

# Behavioral Change Point Analysis in R: The `bcpa` package

Eliezer Gurarie,  
Department of Statistics  
School of Environmental and Forest Sciences  
University of Washington, Seattle  
[eliezg@u.washington.edu](mailto:eliezg@u.washington.edu)

October 2013

## Contents

<b>1</b>	<b>Background</b>	<b>1</b>
<b>2</b>	<b>Summary of method</b>	<b>2</b>
<b>3</b>	<b>Detailed implementation</b>	<b>3</b>
3.1	Estimating auto-correlation / characteristic time-scale for irregular data . . . . .	3
3.2	Speeds and turning angles . . . . .	5
3.3	Obtaining a changepoint . . . . .	7
3.4	Applying the window sweep . . . . .	10
<b>4</b>	<b>Conclusions</b>	<b>15</b>
<b>5</b>	<b>Acknowledgments</b>	<b>15</b>
<b>6</b>	<b>References</b>	<b>16</b>

## 1 Background

The `bcpa` package is designed to streamline the implementation of the “behavioral change point analysis” (BCPA, Gurarie et al. 2009) for any time-stamped movement data, i.e. in which there are  $X$ ,  $Y$  and  $T$

coordinates representing spatial locations and time of observation.

The BCPA was developed in order to identify changes in animal behaviors that were obscured by visual inspection or standard techniques. Specific difficulties associated with movement data include the multi-dimensionality, auto- and cross-correlation, considerable internal structure (reflecting behavioral complexity), and data collection that can be error-ridden or be irregularly sampled. The irregular sampling is a particularly vexing problem for marine organism data, for which locations are usually transmitted only when the animal is at the surface, while most standard approaches to modeling movement (e.g. the Correlated Random Walk) are only meaningful for regularly sampled data.

The paper had attracted considerable interest from biologists and ecologists that collect animal movement data and are eager to identify structure in the behaviors. Unfortunately, the code that was originally posted as supplementary material was incomplete and poorly documented, making the method somewhat inaccessible to practitioners without strong statistical or programming backgrounds. After responding to (about) the hundredth email to share the code and offer some suggestions, and in light of a bevy of improvements to the code itself (e.g. greater speed by dropping some routines into C++, greater flexibility in visualizing and presenting the results, more usable “tuning knobs” and flexible syntax) it became clear that bundling the implementation in an R package would be the most efficient and tidy way to make it more accessible.

## 2 Summary of method

The BCPA uses a likelihood-based method for identifying significant changes in movement parameter values across long, complex datasets by sweeping an analysis window over the timeseries and identifying the most likely changepoints, while simultaneously testing which, if any, of the parameters might have changed at that changepoint.

Implementing the BCPA involves the following steps:

1. Pick a response time-series variable  $X$ . One fairly robust variable is the persistence velocity  $V_p = V \cos(\theta)$  where  $V$  is speed = displacement/time interval and  $\theta$  is turning angle. Another alternative is simply  $V$ .
2. Assume that the observations  $X(t)$  are a observations from a stationary continuous-time Gaussian process with mean  $\mu_i$ , standard deviation  $\sigma_i$  and time-scale of autocorrelation  $\tau_i$ , where  $i \in (1, 2, \dots, N)$  represents an *a priori* unknown number of behavioral states. *Note: the use of a continuous time scale  $\tau > 0$  is a change from the original model, which estimates a discrete autocorrelation  $0 < \rho < 1$ . The time-scale is more biologically meaningful, as it is estimated in units of time: the longer the time-scale the longer the “memory” of the movement.*
3. Obtain a likelihood for  $\mu_i$ ,  $\sigma_i$  and  $\rho_i$  within a given stationary state  $i$  (See Gurarie et al. 2009 for details).
4. Find the location within a window of observations which splits a subset of the data into two sets of the three parameters.
5. Within this window, use a version of BIC to determine which combination (if any) of the three parameters most parsimoniously describes the separation in the data. Often, the null model is returned. I say “modified” because the BIC is usually defined as:  $BIC = -2 \log(L) + k \log(n)$ , where  $L$  is the likelihood,  $k$  is the number of parameters, and  $n$  is the number of data points; however, the 2 is replace

with a constant  $K > 0$ . The smaller this value, the more conservative the analysis, i.e. the more likely it is to select a simpler or null model. This is one of several “tuning knobs” in the BCPA.

