

moveHMM

An R package for animal movement modelling

Michelot T., Langrock R., Patterson T., and Rexstad E.

October 22, 2015

1 Introduction

The analysis of animal movement data has become increasingly important in terrestrial and marine ecology. Improvements in telemetry technology have resulted in an explosion in the volume of high precision data being collected. As a result, there are two challenges which researchers collecting these data regularly face: (1) data volume and (2) employing statistical methods which can accommodate some of the specific features of movement data (Patterson et al., 2009).

A substantial part of the literature on statistical modelling of animal movement data has focused on the intuitive approach of decomposing movement time series into distinct behavioural modes (a.k.a. bouts, states), via the use of so-called state-switching models. These approaches typically involve assuming movements of animals to be driven by stochastically evolving states, such as a slow moving state, which may be indicative of resting or foraging, versus faster movement states which might indicate transits between foraging patches. Associated with changes in movement speeds are changes to the distribution of directional changes in the movement (known as the turning angle — see further description below).

Bayesian methods which employ MCMC approaches have become very popular tools for the analysis of movement data using state-switching models (e.g. Jonsen et al., 2005; Morales et al., 2004). Typically, these have been implemented using WinBUGS (although see McClintock et al., 2012). While these models are relatively straightforward to build and hence fit in WinBUGS, the estimation can be painfully slow due to slow mixing of the MCMC samplers.

However, for an important subset of movement data, namely highly accurate position data (e.g. from GPS) — and more generally all time series of locations where the measurement error is negligible relative to the scale of the movement — the task of statistical classification of behaviour can be done much more efficiently using hidden Markov models (HMMs) and associated frequentist inferential tools. HMMs are increasingly popular in this field, due to their flexibility and to the associated very efficient recursive algorithms available for conducting statistical inference (Patterson et al., 2009; Langrock et al., 2012).

The crucial requirements on movement data in order for HMMs to be suitable are that measurement error in positions is negligible and that there is a regular sampling unit (e.g. one positional observation per hour, or per dive, or any other meaningful unit).

`moveHMM` is an R package which implements HMMs and associated tools for state decoding, model selection etc. specifically tailored to animal movement modelling. Particular attention was paid to computational efficiency with the fitting algorithm implemented in C++. The high computational speed makes it feasible to analyze very large data sets — e.g. tens of thousands of positions collected for each of a dozen individual animals — on standard desktop PCs. The package also allows users to incorporate covariate data into their models, which is particularly useful when inferring the drivers of changes in behaviour.

Our hope is that the `moveHMM` package will provide users who collect movement data with an interface to sophisticated and adequate methods for a statistical analysis of their data. The package is structured so as to allow the users to prepare their data for analysis, fit a variety of HMMs to their data and perform diagnostics on these fitted models.

In this vignette, we briefly introduce HMMs in the context of animal movement. We then provide a detailed example of a typical use of the package (preprocessing movement data, fitting an HMM to the data, and analyzing the fitted model). Finally, we describe more technically the structure of the package and its main functions.

2 A quick overview of the HMM approach to modelling animal movement

For full details of the HMM approach in the context of animal movement modelling, the user should refer to the relevant primary publications (see, e.g., Patterson et al., 2009, Langrock et al., 2012, Zucchini and MacDonald, 2009). Here, we briefly highlight the essential features of the HMM approach.

The standard HMM approach to model an individual animal's movement considers bivariate time series comprising the step length and the turning angle at each time point (see illustration in Figure 1). The associated locations need to be *sampled at equally spaced points in time* (though missing data on an otherwise regular grid can easily be handled) and are assumed to be *observed with no or only negligible error*. The `moveHMM` package is restricted to such discrete time data and involves the assumption that the locations are observed with zero, or at least negligible error.

At each time point, the parameters of the step length distribution (e.g., a gamma distribution) and the parameters of the turning angle distribution (e.g., a von Mises distribution) are determined by an underlying unobserved state. There are finitely many states which provide rough classifications of the movement (e.g. more active vs. less active), often interpreted as proxies for the animal's behavioural states (e.g. transiting vs. foraging). The sequence of

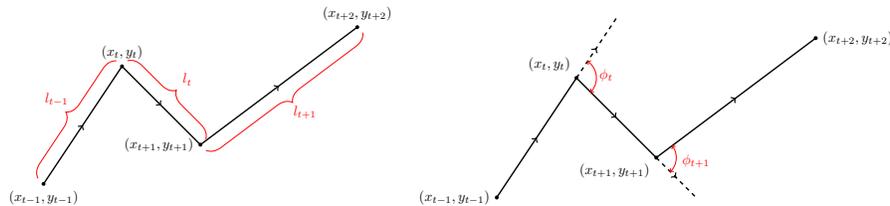


Figure 1: Illustration of step lengths and turning angles

states is assumed to be generated by a Markov chain, usually with a tendency of remaining in a state for some time before switching to another state. The corresponding state transition probabilities, as well as the parameters characterising the state-dependent distributions, are model parameters to be estimated. The number of states (movement modes) is unknown and has to be specified by the user. Typically with movement data one assumes a low number of states (say ≤ 4). For example, an animal may be foraging (low speed, high rates of turning) and transiting (high speeds, low rates of turning). Biologically interesting inference often involves modelling the state transition probabilities as functions of environmental covariates.

For a given set of model parameters, the likelihood of the data can be calculated using a recursive algorithm (the forward algorithm, cf. Patterson et al., 2015), which in a very effective way considers all possible state sequences that might have given rise to the observed time series. This makes numerical maximization of the (log-)likelihood, and hence maximum likelihood estimation, feasible in most cases. Having estimated the model parameters (and examined useful diagnostics of model fit), the user can estimate the most likely sequence of behavioural states and also examine various variables of interest such as the average rate of movement in each state. We encourage the user to consult a good primary text such as Zucchini and MacDonald (2009) to familiarize themselves with the technical details of hidden Markov models.

