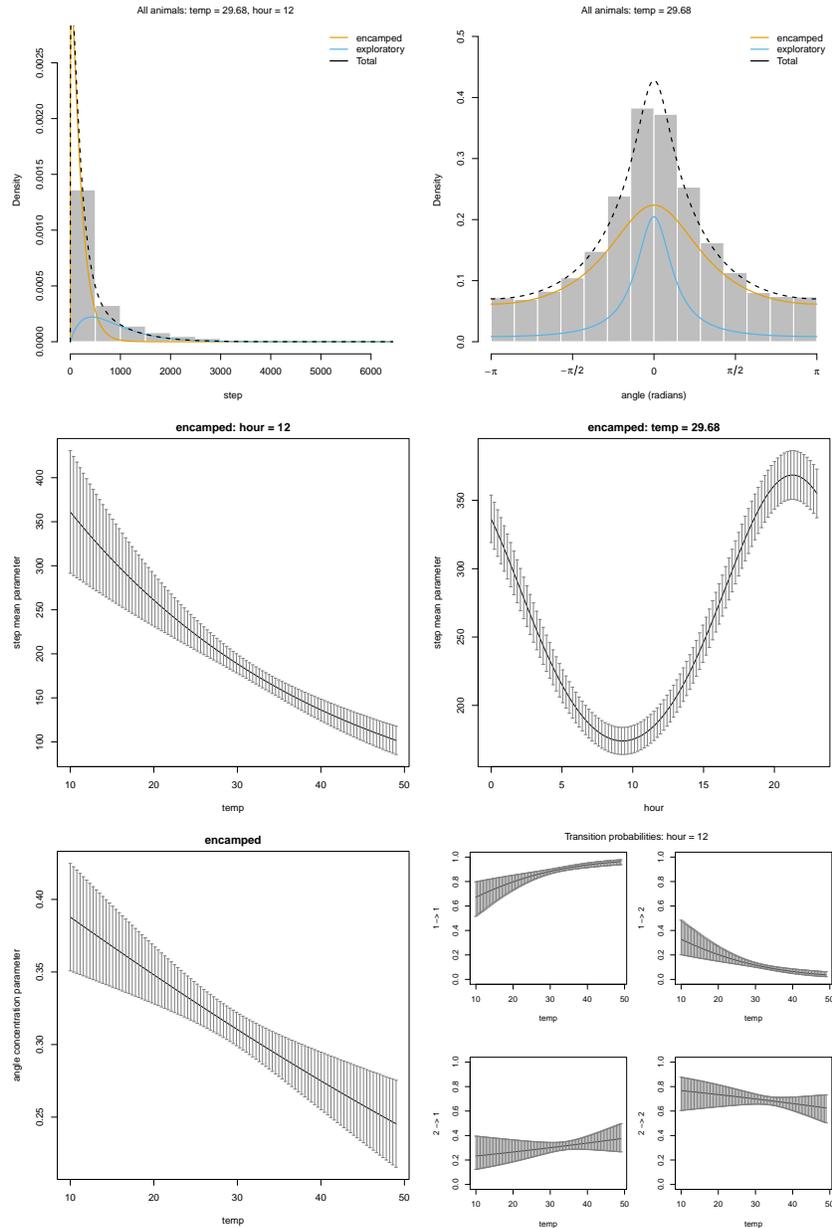


Figure 2. Selected plots for the 2-state (“encamped” and “exploratory”) African elephant example generated using the generic ‘plot’ function. Top panels present histograms of the step length (top-left) and turning angle (top-right) data along with the estimated state-dependent probability distributions based on the mean temperature (temp = 29.7 degrees celsius) at 12:00 GMT (hour = 12). Middle panels present estimates (and 95% confidence intervals) for the step length mean parameter of the “encamped” state as a function of temperature and hour of day. Bottom-left panel presents estimates for the turning angle concentration parameter of the “encamped” state as a function of temperature. Bottom-right panel presents estimated state transition probabilities (1 = “encamped”, 2 = “exploratory”) as a function of temperature at 12:00 GMT.

```

Fixpar <- list(dive = c(-100, NA, NA))
nfsFits <- MIfitHMM(crwOut, nSims = 100, nbStates = nbStates, dist = dist,
                  Par0 = Par0, DM = DM, cons = cons,
                  estAngleMean = list(angle = FALSE),
                  fixPar = fixPar, retryFits = 30,
                  stateNames=stateNames)
plot(nfsFits)

```

Here we specified a gamma distribution for step length (`'step'`), wrapped Cauchy distribution for turning angle (`'angle'`), and Poisson distribution for the number of foraging dives (`'dive'`). The function `getParDM` was used to organize the starting values for the data stream working parameters (`Par0`) in the correct format based on `DM`, `cons`, and estimates of the natural parameters (`Par`) from McClintock *et al.* (2014). The `DM` and `cons` arguments were specified to avoid label switching among the `nSims` imputed data model fits and enforce similar state-dependent probability distribution constraints as McClintock *et al.* (2014), e.g., constraining the Poisson rate parameters such that the “foraging” state tends to have higher numbers of foraging dives than the “transit” state ($\lambda_2 > \lambda_3$). To prohibit foraging dives for the “resting” state, we used the `fixPar` argument to effectively fix the Poisson rate parameter to zero on the natural scale (i.e. $\lambda_1 \approx 0$). To help deal with the problem of convergence to local maxima, the `retryFits` argument allows users to specify the number of times to attempt to re-fit each model using random perturbations of the parameter estimates as the starting values for optimization.