6. Sweep a window of fixed size across the time series and collect all the changepoints, the associated models, and the values of the estimated parameters according to the selected model on either side of the changepoint within the window. Note, the window size is another ”knob” - larger windows are more robust but more coarse, smaller windows are more sensitive but more likely to give slightly spurious results, compensated by adjusting the  $K$ .

These steps summarize the actual analysis. The output of the analysis can be summarized and presented in two ways:

1. Smooth BCPA: A “smooth” output is the one described in the original paper: is the average over all the estimated parameters, and the location of all the change points on the other. This gives interesting output - with, in particular, the opportunity to have parameters change both suddenly and gradually, and to visualize a phase plot of the behavioral shifts - both sudden and gradual.
2. Flat BCPA: A ”flat” output takes the result of the window sweep and finds the most frequently chosen change points, clustering those unique changepoints which are close to each other (within some interval  $dT_c$ ). The three parameters  $\mu$ ,  $\sigma$  and  $\tau$  are then estimated within each section, and the location of these “flat” changepoints is recorded. This output is directly comparable to the BPMM segmentation.

### 3 Detailed implementation

The complete analysis can now be performed in, essentially, two or three lines of code (see final sections of this vignette). Before leaping into them, I review the assumptions and fundamental pieces of the method, using the functions within `bcpa` to facilitate implementation.

#### 3.1 Estimating auto-correlation / characteristic time-scale for irregular data

The BCPA (as currently implemented) analyzes a one-dimensional, arbitrarily sampled autocorrelated Gaussian time-series  $X(t)$ , specified by three parameters, a mean  $\mu$ , a standard deviation  $\sigma$  and a characteristic time-scale  $\tau$  (or autocorrelation coefficient  $\rho$ ), such that:

$$\begin{aligned} EX(t) &= \mu \\ \text{Var}X(t) &= \sigma^2 \\ \text{Cor}X(t)X(t - \Delta t) &= \exp(-\Delta t/\tau) = \rho^{\Delta t}, \end{aligned}$$

The characteristic time-scale is an innovation over the original implementation, which estimated  $\rho$ . The interpretation of  $\rho$ , which ranges from 0 to 1, depends on the units of the time measurement - and can flirt with being uninformatively close to 0 or 1, whereas the time scale is a measurement of the temporal range of correlation in the movement, with somewhat more intuitive biological interpretation (Gurarie and Ovaskainen 2011).


These relationships are used to obtain a likelihood for observations  $X_i$  observed at times  $T_i$ , and the likelihood is maximized to estimate the characteristic time scale. In the `bcpa` package, this is done with the `GetRho` function, which is driven by the `GetL` function (encoded in C++). Examples are given below:

Loading the package

```
library(bcpa)
```


Simulating a gappy, Gaussian, time series

```
rho <- 0.8
x.full <- arima.sim(1000, model = list(ar = rho))
t.full <- 1:1000
keep <- sort(sample(1:1000, 200))
x <- x.full[keep]
t <- t.full[keep]
plot(t, x, type = "l")
```



Obtaining the likelihood function for different values of  $\rho$ .

```
rhos <- seq(0, .99, .01)
L <- rep(NA, length(rhos))
for(i in 1:length(rhos))
  L[i] <- GetL(x,t,rhos[i])
# plot likelihood profile
plot(rhos, L, type="l")
abline(v = rho, lty=2, lwd=2); abline(v = rhos[L == max(L)], lty=3, lwd=2)
legend("bottomleft", legend=c("true value", "MLE"), lty=2:3, lwd=2)
```



Using the `GetRho` function to estimate  $\rho$  or  $\tau$ :

```
GetRho(x, t, tau = FALSE)
```

```
## rho.hat      LL
##  0.8209 -357.6468
```

```
GetRho(x, t, tau = TRUE)
```

```
## rho.hat      LL
##  5.067 -357.647
```