3 Example application

Before we provide a detailed description of the various features of the `moveHMM` package in the subsequent section, we illustrate a typical HMM-based analysis of movement data using the main functions of the package, via an example. We use the data from Morales et al. (2004), collected on four elk in Canada.

3.1 Movement data

The input data need to have the correct format for subsequent processing and analysis. The data need to be provided as a `data.frame`, with two mandatory columns:

- Easting or longitude (default name: `x`)

- Northing or latitude (default name: y)

It is possible to have a column “ID”, which contains the identifiers of the observed animals. If no column named “ID” is provided, all the observations will be considered to belong to a single animal. Additional columns are considered as covariates. Note that covariates need to have numerical values, rather than character values.

3.1.1 Load and format the tracking data

The elk data considered in Morales et al. (2004) are available online at:

http://www.esapubs.org/archive/ecol/E085/072/elk_data.txt

We load the relevant rows and columns into a data frame using the `read.table` command.

```
trackData <- read.table(
  "http://www.esapubs.org/archive/ecol/E085/072/elk_data.txt",
  sep="\t", header=TRUE) [1:735, c(1,2,3,7)]
```

The dataframe `trackData` now has four columns: “Individual”, “Easting”, “Northing”, and “dist_water..meters.”. The last one is the distance of the animal to water, which for illustration purposes we want to include in the model as a covariate.

```
> head(trackData)
  Individual Easting Northing dist_water..meters.
1   elk-115  769928  4992847           200.00
2   elk-115  766875  4997444           600.52
3   elk-115  765949  4998516           561.81
4   elk-115  765938  4998276           550.00
5   elk-115  766275  4998005           302.08
6   elk-115  766368  4998051           213.60
```

The animals’ identifiers need to be stored in a column named “ID”, before we can preprocess the data. Thus, we modify it accordingly. We also shorten the name of the covariate.

```
colnames(trackData)[1] <- "ID"
colnames(trackData)[4] <- "dist_water"
```

Besides, we decide that we want to deal with distances in kilometers, instead of meters, for the step lengths.

```
trackData$Easting <- trackData$Easting/1000
trackData$Northing <- trackData$Northing/1000
```

As a result, this is what the data looks like:

```
> head(trackData)
  ID Easting Northing dist_water
1 elk-115 769.928 4992.847    200.00
2 elk-115 766.875 4997.444    600.52
```

```

3 elk-115 765.949 4998.516 561.81
4 elk-115 765.938 4998.276 550.00
5 elk-115 766.275 4998.005 302.08
6 elk-115 766.368 4998.051 213.60

```

3.1.2 Use prepData

With the data in the proper format, it can be processed using `prepData` to compute step lengths and angles. We choose the arguments carefully:

- `type` specifies whether the coordinates are easting/northing (`type="UTM"`) or longitude/latitude (`type="LL"`) values. The latter is the default, so we need to call the function with the argument `type="UTM"`, to indicate that UTM coordinates are provided.
- `coordNames` are the names of the coordinates in the input dataframe. The default is “x” and “y”, so we need to call the function with the argument `coordNames=c("Easting", "Northing")`.

The call to the function for these data is,

```

data <- prepData(trackData, type="UTM", coordNames=c("Easting",
  Northing"))

```

The step lengths and turning angles are computed, and the returned object is a data frame.

```

> head(data)
  ID      step      angle      x      y dist_water
1 elk-115 5.5184434      NA 769.928 4992.847    200.00
2 elk-115 1.4165663 0.1262112 766.875 4997.444    600.52
3 elk-115 0.2397525 2.3832412 765.949 4998.516    561.81
4 elk-115 0.4327600 0.9385238 765.938 4998.276    550.00
5 elk-115 0.1037545 1.1375066 766.275 4998.005    302.08
6 elk-115 12.4164659 -0.9687435 766.368 4998.051    213.60

```

Note that the coordinates have been renamed “x” and “y”. This makes the processing of the data simpler.

If the data set contains covariates which have missing values, then those are imputed using the closest non-missing value, by default the previous one if it is available. This is arbitrary and might not be appropriate in all situations.

It is also possible to print summary information about the data, using the function `summary`, e.g.

```

> summary(data)
Movement data for 4 animals:
elk-115 -- 194 observations
elk-163 -- 159 observations
elk-287 -- 164 observations
elk-363 -- 218 observations

```

```
Covariate(s):
dist_water
  Min.      25%      Median      Mean      75%      Max.
0.0000  213.6000  477.6200  773.6457  1169.4950  3781.0400
```

3.1.3 Use plot.moveData

Once the data have been preprocessed, they can be plotted using the generic function `plot`. This displays maps of the animals' tracks, times series of the steps and angles, and histograms of the steps and angles. A few plotting options are available. These are described in the documentation. To plot all animals' tracks on a single map, we call:

```
plot(data, compact=T)
```

The resulting map, and the steps and angles graphs for the first animal, are displayed in Figure 2 (we omit the graphs for the three other animals, which are also displayed when using the above command).