The results are very similar to those of the discrete-time model of McClintock *et al.* (2014), with periods of foraging often followed by resting (Figure 3). The “activity budgets” (i.e. the proportion of time steps allocated to each state) calculated by `MIpool` based on the estimated state sequences for each imputation were 0.3 (95% CI: 0.22–0.39) for “resting”, 0.29 (95% CI: 0.22–0.36) for “foraging”, and 0.41 (95% CI: 0.32–0.52) for “transit”.

3.3 Loggerhead turtle

For our third example, we demonstrate how to model movement direction and step length as a function of angular covariates using hitherto unpublished loggerhead turtle (*Caretta caretta*) data for a captive-raised juvenile released in 2012 on the coast of North

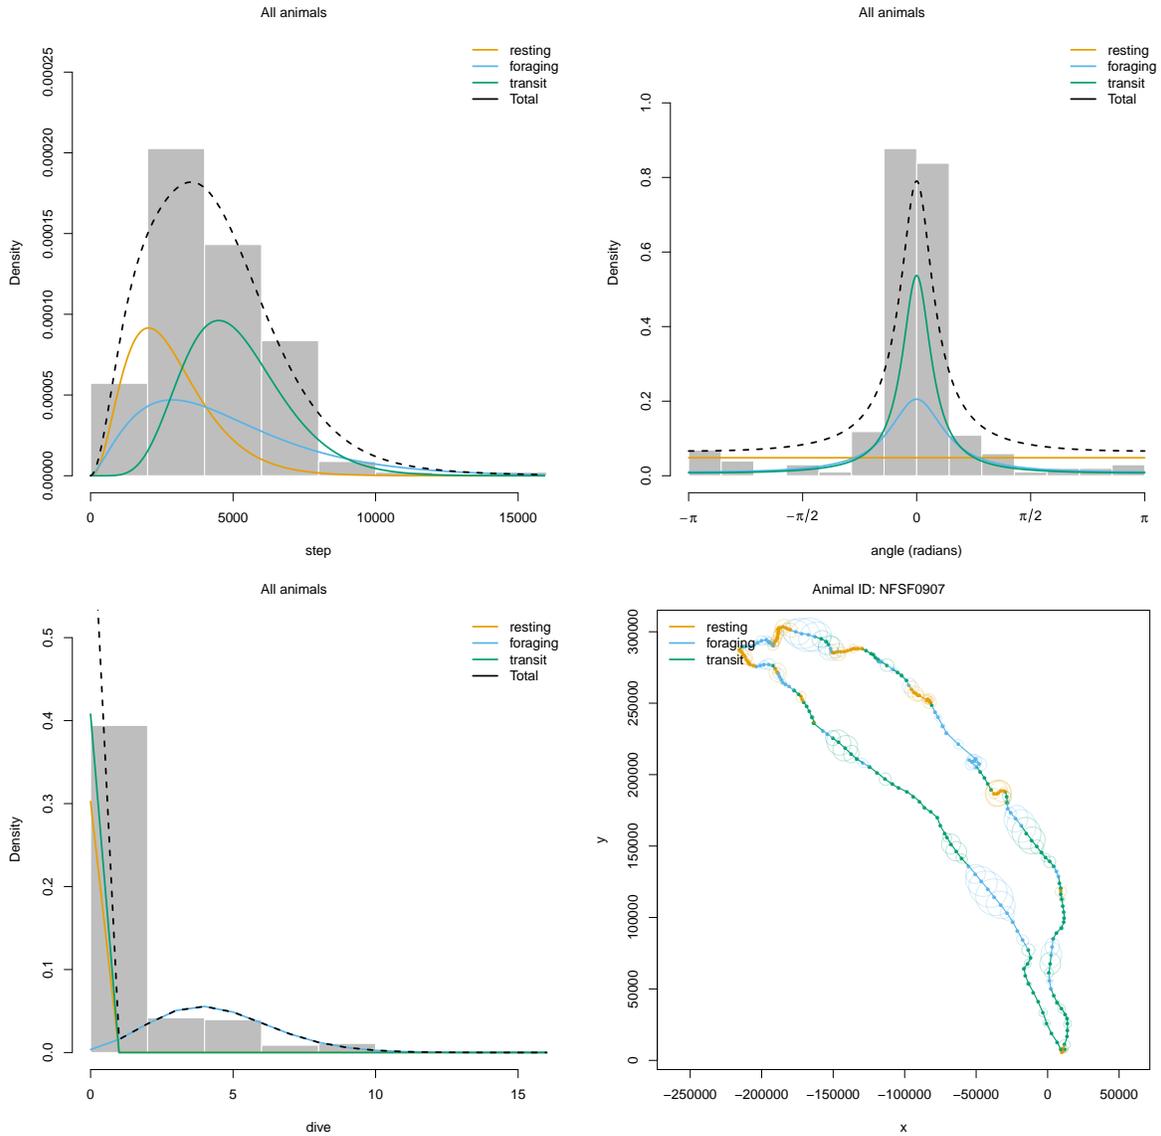


Figure 3. Plots of the northern fur seal example results generated using the generic ‘plot’ function. The estimated probability distributions for step length (top-left panel), turning angle (top-right panel), and number of foraging dives (bottom-left panel) for the 3-state (“resting”, “foraging”, and “transit”) model are plotted along with histograms of these data streams. The temporally-regular predicted locations (and 95% ellipsoidal confidence bands) and estimated states are plotted in the bottom-right panel. All estimates are pooled across multiple imputations of the position process and thus reflect uncertainty attributable to location measurement error and temporally-irregular observations.