**Future work:** It is straightforward to obtain confidence intervals around this estimate.

### 3.2 Speeds and turning angles

To apply the method to movement data of the form  $\mathbf{Z}_i, T_i$ , where  $\mathbf{Z}_i$  represents the location vector at time  $T_i$ , it is necessary to extract a one-dimensional time series that conforms to the assumption. Examples which generally conform to the assumptions are the “persistence” velocity  $V_p(t)$  and turning velocity  $V_t(t)$ :

$$V_p(T_i) = V(T_i) \cos(\Theta(T_i)) \quad (1)$$

$$V_t(T_i) = V(T_i) \sin(\Theta(T_i)) \quad (2)$$

where  $V(T_i) = \|\mathbf{Z}_i - \mathbf{Z}_{i-1}\| / (T_i - T_{i-1})$  is the scalar speed at time  $T_i$ .  $V_p$  captures the tendency and magnitude of a movement to persist in a given direction while  $V_t$  captures the tendency of movement to head


in a perpendicular direction in a given time interval. Thus, the primary descriptive features of movement, namely speed, directional persistence, and variability are captured in these variables. Alternatively, the log of step lengths can have a roughly Gaussian distribution.

For demonstration purposes, we include a simulated movement data set called `Simp`:

```
data(Simp)
head(Simp)

##   Time    X     Y
## 1 0.18 23.75 -9.064
## 2 0.22 26.75 -11.325
## 3 0.74 13.25  2.946
## 4 0.88 28.46 26.634
## 5 1.40 96.78 121.783
## 6 1.41 97.86 123.348

plot(Simp)
```



The `Simp` objects is of class “track”, which is simply a data frame with columns  $X$ ,  $Y$  and  $T$ , and the `bcpa`

package contains a plotting method for a track of this form with a green circle illustrating the start and a red rhombus indicating the end of the track.

To obtain the step length and turning angles, use the `GetVT` function, which decomposes the data into single step and all the relevant statistics:

```
Simp.VT <- GetVT(Simp)
head(Simp.VT)
```


##	Z.start	Z.end	S	Phi	Theta	T.start	T.end	T.mid
## 2	26.75- 11.33i	13.25+ 2.95i	19.64	2.3282	2.97445	0.22	0.74	0.480
## 3	13.25+ 2.95i	28.46+ 26.63i	28.15	1.0000	-1.32815	0.74	0.88	0.810
## 4	28.46+ 26.63i	96.78+121.78i	117.13	0.9481	-0.05195	0.88	1.40	1.140
## 5	96.78+121.78i	97.86+123.35i	1.90	0.9674	0.01930	1.40	1.41	1.405
## 6	97.86+123.35i	101.47+135.40i	12.59	1.2797	0.31228	1.41	1.48	1.445
## 7	101.47+135.40i	117.74+193.35i	60.19	1.2971	0.01739	1.48	2.00	1.740

##	dT	V	T.POSIX
## 2	0.52	37.77	0.480
## 3	0.14	201.08	0.810
## 4	0.52	225.26	1.140
## 5	0.01	189.97	1.405
## 6	0.07	179.80	1.445
## 7	0.52	115.75	1.740

The overall persistence of the movement and distribution of step lengths:

```
par(mfrow = c(1, 2))
hist(Simp.VT$V, breaks = 20, col = "grey")
hist(Simp.VT$Theta, breaks = 20, col = "grey")
```



### 3.3 Obtaining a changepoint

A single changepoint in a time-series where the parameters change at some unknown timepoints  $t^*$  is done by simply sweeping all possible breaks, and finding the most-likely changepoint according to the likelihood. This is illustrated below:

```
par(bty="l")
mu1 <- 5; mu2 <- 3
sigma1 <- 2; sigma2 <- 1
rho1 <- 0.5; rho2 <- 0.5

SimTS <- function(n, mu, rho, sigma)
{
  X.standard <- arima.sim(n, model=list(ar = rho))
  X.standard/sd(X.standard)*sigma + mu
}

# create time series with break at 500
t.full <- 1:1000
t.break <- 500
x.full <- c(SimTS(t.break, mu1, rho1, sigma1),
           SimTS(max(t.full)-t.break+1, mu2, rho2, sigma2))


# subsample 100 observations and estimate
keep <- sort(sample(1:length(x.full), 100))
x <- x.full[keep]
t <- t.full[keep]
(BB <- GetBestBreak(x,t, tau=FALSE))

##      bb.index      bb.time      mu1      s1      rho.hat      LL
##      51.0000     482.0000     5.1092     2.0183     0.5548    -104.1158
##           mu2          s2      rho.hat      LL ll.total.LL
##           3.0158      0.9065     0.7785    -58.0164    -162.1321
```