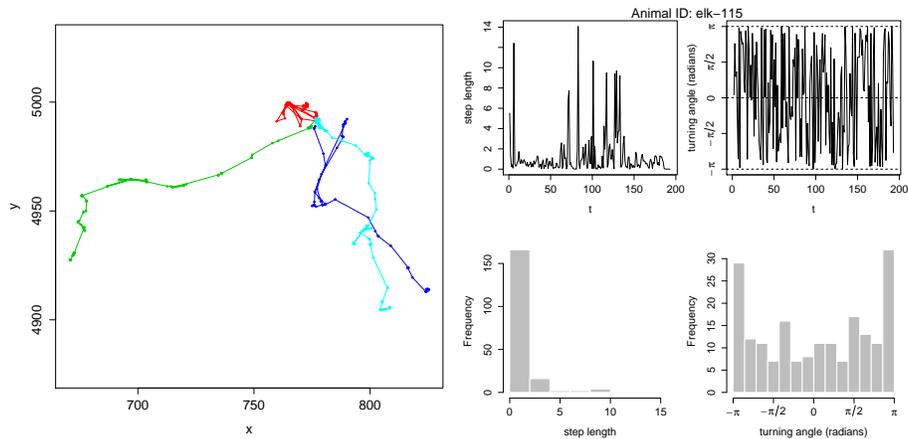


Figure 2: Map of the animals' tracks (left) – each color represents an animal. Graphs of the step lengths and turning angles for one individual, “elk-115” (right).

The time series of step lengths is one way to check the data for outliers.

3.2 Fitting the model

3.2.1 fitHMM

The function `fitHMM` is used to fit an HMM to the data. Its arguments are described in the documentation. Here are a few choices we make:

- `nbStates=2`, i.e. we fit a 2-state HMM to the data;

- `beta0=NULL` and `delta0=NULL`, i.e. we use the default values for the initial values `beta0` and `delta0`;
- `formula=~dist_water`, i.e. the transition probabilities are functions of the covariate “dist_water”;
- `stepDist="gamma"`, to model the step lengths with the gamma distribution (note that it is the default, so we do not need to explicitly specify it);
- `angleDist="vm"`, to model the turning angles with the von Mises distribution (default);
- `angleMean=NULL`, because we want to estimate the mean of the angle distribution (default);
- `stationary=FALSE`, as due to the covariates the process is not stationary (default).

We also need to specify initial values for the parameters of the state-dependent distributions, to be used by the optimization function. Note that this choice is crucial, and that **the algorithm might not find the global optimum of the likelihood function if the initial parameters are poorly chosen**. The initial parameters should be specified in two vectors, `stepPar0` (for the step distribution) and `anglePar0` (for the angle distribution). The necessary parameters of each distribution are detailed in Section 4.1.1.

Zero-inflation must be included in the step length distribution if some steps are of length exactly zero (which is the case for the elk data). To do so, another parameter is added to the step distribution: its mass on zero.

```
# initial parameters for gamma and von Mises distributions
mu0 <- c(0.1,1) # step mean (two parameters: one for each state)
sigma0 <- c(0.1,1) # step SD
zeromass0 <- c(0.05,0) # step zero-mass
stepPar0 <- c(mu0,sigma0,zeromass0)
angleMean0 <- c(pi,0) # angle mean
kappa0 <- c(1,1) # angle concentration
anglePar0 <- c(angleMean0,kappa0)

# call to fitting function
m <- fitHMM(data=data,nbStates=2,stepPar0=stepPar0,
            anglePar0=anglePar0,formula=~dist_water)
```

The returned object, `m`, is of the class `moveHMM`. It can be printed in order to obtain the maximum likelihood estimates of all model parameters.

```
> m
Value of the maximum log-likelihood: -1893.092

Step length parameters:
-----
mean
[1] 0.3580483 3.3556845
```

```

sd
[1] 0.3813433 4.3565876
zero-mass
[1] 1.943569e-03 9.999988e-09

Turning angle parameters:
-----
mean
[1] -2.9940947 0.1461898
concentration
[1] 0.6011242 0.2284411

Regression coeffs for the transition probabilities:
-----
              1 -> 2          2 -> 1
intercept  -1.7869761202 -1.883796894
dist_water -0.0004087263  0.001450852

Initial distribution:
-----
[1] 0.4256565 0.5743435

```

3.2.2 Dealing with numerical instability

As mentioned above, the numerical maximization routine might not identify the global maximum of the likelihood function, or even fail to converge altogether, for poorly chosen initial values of the parameters. In such a case, the optimization routine `nlm` might produce an error such as:

```

Error in nlm(nLogLike, wpar, nbStates, bounds, parSize, data,
  stepDist, :
  non-finite value supplied by 'nlm'

```

The best way to deal with such numerical problems is to test different sets of initial values, possibly chosen randomly. By comparing the resulting estimates for the different initial values used, one usually obtains a good feeling for any potential sensitivity of the numerical search to its chosen starting point. Note, however, that in any case there will usually be no certainty that the global maximum of the likelihood, i.e. the maximum likelihood estimate, has been identified.

During preliminary tests on the elk data, we noticed that, in this example, the numerical search is highly sensitive to the choice of the initial parameters `beta0`. This is due to the high values of the covariate: a small change in the associated regression coefficients can make a big difference in the likelihood function. In such cases, it is advisable to standardize the covariate values before fitting the model, for example by calling:

```

data$dist_water <-
  (data$dist_water - mean(data$dist_water)) / sd(data$dist_water)

```

This allows for greater numerical stability, with the convergence of the fitting function depending less on the choice of initial values. The value of the

maximum log-likelihood is not affected by the standardization of the covariate values, only the maximum likelihood estimate of `beta` is.

3.2.3 Confidence intervals

Confidence intervals for the model parameters can be computed with the function `CI`, passing as an argument the object created by `fitHMM`. It is possible to give the significance level of the desired confidence interval as an argument, e.g. 0.99 for 99% confidence intervals. By default, 95% confidence intervals are returned.