Carolina, USA. The data consist of 165 temporally-irregular Argos locations subject to measurement error and rasters of daily ocean surface currents collected between 20 November and 19 December 2012. Assuming a gamma distribution for step length (l_t) and a wrapped Cauchy distribution for turning angle (ϕ_t), we model the mean step length parameter (μ_t^l) as a function of ocean surface current speed (w_t) and direction (r_t) relative to the bearing of movement (b_t):

$$\mu_t^l = \exp(\beta_0^l + \beta_1^l w_t \cos(b_t - r_t)), \quad (3)$$

and the turning angle mean parameter (μ_t^ϕ) as a trade-off between short-term directional persistence and bias in the direction of ocean surface currents using the circular-circular regression link function:

$$\mu_t^\phi = \text{atan2}(\sin(d_t)\beta^\phi, 1 + \cos(d_t)\beta^\phi), \quad (4)$$

where $d_t = \text{atan2}(\sin(r_t - b_{t-1}), \cos(r_t - b_{t-1}))$.

We wish to fit a 2-state HMM to the turtle data, with a “foraging” state unaffected by currents and a “transit” state potentially influenced by ocean surface currents as in Eqs. 3 and 4. We used `crawlWrap` to predict $T = 350$ temporally-regular locations at 2 hr time steps assuming a bivariate normal measurement error model that accounts for the Argos location quality class (i.e. 3,2,1,0,A,B) of each observation. We then again used multiple imputation to account for location uncertainty by repeatedly fitting the HMM to `nSims` realizations of the position process using `MIfitHMM`. We first draw `nSims` realizations of the position process and extract the corresponding spatial covariates from the raster bricks for ocean surface current speed (“speedBrick”) and direction (“dirBrick”) using `MIfitHMM` with `fit=FALSE`:

```
miTurtleData <- MIfitHMM(crwOut, nSims = 100, fit=FALSE,
  spatialCovs = list(w = speedBrick, d = dirBrick, r = dirBrick),
  angleCovs = "d")
```

When the `fit` argument is `FALSE`, `MIfitHMM` returns a list of length `nSims` composed of `momentuHMMData` objects (`miData`). For convenience and ease of interpretation, we manually added an additional covariate ($angle_osc = \cos(b_t - r_t)$) to each of the imputed data sets and fitted the 2-state HMM using Eqs. 3 and 4 for state 2 (“transit”):

```

nbStates<-2
dist <- list(step = "gamma", angle = "wrpcauchy")
DM <- list(step = list(mean = ~state2(w:angle_osc), sd = ~1),
           angle = list(mean = ~state2(d), concentration= ~1))
turtleFits <- MifitHMM(miTurtleData$miData, nbStates = nbStates, dist = dist,
                      Par0 = Par0, DM = DM,
                      estAngleMean = list(angle = TRUE),
                      circularAngleMean = list(angle = TRUE))
plot(turtleFits, plotCI = TRUE, covs = data.frame(angle_osc = cos(0)))

```

Note that the `state2` special function in `DM` indicates the covariate formulas are specific to state 2 (“transit”) and the `circularAngleMean` argument indicates that circular-circular regression link function is to be used on the mean turning angle parameter as in Eq. 4.

For the “transit” state, pooled parameter estimates indicated step lengths increased with ocean surface current speed and as the bearing of movement aligned with ocean surface current direction ($\beta_1^l = 0.43$, 95% CI: 0.1 – 0.77; Figure 4). The estimated wrapped Cauchy distribution for turning angle had mean angles (μ_t^ϕ) biased towards the direction of ocean surface currents for each time step ($\beta^\phi = 0.24$, 95% CI: $-0.01 - 0.5$), with concentration parameter $\rho_2^\phi = 0.85$ (95% CI: 0.77–0.92) indicating turning angles were concentrated at μ_t^ϕ . Thus movement during the “transit” state appears to strongly follow ocean surface currents (mean $angle_osc = 0.87$, $SD = 0.23$), while movement during the “foraging” state exhibited shorter step lengths ($\mu_1^l = 2996\text{m}$, 95% CI: 1791 – 4202) perpendicular to ocean surface currents (mean $angle_osc = 0.07$, $SD = 0.26$), with no directional persistence ($\rho_1^\phi = 0.03$). The turtle spent 0.53 (95% CI: 0.34–0.71) of the 2 hr time steps in the “foraging” state and 0.47 (95% CI: 0.29–0.66) of time steps in the “transit” state as it travelled northeast along a predominant current until it (presumably) found an attractive foraging patch (Figure 4).

3.4 Grey seal

For our last example, we perform a similar analysis of a grey seal (*Halichoerus grypus*) track that was originally conducted by McClintock *et al.* (2012) using Bayesian methods and (computationally-intensive) Markov chain Monte Carlo. The data consist of 1045 temporally-irregular Fastloc GPS locations collected in the North Sea between 9 April and 11 August 2008. Because the seal repeatedly visited the same haulout and foraging