The estimates should be fairly good (note that we are choosing to estimate  $\rho$  rather than  $\tau$ ).

```
plot(t,x, type="l")
abline(v = 500, col=2, lwd=2, lty=2); abline(v = BB[2], col=2, lwd=2, lty=3)
legend("topright", legend=c("true break", "estimated break"), col=2, lwd=2, lty=2:3)
```





The likelihood is used to obtain a BIC value for all of the possible models. The possible models are numbered M0 to M7, corresponding to no significant changes (M0), only  $\mu$ ,  $\sigma$  or  $\rho$  changing (M1, M2, M3, respectively), both of  $\mu$  and  $\sigma$ ,  $\mu$  and  $\rho$  and  $\sigma$  and  $\rho$  changing (M4, M5, M6), and all three parameters changing (M7). The model are compared with the `GetModels` function:

```
GetModels(x, t, BB[1], tau = FALSE)
```

```
##      Model    LL  bic  mu1   s1  rho1  mu2   s2  rho2
## [1,]    0 -193.7 401.3 4.061 1.888 0.7661 4.061 1.8885 0.7661
## [2,]    1 -178.9 380.8 5.109 1.553 0.5540 2.969 1.5526 0.5540
## [3,]    2 -193.6 410.1 4.061 2.018 0.8683 4.061 0.8533 0.8683
## [4,]    3 -191.1 405.2 4.061 1.888 0.5548 4.061 1.8885 0.7533
## [5,]    4 -161.8 351.2 5.109 2.018 0.6433 2.969 0.8533 0.6433
## [6,]    5 -177.0 381.7 5.109 1.553 0.5548 2.969 1.5526 0.7533
## [7,]    6 -192.5 412.6 4.061 2.018 0.5548 4.061 0.8533 0.7533
## [8,]    7 -161.2 354.6 5.109 2.018 0.5548 2.969 0.8533 0.7533
```

The model selection should select model 4 ( $\mu$  and  $\sigma$  change) as having the lowest BIC.

This is with the default sensitivity parameter  $K = 2$ , which comes from the definition of the BIC =  $-Kn \log(L) + k \log(n)$  (see summary in section 2). If we lower this value, a simpler model is more likely to be selected:

```
GetModels(x, t, BB[1], tau = FALSE, K = 0.5)
```

```
##      Model    LL  bic  mu1   s1  rho1  mu2   s2  rho2
## [1,]    0 -193.7 110.7 4.061 1.888 0.7661 4.061 1.8885 0.7661
## [2,]    1 -178.9 112.5 5.109 1.553 0.5540 2.969 1.5526 0.5540
## [3,]    2 -193.6 119.8 4.061 2.018 0.8683 4.061 0.8533 0.8683
```

```
## [4,] 3 -191.1 118.6 4.061 1.888 0.5548 4.061 1.8885 0.7533
## [5,] 4 -161.8 108.5 5.109 2.018 0.6433 2.969 0.8533 0.6433
## [6,] 5 -177.0 116.1 5.109 1.553 0.5548 2.969 1.5526 0.7533
## [7,] 6 -192.5 123.9 4.061 2.018 0.5548 4.061 0.8533 0.7533
## [8,] 7 -161.2 112.8 5.109 2.018 0.5548 2.969 0.8533 0.7533
```

And if we increase it, more complex models are likely to be selected


```
GetModels(x, t, BB[1], tau = FALSE, K = 5)
```

```
##      Model    LL    bic  mu1    s1  rho1  mu2    s2  rho2
## [1,] 0 -193.7 982.4 4.061 1.888 0.7661 4.061 1.8885 0.7661
## [2,] 1 -178.9 917.4 5.109 1.553 0.5540 2.969 1.5526 0.5540
## [3,] 2 -193.6 990.8 4.061 2.018 0.8683 4.061 0.8533 0.8683
## [4,] 3 -191.1 978.5 4.061 1.888 0.5548 4.061 1.8885 0.7533
## [5,] 4 -161.8 836.5 5.109 2.018 0.6433 2.969 0.8533 0.6433
## [6,] 5 -177.0 912.7 5.109 1.553 0.5548 2.969 1.5526 0.7533
## [7,] 6 -192.5 990.1 4.061 2.018 0.5548 4.061 0.8533 0.7533
## [8,] 7 -161.2 838.2 5.109 2.018 0.5548 2.969 0.8533 0.7533
```