Below we show the 95% confidence intervals for the parameters of the 2-state model fitted to the elk data. `CI(m)$stepPar` corresponds to the bounds of the confidence intervals for the step parameters (`CI(m)$stepPar$lower` and `CI(m)$stepPar$upper`). `CI(m)$anglePar` and `CI(m)$beta` are respectively the bounds of the confidence intervals for the angle parameters and the regression coefficients of the transition probabilities.

```
> CI(m)
$stepPar
$stepPar$lower
      [,1]      [,2]
[1,] 0.3008217103 2.635316e+00
[2,] 0.3101405634 3.478951e+00
[3,] 0.0002682572 9.994941e-09

$stepPar$upper
      [,1]      [,2]
[1,] 0.42616138 4.272967e+00
[2,] 0.46889284 5.455626e+00
[3,] 0.01393564 1.000504e-08

$anglePar
$anglePar$lower
      [,1]      [,2]
[1,] -3.232034 -1.21994842
[2,]  0.455092  0.05797161

$anglePar$upper
      [,1]      [,2]
[1,] -2.7556474 1.7238709
[2,]  0.7558124 0.5103924

$beta
$beta$lower
      [,1]      [,2]
[1,] -2.606487342 -3.0489598624
[2,] -0.000968507  0.0002106212

$beta$upper
      [,1]      [,2]
[1,] -0.9674648979 -0.718633927
[2,]  0.0001510544  0.002691082
```

Note that a warning message is also output:

```
Warning message:
In CI(m) :
Some of the parameter estimates seem to lie close to the boundaries
of their parameter space. The associated CIs are probably
unreliable (or might not be computable).
```

It here refers to the zero-inflation parameter in the second state. Its estimate is very close to zero, the inferior boundary of its range (this parameter is in the interval $[0,1]$), and this causes the corresponding confidence interval to be unreliable. The function can sometimes fail to compute such a confidence interval, and would return `NA` instead.

3.3 Further inference tools and visualization of the model

Various options are available for the class `moveHMM`, and here we explain how to use them in the `elk` example.

3.3.1 Plot the model

The fitted model can be plotted, using the generic function `plot`. A few graphical options are available and listed in the documentation. Here, we call:

```
plot(m)
```

This outputs:

- an histogram of step lengths of all animals, with the fitted state-dependent densities,
- an histogram of turning angles of all animals, with the fitted state-dependent densities,
- plots of the transition probabilities as functions of the covariate considered,
- a map of each animal's track, colored by states.

Figure 3 displays those plots, but showing only one of the plotted maps, namely the one corresponding to the first animal, “`elk-115`”. The state-dependent densities are weighted by the relative frequency of each state in the most probable state sequence (decoded with the Viterbi algorithm, see Section 3.3.2). For example, if according to the most probable state sequence one third of the observations is allocated to the first state, and two thirds to the second state, the plots of the densities in the first state are weighted with a factor $1/3$, and in the second state with a factor $2/3$.

The first state (in red on the plots) corresponds to short steps, and angles centered around π , and the second state (in green on the plots) corresponds to longer steps, and angles centered around 0.

The plots of the transition probabilities as function of the covariate indicate that animals tend to switch from the second state to the first state when they are far from water, whereas they stay in the second state when closer to water.

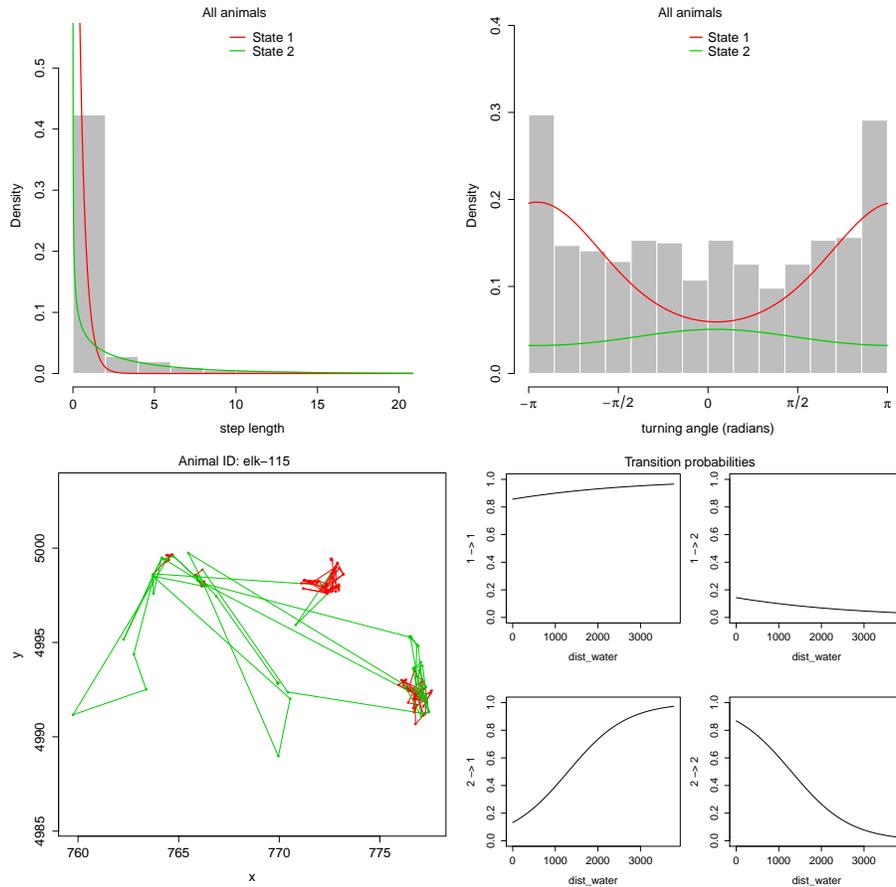


Figure 3: Output of `plot.moveHMM`. Histogram of step lengths with fitted distributions (top-left), histogram of turning angles with fitted distributions (top-right), map of decoded track for the first animal (bottom-left), and transition probabilities as functions of “`dist_water`” (bottom-right).