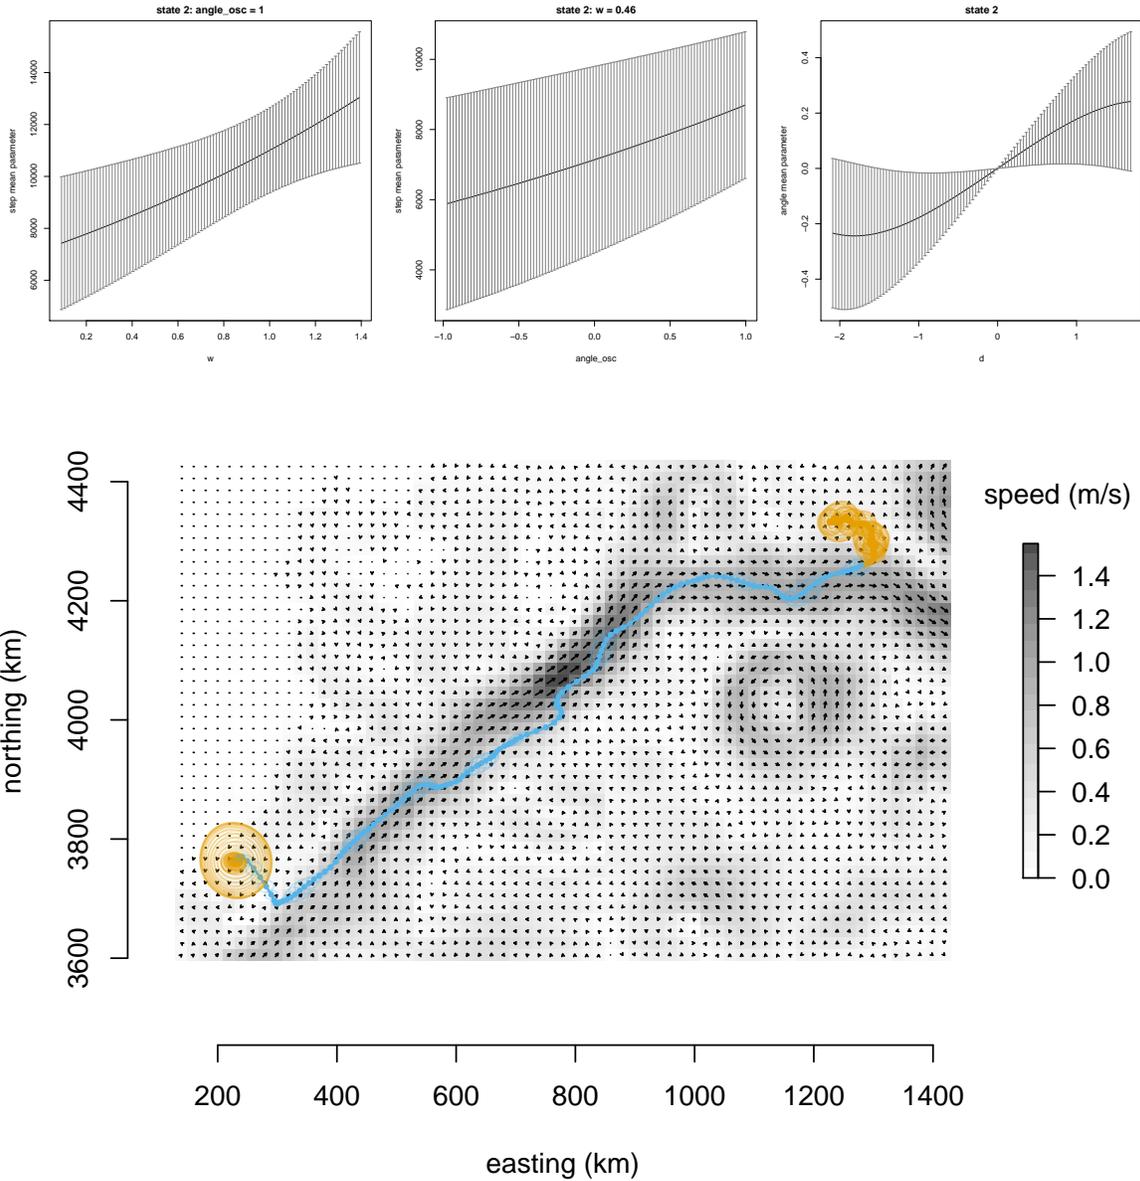


Figure 4. Selected results from the loggerhead turtle example. Top panels include estimates and 95% confidence intervals for the mean step length parameter as a function of ocean surface current speed (w) when ocean surface current direction (r_t) is the same as the bearing (b_t) of movement (i.e. $\text{angle_osc} = \cos(b_t - r_t) = 1$; top-left panel), mean step length parameter as a function of angle_osc at the mean ocean surface current speed ($w = 0.46$ m/s; top-middle panel), and mean turning angle parameter as a function of $d_t = \text{atan2}(\sin(r_t - b_{t-1}), \cos(r_t - b_{t-1}))$ (top-right panel). Bottom panel plots the pooled track, 95% error ellipse confidence bands, and state (orange = “foraging”, blue = “transit”) estimates based on multiple imputations of the position process relative to ocean surface current speed (m/s) and direction on 2 December 2012.

locations, it provides a nice example for demonstrating how to implement biased movements relative to activity centers using `momentuHMM`. McClintock *et al.* (2012) fitted a 5-state model to these data that included three center of attraction states, with movement biased towards two haulout sites (“Abertay” and “Farne Islands”) and a foraging area (“Dogger Bank”), and two “exploratory” states (“low speed”, “high speed”) that were unassociated with an activity center. After using `crawlWrap` to predict $T = 1515$ temporally-regular locations at 2 hr time steps including a bivariate normal measurement error model, we can perform a very similar analysis to McClintock *et al.* (2012) in `momentuHMM` by using the `centers` argument and state-specific functions for the probability distribution parameters. A cluster analysis on the observed locations using the R package `dtwclust` (Sarda-Espinosa 2017) identified three centroids with coordinates that were nearly identical to the three activity centers (“Abertay”, “Farne Islands”, and “Dogger Bank”) identified by McClintock *et al.* (2012). We use these coordinates to derive covariates relative to the activity centers when drawing `nSims` realizations of the position process:

```
crwSim <- MIfitHMM(crwOut, nSims = 100, fit=FALSE,
                  center = centers)
```