3.3.2 State decoding

Two functions can be used to decode the state process.

Viterbi algorithm To globally decode the state process, the Viterbi algorithm (Zucchini and MacDonald, 2009) is implemented in the function `viterbi`. This function outputs the most likely sequence of states to have generated the observation, under the fitted model. Below are the most probable states for the first 25 observations of the first individual:

```
> states <- viterbi(m)
> states[1:25]
[1] 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
```

State probabilities To get more accurate information on the state process, it is possible to compute the state probabilities for each observation, using `stateProbs`. This returns a matrix with as many columns as there are states in the model, and as many rows as there are observations (stacking all animals' observations). The elements of the matrix are defined as

$$\text{stateProbs}(m)[t,j] = \Pr(S_t = j)$$

where $\{S_t\}$ is the state process.

For example:

```
> sp <- stateProbs(m)
> head(sp)
      [,1]      [,2]
[1,] 1.632239e-05 0.9999837
[2,] 9.911030e-02 0.9008897
[3,] 6.024606e-01 0.3975394
[4,] 5.647906e-01 0.4352094
[5,] 4.224738e-01 0.5775262
[6,] 4.447990e-11 1.0000000
```

The function `plotStates` can be used to visualize the results of `viterbi` and `stateProbs`. Figure 4 shows the plots of the most likely state sequence decoded by the Viterbi algorithm, as well as both columns of the matrix of state probabilities, for one individual, “elk-115”. It was obtained with the following command:

```
plotStates(m, animals="elk-115")
```

3.3.3 Model selection with AIC

The generic method AIC is available to compare `moveHMM` models. For example, we now fit a 3-state HMM to the data, and want to compare the AICs of the 2-state and 3-state models.

```
# initial parameters
mu0 <- c(0.1,0.5,3)
sigma0 <- c(0.05,0.5,1)
zeromass0 <- c(0.05,0,0)
stepPar0 <- c(mu0,sigma0,zeromass0)
angleMean0 <- c(pi,pi,0)
kappa0 <- c(1,1,1)
anglePar0 <- c(angleMean0,kappa0)

# fit the 3-state model
m3 <- fithMM(data=data,nbStates=3,stepPar0=stepPar0,
             anglePar0=anglePar0,formula=~dist_water)

# compare the AIC of the models
print(AIC(m,m3))
```

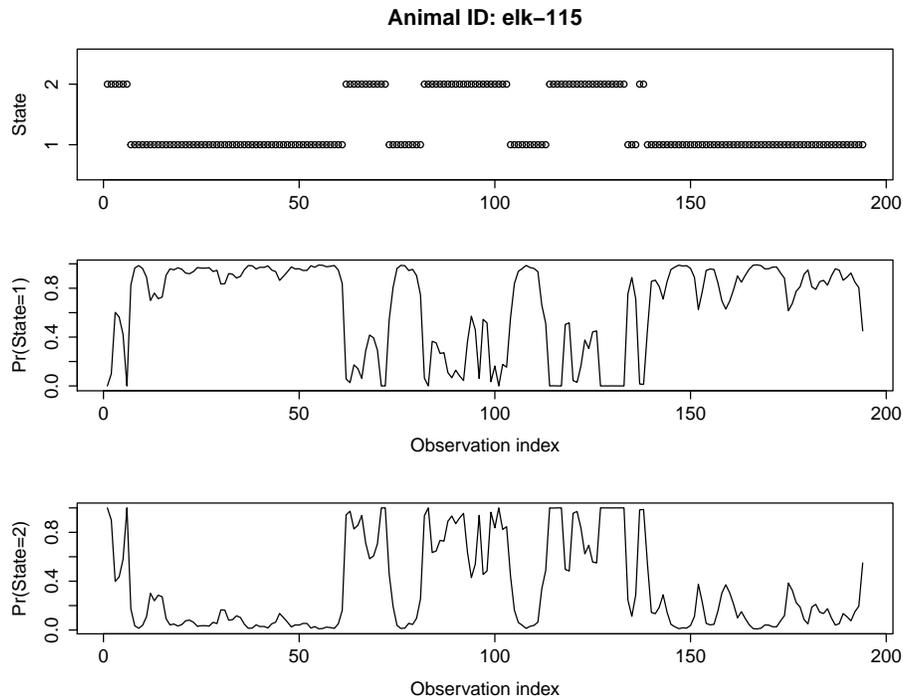


Figure 4: Decoded states sequence (top row), and state probabilities of observations (middle and bottom rows) for elk-115

This will output:

```

Model      AIC
1      m3 3673.066
2       m 3816.185

```

3.3.4 Model checking

The pseudo-residuals (a.k.a. quantile residuals) of the model can be computed with `pseudoRes`. These follow a standard normal distribution if the fitted model is the true data-generating process. In other words, a deviation from normality indicates a lack of fit. For more theoretical background on pseudo-residuals, see Zucchini and MacDonald (2009). The pseudo-residuals of the 2-state model fitted to the elk data are displayed in Figure 5. They can be computed and plotted with the following commands.

```

# compute the pseudo-residuals
pr <- pseudoRes(m)

# time series, qq-plots, and ACF of the pseudo-residuals