Specifying the `centers` argument results in the calculation of two covariates for each activity center: the distance (with ‘.dist’ suffix) and angle (with ‘.angle’ suffix) from each location at time t . These covariates can then be used to model parameters as a function of the distance and angle to activity centers for each time step:

```
dist <- list(step = "weibull", angle = "wrpcauchy")
distFormula <- ~state1(I(Abertay.dist>2500)) + state2(I(Farne.dist>2500))
              + state3(I(Dogger.dist>15000))
angleFormula <- ~state1(Abertay.angle) + state2(Farne.angle)
              + state3(Dogger.angle)
stepDM <- list(shape = distFormula, scale = distFormula)
angleDM <- list(mean = angleFormula, concentration = distFormula)
DM <- list(step = stepDM, angle = angleDM)
```

Similar to McClintock *et al.* (2012), we assume a Weibull distribution for step length where both the shape and scale parameter depend on the distance from the location at time t to each activity center. For the activity centers on land (“Abertay” and

“Farne”), we allow the (state-dependent) step length parameters to change when the seal is beyond 2500m of the haulout. For the “Dogger” activity center, we allow the parameters to change when the seal is beyond 15000m of this (presumably) foraging area. We thus allow the movement behavior to change within these activity center states upon entering or leaving the vicinity of these sites. We assume a wrapped Cauchy distribution for turning angle with (state-dependent) mean angle derived from the direction to each activity center at time t , and the concentration parameter is modeled similarly to the step length parameters. For the two “exploratory” states, we assumed they are simple random walks unaffected by proximity to activity centers. To complete our model specification, we use the `knownStates` argument to assign the seal to the corresponding activity center state whenever it was within the 2500m (haulout area) or 15000m (foraging area) thresholds for each imputed data set:

```
greySealFits <- MIfitHMM(miDat, nSims = 400,
                        nbStates = 5, dist = dist,
                        Par0 = Par0, beta0 = beta0, fixPar = fixPar,
                        formula = distFormula,
                        estAngleMean = list(angle=TRUE),
                        circularAngleMean = list(angle=TRUE),
                        DM = DM, knownStates = knownStates)
plot(greySealFits, plotCI = TRUE)
```

As with the step length and turning angle concentration parameters, the state transition probabilities are also allowed to change as a function of distance to activity centers (as specified by the `formula` argument). The starting values (`Par0` and `beta0`) for each imputation were extracted from a single HMM fitted to the best predicted locations from `crawlWrap`, and `fixPar` was used to remove short-term directional persistence (and thus formulate the model as a mixture of biased and simple random walks).

Estimated activity budgets for the 5 states of this multiple imputation HMM were 0.28 (0.27 – 0.3) for the “Abertay” haul-out state, 0.12 (0.11 – 0.14) for the “Farne Islands” haul-out state, 0.37 (0.35 – 0.38) for the “Dogger Bank” foraging state, 0.09 (0.03 – 0.2) for a low-speed “exploratory” state, and 0.14 (0.08 – 0.23) for a high-speed “exploratory” state. All three activity center states exhibited shorter step lengths and less biased movements when within the vicinity of these targets (Figure 5). Results from this analysis were thus very similar to those of McClintock *et al.* (2012), but

this implementation required far less computation time and no custom model-fitting algorithms.

The `simData` function can be used to simulate tracks from a fitted model:

```
greySealSim<-simData(model = greySealFits, centers = centers,  
                    initialPosition = centers[1,],  
                    obsPerAnimal = 1515)
```

A simulated track is presented along with the fitted track in Figure 6. While potentially useful for study design, power analysis, and prediction, the `simData` function can also be helpful in assessing goodness of fit by repeatedly drawing simulated data sets from a fitted model and comparing them to observed properties of the data (e.g. Morales *et al.* 2004).

4 Discussion

Here we have introduced version 1.0.0 of the R package `momentuHMM` and demonstrated some of its capabilities for conducting multivariate HMM analyses with animal location, auxiliary biotelemetry, and environmental data. The package allows for fitting (and simulating from) a suite of biased and correlated random walk movement process models (e.g. McClintock *et al.* 2012), can be used for an unlimited number of data streams and latent behavior states, includes multiple imputation methods to account for measurement error, temporal irregularity, and other forms of missing data that would otherwise be prohibitive to maximum likelihood analysis, and integrates seamlessly with rasters to facilitate spatio-temporal covariate modelling. Because the package incorporates biased random walks, it can also be used to implement group dynamic models (e.g. Langrock *et al.* 2014). The package therefore greatly expands on available software and facilitates the incorporation of more ecological and behavioral realism for hypothesis-driven analyses of animal movement that account for many of the challenges commonly associated with telemetry data. While many of the features of `momentuHMM` were motivated by animal movement data, we note that the package is not limited to location data and can be used for analyzing any type of data that is amenable to (multivariate) HMMs.