```

```
plotPR(m)
```

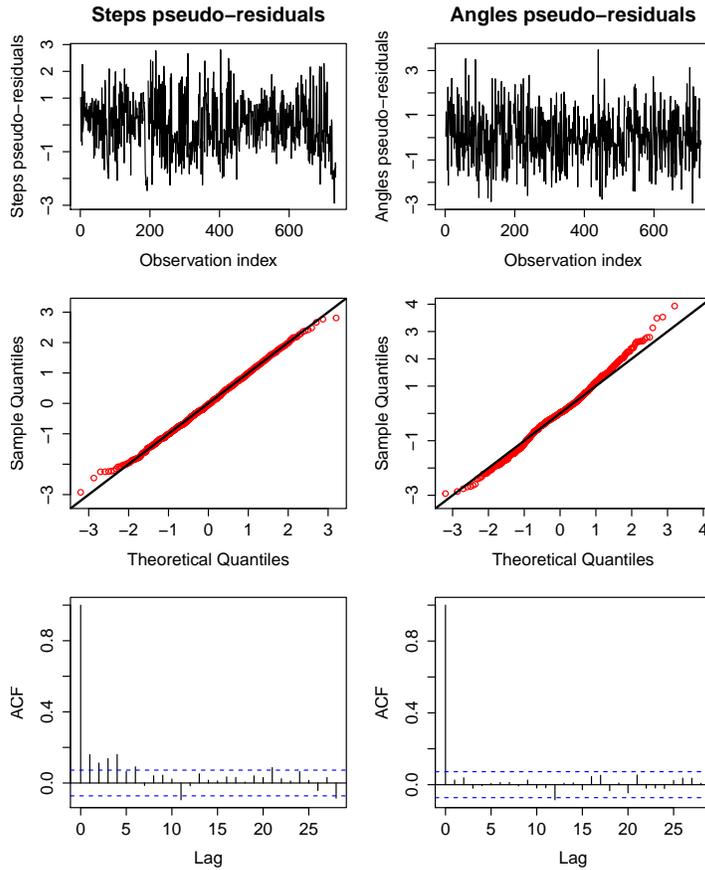


Figure 5: Time series, qq-plots, and autocorrelation functions of the pseudo-residuals of the 2-state model.

The following message is also output when calling one of the above commands on the elk model:

```
Note: Some angles are equal to pi, and the corresponding pseudo-residuals are not included.
```

The reason why those pseudo-residuals are not included is that they are equal to $+\infty$. Indeed, π is the upper bound of the definition set of the circular distributions that are included in the package. Thus, the pseudo-residual associated with an angle of π is the upper bound of the definition set of the normal distribution, i.e. $+\infty$.

4 Package features

In this section, we describe the global structure of the package, and then describe in more detail the main functions required to fit an HMM to movement data.

The package is articulated in terms of two S3 classes: `moveData` and `moveHMM`. The first extends the native R data frame, essentially gathering time series of the movement metrics of interest, namely the step lengths and turning angles, as well as the covariate values. A `moveHMM` object is a fitted model, which stores in particular the values of the MLE of the parameters.

To create a `moveData` object, the function `prepData` is called on the tracking data (track points coordinates). Then, the function `fitHMM` is called on the `moveData`, and returns a `moveHMM`.

Both classes can be used through their methods (e.g. `plot.moveData`, `AIC.moveHMM`), and a variety of other functions can be called on `moveHMM` objects. All functions are described in more detail in Section 4.2 and their use is explained on an example in Section 3. Figure 6 illustrates the links between the main components of the package.

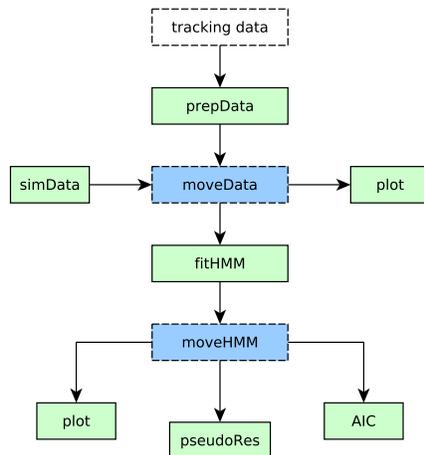


Figure 6: Structure of the main components of the package. The blue boxes are S3 classes, and the green boxes are functions. The arrows indicate input and output of data.

Note we will occasionally omit the class to which a method belongs, if the context makes it clear. Besides, as illustrated in the example in Section 3, it is not necessary to specify the class when calling the R function (e.g. calling `plot` on a `moveHMM` object automatically refers to `plot.moveHMM`).

4.1 Model options

4.1.1 Distributions

Here is the list of distributions included, with the names they have in the package.

- Step length: gamma (“gamma”), Weibull (“weibull”), exponential (“exp”) and log-normal (“lnorm”).
- Turning angle: von Mises (“vm”) and wrapped-Cauchy (“wrpcauchy”). It is also possible to specify `angleDist="none"`, if the angles are not modelled.

The parameters depend on the distribution used. The gamma distribution expects the mean and standard deviation, and all other distributions expect the same parameters as the corresponding R density function, i.e.

Distribution	Parameters	
gamma	mean > 0	standard deviation > 0
Weibull	shape > 0	scale > 0
log-normal	location $\in \mathbb{R}$	scale > 0
exponential	rate > 0	
von Mises	mean $\in (-\pi, \pi]$	concentration > 0
wrapped Cauchy	mean $\in (-\pi, \pi]$	concentration $\in (0, 1)$

For the gamma distribution, the link between the mean/standard deviation (expected by `fitHMM`) and shape/scale (expected by `dgamma`) is given by:

$$\begin{aligned} \text{shape} &= \frac{\text{mean}^2}{\text{sd}^2} \\ \text{scale} &= \frac{\text{sd}^2}{\text{mean}} \end{aligned}$$

4.1.2 Zero-inflation

If some steps are exactly equal to zero, then strictly positive distributions such as the gamma are inadequate. In such cases, zero-inflated distributions can be considered. A zero-inflated step length distribution simply assumes that there is a probability z of observing a 0 and a probability of $1 - z$ of observing a positive value distributed according to a standard positive distribution (e.g. a gamma). Within the package `moveHMM`, zero-inflation will automatically be included if there are zero steps. In that case the (state-dependent) values z will be estimated, with the remaining positive distribution, weighted by $1 - z$, specified as one of the available standard step length distributions, listed in Section 4.1.1.