Model fitting in `momentuHMM` is relatively fast because the forward algorithm (Eq. 1) is coded in C++. Because multiple imputations are completely parallelizable, with

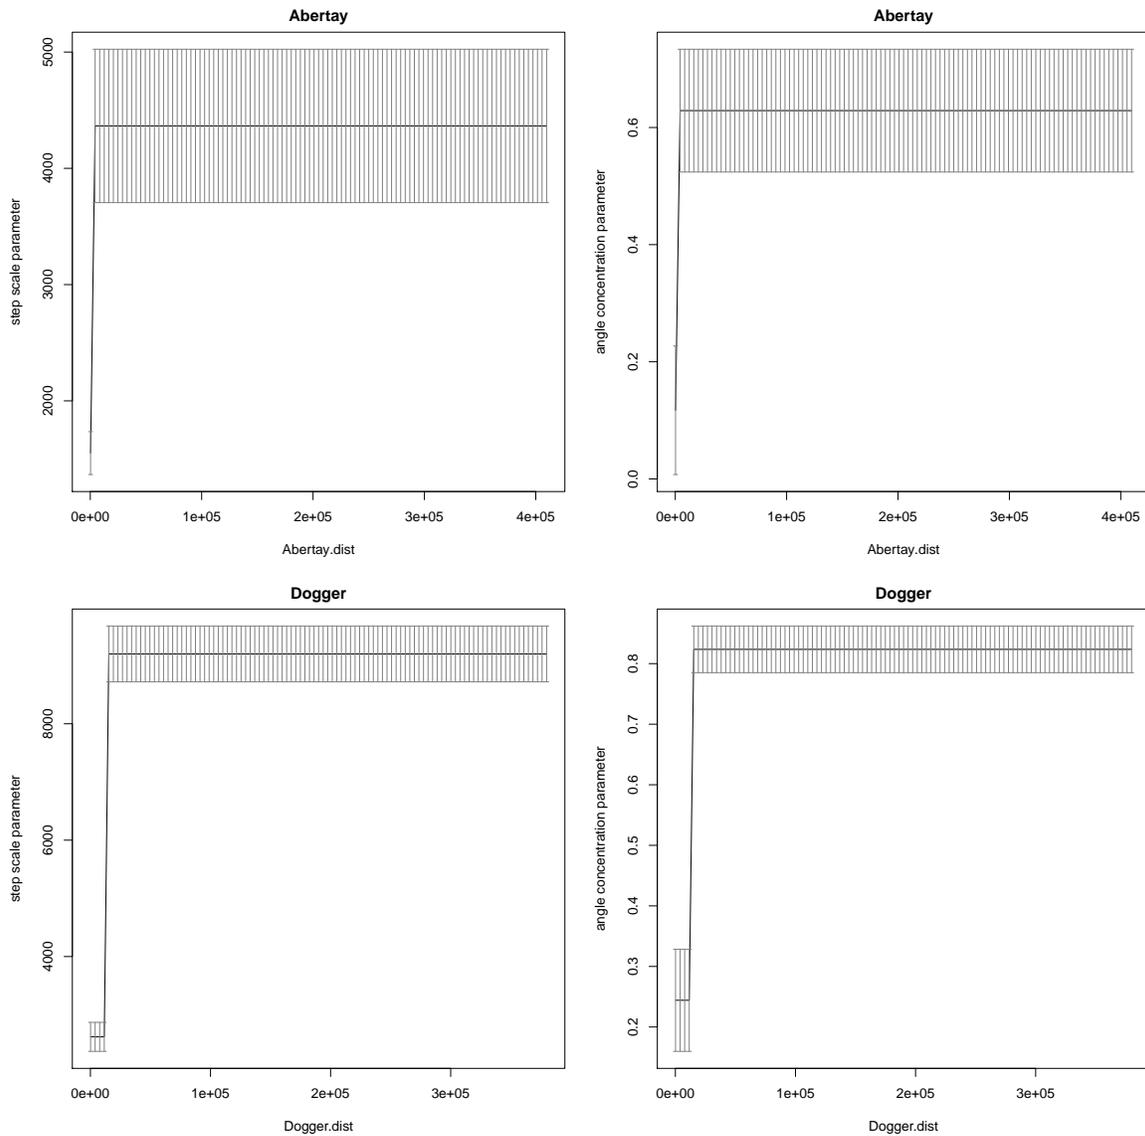


Figure 5. Selected results from the grey seal example. Panels include estimates and 95% confidence intervals for the “Abertay” haul-out state step length scale parameter as a function of distance in meters (‘Abertay.dist’; top-left panel), “Abertay” haul-out state turning angle concentration parameter as a function of distance (top-right panel), “Dogger Bank” foraging state step length scale parameter as a function of distance (‘Dogger.dist’; bottom-left panel), and the “Dogger Bank” foraging state turning angle concentration parameter as a function of distance (bottom-right panel).

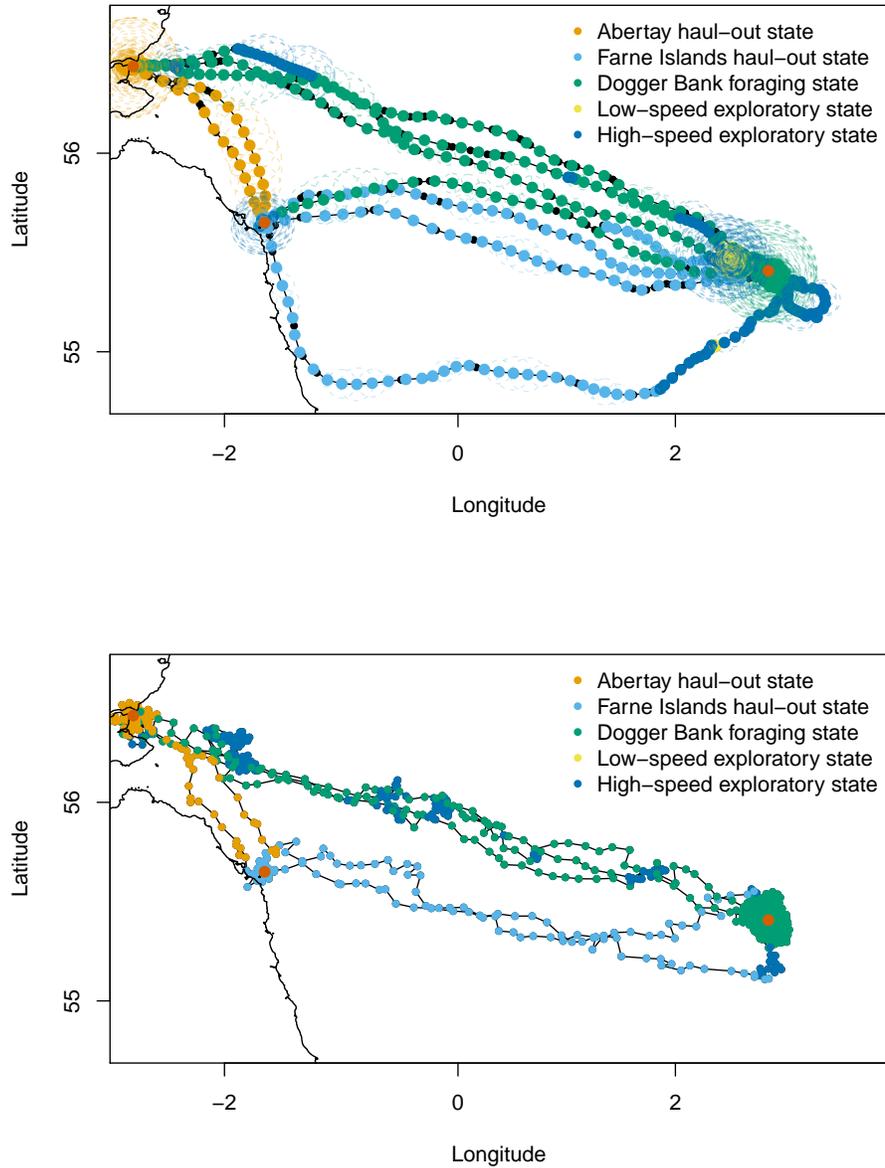


Figure 6. Fitted and simulated tracks from the grey seal example. This seal tended to move in a clockwise fashion between two haul-out locations (“Abertay” and “Farne Islands”) and a foraging area (“Dogger Bank”) in the North Sea. Top panel plots the pooled track, 95% error ellipse confidence bands, and state estimates based on the 5-state HMM fitted to multiple imputations of the position process. Red points indicate the locations of the three activity centers. Black points indicate the (temporally-irregular) observed locations. Bottom panel presents the locations and states for a track simulated from the fitted model using the ‘simData’ function.

sufficient processing power computation times for analyses that account for measurement error, temporal irregularity, or other forms of missing data need not be longer than that required to fit a single HMM. However, computation times will necessarily be slower as the number of states and/or parameters increase. For example, `momentuHMM` required about 1 hr to fit a single HMM with $N = 6$ states, seven data streams, and $T = 7414$ time steps (McClintock 2017).

As in any maximum likelihood analysis based on numerical optimization, computation times will also depend on the starting values (`Par0` and `beta0`). Specifying “good” starting values is arguably the most challenging aspect of model fitting in `momentuHMM`, particularly for the working scale coefficients when using covariates. The `getPar`, `getPar0`, and `getParDM` functions are designed to help with the specification of starting values, and the `retryFits` argument in `crawlWrap`, `fitHMM`, and `MIfitHMM` will re-optimize based on random perturbations of the parameters to help explore the likelihood surface and diagnose convergence to local maxima. Optimization for the circular-linear regression link function (`tan(mean/2)`; see Table 2) in particular can be prone to local minima, so users are encouraged to explore a range of starting values when fitting these models.

While `momentuHMM` includes functions for drawing realizations of the position process based on the CTCRW model of Johnson *et al.* (2008), this is but one of many methods for performing the first stage of multiple imputation. Realizations of the position process from any movement model that accounts for measurement error and/or temporal irregularity (e.g. Calabrese *et al.* 2016; Gurarie *et al.* 2017) could be passed to `MIfitHMM` for HMM-type analyses in the second stage. Multiple imputation methods also need not be limited to these telemetry error scenarios. For example, conventional missing data could also be imputed using standard techniques (Rubin & Schenker 1986), thereby allowing the investigation of non-random mechanisms for missingness that can be problematic if left unaccounted for in HMMs.

There remain many potential avenues for refining and extending the capabilities of `momentuHMM`. Computation times could likely be improved by further optimizing the R and C++ code for speed. Notable extensions include hidden semi-Markov models and random effects on data stream probability distribution and state transition probability parameters (Zucchini *et al.* 2016). We would also like to incorporate additional parameters for change-point thresholds and the locations of activity centers instead

of requiring that they be pre-specified (and potentially compared using AIC or other model selection criteria) as in grey seal example. Lastly, it is relatively straightforward to add additional probability distributions, and we are pleased to do so upon request. Practitioners interested in additional features for `momentuHMM` are encouraged to contact the authors.

Acknowledgments

The findings and conclusions in the manuscript are those of the author(s) and do not necessarily represent the views of the National Marine Fisheries Service, NOAA. Any use of trade, product, or firm names does not imply an endorsement by the US Government.