4.1.3 Covariates

In practice it is often of interest to model the state transition probabilities as functions of time-varying covariates. This can be done by assuming the Markov chain to be time-varying, with transition probability matrix $\mathbf{\Gamma}^{(t)} = \left(\gamma_{ij}^{(t)} \right)$, linking the transition probabilities to the covariate(s) via the multinomial logit link. In the general case of N states,

$$\gamma_{ij}^{(t)} = \Pr(S_t = j | S_{t-1} = i) = \frac{\exp(\eta_{ij})}{\sum_{k=1}^N \exp(\eta_{ik})},$$

where

$$\eta_{ij} = \begin{cases} \beta_0^{(ij)} + \sum_{l=1}^p \beta_l^{(ij)} w_{lt} & \text{if } i \neq j, \\ 0 & \text{otherwise,} \end{cases}$$

for $i, j = 1, \dots, N$. Here $\{S_t\}$ is the state process, w_{lt} is the l -th covariate at time t and p is the number of covariates considered. The β parameters directly affect the off-diagonal elements in $\mathbf{\Gamma}^{(t)}$ — with an increase in the linear predictor η_{ij} resulting in an increase in $\gamma_{ij}^{(t)}$ — and hence also the diagonal entries due to the row constraints (with the entries in each row summing to one). Note in particular that we have to fix $\eta_{ii} = 0$ for all i since otherwise the model would be overparameterized (not identifiable).

Within `moveHMM`, the β coefficients for the off-diagonal transition probabilities are stored in an $(p + 1) \times (n \cdot (n - 1))$ matrix. For example, for a 3-state HMM with two covariates, the matrix `beta` is

$$\begin{pmatrix} \beta_0^{(12)} & \beta_0^{(13)} & \beta_0^{(21)} & \beta_0^{(23)} & \beta_0^{(31)} & \beta_0^{(32)} \\ \beta_1^{(12)} & \beta_1^{(13)} & \beta_1^{(21)} & \beta_1^{(23)} & \beta_1^{(31)} & \beta_1^{(32)} \\ \beta_2^{(12)} & \beta_2^{(13)} & \beta_2^{(21)} & \beta_2^{(23)} & \beta_2^{(31)} & \beta_2^{(32)} \end{pmatrix}$$

Here the first row corresponds to the intercept terms and the other two rows to the slope coefficients associated with the two covariates. There are as many columns as there are off-diagonal entries in the 3×3 transition probability matrix, and that matrix is filled row-wise (i.e. column 1 in `beta` is linked to $\gamma_{12}^{(t)}$, column 2 is linked to $\gamma_{13}^{(t)}$, column 3 is linked to $\gamma_{21}^{(t)}$, etc.).

In practice, many movement models involve only two states, in which case

the above equations boil down to

$$\begin{aligned} \mathbf{\Gamma}^{(t)} &= \begin{pmatrix} \frac{1}{1 + \exp\left(\beta_0^{(12)} + \sum_{l=1}^p \beta_l^{(12)} w_{lt}\right)} & \frac{\exp\left(\beta_0^{(12)} + \sum_{l=1}^p \beta_l^{(12)} w_{lt}\right)}{1 + \exp\left(\beta_0^{(12)} + \sum_{l=1}^p \beta_l^{(12)} w_{lt}\right)} \\ \frac{\exp\left(\beta_0^{(21)} + \sum_{l=1}^p \beta_l^{(21)} w_{lt}\right)}{1 + \exp\left(\beta_0^{(21)} + \sum_{l=1}^p \beta_l^{(21)} w_{lt}\right)} & \frac{1}{1 + \exp\left(\beta_0^{(21)} + \sum_{l=1}^p \beta_l^{(21)} w_{lt}\right)} \end{pmatrix} \\ &= \begin{pmatrix} 1 - \text{logit}^{-1}\left(\beta_0^{(12)} + \sum_{l=1}^p \beta_l^{(12)} w_{lt}\right) & \text{logit}^{-1}\left(\beta_0^{(12)} + \sum_{l=1}^p \beta_l^{(12)} w_{lt}\right) \\ \text{logit}^{-1}\left(\beta_0^{(21)} + \sum_{l=1}^p \beta_l^{(21)} w_{lt}\right) & 1 - \text{logit}^{-1}\left(\beta_0^{(21)} + \sum_{l=1}^p \beta_l^{(21)} w_{lt}\right) \end{pmatrix} \end{aligned}$$

The inverse logit link function is applied in order to map the real-valued predictor onto the interval $[0, 1]$ (with the above multinomial logit link representing a generalization of this approach to the case of $N > 2$ states). In the case of two states, the matrix β in `moveHMM` is structured as follows:

$$\begin{pmatrix} \beta_0^{(12)} & \beta_0^{(21)} \\ \beta_1^{(12)} & \beta_1^{(21)} \\ \vdots & \vdots \\ \beta_p^{(12)} & \beta_p^{(21)} \end{pmatrix}.$$

4.1.4 Stationarity

The function `fitHMM` includes the option of fitting a stationary model (using the option `stationary=TRUE`, with the default being `stationary=FALSE`). This is only possible if no covariates are incorporated into the model. (Otherwise the transition probabilities will be time-dependent, such that the Markov chain is non-homogeneous and in particular cannot be stationary.) When no covariates are considered and the option `stationary=TRUE` is selected, then the initial state distribution of the Markov chain will automatically be chosen as the stationary distribution (a.k.a. steady-state distribution) implied by the estimated transition probability matrix (as opposed to being estimated when `stationary=FALSE`). This stationary distribution is the vector δ that solves the equation $\delta = \delta\mathbf{\Gamma}$ subject to $\sum_{i=1}^N \delta_i = 1$. In practice, this solution almost always exists.