References

- Beyer, H.L., Morales, J.M., Murray, D. & Fortin, M.J. (2013) The effectiveness of bayesian state-space models for estimating behavioural states from movement paths. *Methods in Ecology and Evolution*, **4**, 433–441.
- Calabrese, J.M., Fleming, C.H. & Gurarie, E. (2016) `ctmm`: an R package for analyzing animal relocation data as a continuous-time stochastic process. *Methods in Ecology and Evolution*, **7**, 1124–1132.
- Cornelissen, G. (2014) Cosinor-based rhythmometry. *Theoretical Biology and Medical Modelling*, **11**, 16.
- Costa, D.P., Robinson, P.W., Arnould, J.P., Harrison, A.L., Simmons, S.E., Hassrick, J.L., Hoskins, A.J., Kirkman, S.P., Oosthuizen, H., Villegas-Amtmann, S. *et al.* (2010) Accuracy of argos locations of pinnipeds at-sea estimated using fastloc gps. *PloS one*, **5**, e8677.
- DeRuiter, S.L., Langrock, R., Skirbutas, T., Goldbogen, J.A., Chalambokidis, J., Friedlaender, A.S. & Southall, B.L. (2017) A multivariate mixed hidden markov model to analyze blue whale diving behaviour during controlled sound exposures. *The Annals of Applied Statistics*, **11**, 362–392.
- Duchesne, T., Fortin, D. & Rivest, L.P. (2015) Equivalence between step selection func-

- tions and biased correlated random walks for statistical inference on animal movement. *PloS one*, **10**, e0122947.
- Gilbert, P. & Varadhan, R. (2016) *numDeriv: Accurate Numerical Derivatives*. R package version 2016.8-1.
- Gurarie, E., Fleming, C.H., Fagan, W.F., Laidre, K.L., Hernández-Pliego, J. & Ovaskainen, O. (2017) Correlated velocity models as a fundamental unit of animal movement: synthesis and applications. *Movement Ecology*, **5**, 13.
- Hijmans, R.J. (2016) *raster: Geographic Data Analysis and Modeling*. R package version 2.5-8.
- Hooten, M.B., Johnson, D.S., McClintock, B.T. & Morales, J.M. (2017) *Animal Movement: Statistical Models for Telemetry Data*. CRC Press.
- Johnson, D.S. (2017) *crawl: Fit Continuous-Time Correlated Random Walk Models to Animal Movement Data*. R package version 2.1.2.
- Johnson, D.S., London, J.M., Lea, M.A. & Durban, J.W. (2008) Continuous-time correlated random walk model for animal telemetry data. *Ecology*, **89**, 1208–1215.
- Jonsen, I.D., Flemming, J.M. & Myers, R.A. (2005) Robust state–space modeling of animal movement data. *Ecology*, **86**, 2874–2880.
- Langrock, R., Hopcraft, G., Blackwell, P., Goodall, V., King, R., Niu, M., Patterson, T., Pedersen, M., Skarin, A. & Schick, R. (2014) Modelling group dynamic animal movement. *Methods in Ecology and Evolution*, **5**, 190–199.
- Langrock, R., King, R., Matthiopoulos, J., Thomas, L., Fortin, D. & Morales, J.M. (2012) Flexible and practical modeling of animal telemetry data: hidden markov models and extensions. *Ecology*, **93**, 2336–2342.
- McClintock, B.T. (2017) Incorporating telemetry error into hidden markov models of animal movement using multiple imputation. *Journal of Agricultural, Biological, and Environmental Statistics*, **in press**.

- McClintock, B.T., Johnson, D.S., Hooten, M.B., Ver Hoef, J.M. & Morales, J.M. (2014) When to be discrete: the importance of time formulation in understanding animal movement. *Movement Ecology*, **2**, 21.
- McClintock, B.T., King, R., Thomas, L., Matthiopoulos, J., McConnell, B.J. & Morales, J.M. (2012) A general discrete-time modeling framework for animal movement using multistate random walks. *Ecological Monographs*, **82**, 335–349.
- McClintock, B.T., London, J.M., Cameron, M.F. & Boveng, P.L. (2017) Bridging the gaps in animal movement: hidden behaviors and ecological relationships revealed by integrated data streams. *Ecosphere*, **8**, e01751.
- McClintock, B.T., Russell, D.J., Matthiopoulos, J. & King, R. (2013) Combining individual animal movement and ancillary biotelemetry data to investigate population-level activity budgets. *Ecology*, **94**, 838–849.
- Michelot, T., Langrock, R. & Patterson, T.A. (2016) moveHMM: An R package for the statistical modelling of animal movement data using hidden markov models. *Methods in Ecology and Evolution*, **7**, 1308–1315.
- 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.
- R Core Team (2016) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rubin, D.B. & Schenker, N. (1986) Multiple imputation for interval estimation from simple random samples with ignorable nonresponse. *Journal of the American Statistical Association*, **81**, 366–374.
- Sarda-Espinosa, A. (2017) *dtwclust: Time Series Clustering Along with Optimizations for the Dynamic Time Warping Distance*. R package version 3.1.2.
- Wall, J., Wittemyer, G., LeMay, V., Douglas-Hamilton, I. & Klinkenberg, B. (2014) Elliptical time-density model to estimate wildlife utilization distributions. *Methods in Ecology and Evolution*, **5**, 780–790.

- Whoriskey, K., Auger-Méthé, M., Albertsen, C.M., Whoriskey, F.G., Binder, T.R., Krueger, C.C. & Mills Flemming, J. (2017) A hidden markov movement model for rapidly identifying behavioral states from animal tracks. *Ecology and Evolution*, **7**, 2112–2121.
- Zucchini, W., MacDonald, I.L. & Langrock, R. (2016) *Hidden Markov Models for Time Series: An Introduction Using R*. CRC Press.