4.2 Main functions

4.2.1 prepData

Tracking data usually consists of time series of either easting-northing coordinates or longitude-latitude values. However, with the HMM approach the derived quantities step lengths and turning angles are modelled.

The function `prepData` computes the steps and angles from the coordinates. As input, this function takes an R data frame with columns “x” (either easting or longitude) and “y” (either northing or latitude). If the names of the coordinates columns are not “x” and “y”, then the argument `coordNames` should specify them. If several animals were observed, there should also be a column “ID” which identifies the animal being observed. If there is no “ID” column, all observations will be considered to be associated with a single animal. All additional columns are considered as covariates. In addition to the data frame, `prepData` takes an argument `type`, which can either be “LL” (longitude-latitude, the default) or “UTM”. The former indicates that the coordinates are longitude-latitude values, and the latter that they are easting-northing values.

To compute the step lengths, `prepData` calls the function `spDistN1` from the package `sp`. The step lengths are in the unit of the input if easting/northing are provided, and in kilometres if longitude/latitude are provided.

`prepData` outputs a data frame, with the same columns as the input, plus columns “step” and “angle”. This object is of the class `moveData`, and can be plotted using the generic function `plot`.

4.2.2 fitHMM

Using the function `fitHMM`, an HMM can be fitted to an object of class `moveData`, via numerical maximum likelihood. The list of the arguments of `fitHMM` is detailed in the documentation. The maximum likelihood estimation is carried out using the R function `nlm`.

This function outputs a list of information about the model. Most elements of that list are only meant to be used by the `moveHMM` functions (see Sections 4.2.3 and 4.2.4), but a few can be informative *per se*:

- `mle` contains the estimates of the parameters of the model;
- `mod` contains the output of the optimization function `nlm`, including `mod$minimum` (minimum of the negative log-likelihood) and `mod$hessian`, the Hessian of the negative log-likelihood function at its minimum.

4.2.3 Generic methods

Methods (i.e. class functions) are available for both `moveData` and `moveHMM` objects, to operate on them. For details on the options see the documentation, and for an example of their use, see Section 3.

- `plot.moveData` plots a few graphs to illustrate the data: a map of each animal’s track, time series of the steps and angles, histograms of the steps and angles.
- `summary.moveData` outputs some summary information about a `moveData` object: the number of animals, the number of observations for each animal, and quantiles of the covariate values.

- `plot.moveHMM` plots a few graphs to illustrate the fitted model: a map of each animal's track, colored by states, plots of the estimated density functions, plots of the transition probabilities as functions of the covariates.
- `print.moveHMM` prints the value of the maximum log-likelihood, and the maximum likelihood estimates of the parameters of the model.
- `AIC.moveHMM` returns the AIC of one or several fitted models.

4.2.4 Other operations on `moveHMM`

Other functions can be called on a `moveHMM` object, for further analysis.

- `CI` computes confidence intervals for the step length distribution parameters, for the turning angle distribution parameters, and for the regression coefficients of the transition probabilities.
- `pseudoRes` computes the pseudo-residuals of the model. These can be used to assess the goodness of fit. If the model is the true data-generating process, then the pseudo-residuals follow a standard normal distribution (Zucchini and MacDonald, 2009).
- `stateProbs` computes the probabilities of the underlying Markov chain being in the different states, at each observation, under the fitted model.
- `viterbi` computes the sequence of most probable states, under the fitted model, using the Viterbi algorithm (Zucchini and MacDonald, 2009).
- `plotStates` plots the most probable state sequence (as decoded with `viterbi`), and the state probabilities (as computed with `stateProbs`).
- `plotPR` plots time series, qq-plots, and the sample autocorrelation (ACF) functions of the pseudo-residuals of the fitted model (Zucchini and MacDonald, 2009). The qq-plots can be used to visually assess whether or not the pseudo-residuals are standard normally distributed: the points in the qq-plots will be close to the straight line if the model fits the data well, and any notable deviation from the straight line is an indication of a lack of fit. If the sample ACFs display a residual autocorrelation, then this is an indication that the model might not have captured all relevant correlation structure in the data.

4.2.5 `simData`

The function `simData` simulates movement data from an HMM, given its parameters. The returned object is of the class `moveData`, and can be visualized using `plot`, or fitted using `fitHMM`. The arguments of `simData` are detailed in the documentation.

References

- JONSEN, I.D., FLEMMING, J.M., MYERS, R.A. (2005), “Robust state-space modeling of animal movement data”, *Ecology*, 86 (11), 2874–2880.
- 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 (11), 2336–2342.
- 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 (3), 335–349.
- 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 (9), 2436–2445.
- PATTERSON T.A., BASSON M., BRAVINGTON M.V., GUNN J.S. (2009), “Classifying movement behaviour in relation to environmental conditions using hidden Markov models” *Journal of Animal Ecology*, 78 (6), 1113–1123.
- PATTERSON T.A., LANGROCK, R., BLACKWELL, P.G., THOMAS, L., PARTON, A. (2015), “Statistical modelling of animal movement: a myopic review and a discussion of good practice” *Preprint*.
- ZUCCHINI, W. AND MACDONALD, I.L. (2009). *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall (London).