Below, we change only the autocorrelation parameter, which is quite a bit more difficult to detect by eye:

```
mu1 <- 0
mu2 <- 0
sigma1 <- 1
```

```
## Warning: NA/Inf replaced by maximum positive value
```



The model selection should select model 4 (only  $\rho$  changes):

```
GetModels(x, t, BB[1], tau = FALSE)
```

```
##      Model      LL      bic      mu1      s1      rho1      mu2      s2      rho2
## [1,]      0 -133.3 280.5  0.01470 0.9400 0.5022 0.0147 0.9400 0.5021875
## [2,]      1 -131.7 286.4 -0.07943 0.9364 0.4731 0.0857 0.9364 0.4730730
## [3,]      2 -131.3 285.6  0.01470 0.9061 0.5475 0.0147 0.9666 0.5474696
## [4,]      3 -122.1 267.2  0.01470 0.9400 0.8828 0.0147 0.9400 0.0005311
## [5,]      4 -131.2 289.9 -0.07943 0.9061 0.5308 0.0857 0.9666 0.5307844
## [6,]      5 -121.8 271.3 -0.07943 0.9364 0.8828 0.0857 0.9364 0.0005311
## [7,]      6 -122.0 271.6  0.01470 0.9061 0.8828 0.0147 0.9666 0.0005311
## [8,]      7 -121.7 275.7 -0.07943 0.9061 0.8828 0.0857 0.9666 0.0005311
```

### 3.4 Applying the window sweep

The main wrapper function for the complete analysis is `WindowSweep`. We select a window size and sensitivity parameter  $K$  and sweep analysis windows across the entire time series:

```
Simp.ws <- WindowSweep(Simp.VT, "V*cos(Theta)", windowsize = 50, progress = FALSE,
  K = 2)
```

Note that the second argument of the function is a character string within which any function of the columns of the VT table can be analyzed as a response.

The key portion of the output of this function is the “windowsweep” data frame, which contains the proposed break (last column), the parameters to the left and right of the break, and the selected model:

```
head(Simp.ws$ws)
```


```
##      Model      LL      bic      mu1      s1      rho1      mu2      s2      rho2 Break.bb.time
## 1      1 -288.6 596.8 110.5 142.6 0.6908 348.4 142.6 0.6908          7.625
## 2      1 -288.3 596.2 112.3 145.3 0.7268 329.1 145.3 0.7268          7.405
## 3      1 -288.0 595.6 114.5 146.0 0.7487 323.4 146.0 0.7487          7.405
## 4      1 -289.8 599.2 111.7 142.9 0.5501 331.5 142.9 0.5501          7.565
## 5      1 -285.2 590.1 117.9 111.6 0.3901 427.4 111.6 0.3901         10.130
## 6      1 -285.3 590.2 116.4 111.3 0.3851 426.5 111.3 0.3851         10.130
```

Note that in this example, the first 4 windows detect no changes in the parameter values (Model 0) so the values are the same to the left and to the right of each changepoint. (A minor note: the `rho1` and `rho2` columns correspond to  $\tau_1$  and  $\tau_2$  - i.e. the time-scales.)

The following functions plot the output of the “smooth” summary, i.e. the summary in which all the windows are averaged to obtain the “smooth” model. In these plots, the vertical lines represent the significant change points, the width of the lines is proportional to the number of time that change point was selected, the black

and red lines represent the mean and standard deviation estimate, and the colors reflect the autocorrelation time-scale (bluer colors have smaller autocorrelation time scales).


```
plot(Simp.ws, type = "smooth")
plot(Simp.ws, type = "smooth", threshold = 7)
```



The `threshold` parameter indicates how many of the windows that were swept over the data must have selected that changepoint for it to be considered significant. The changepoints selected (at that threshold) are almost exactly the ones that were encoded in the original simulation.

The “flat” analysis first selects changepoints that it deems significant by clustering neighboring changepoints, and then estimates a homogeneous behavior between those changepoints. Note that by increasing the clusterwidth to 3, many of the more minor changepoints are filtered away and the resulting profile is fairly uniform.

```
plot(Simp.ws, type = "flat")
plot(Simp.ws, type = "flat", clusterwidth = 3)
```



A summary of the flat changepoints can be obtained as follows:

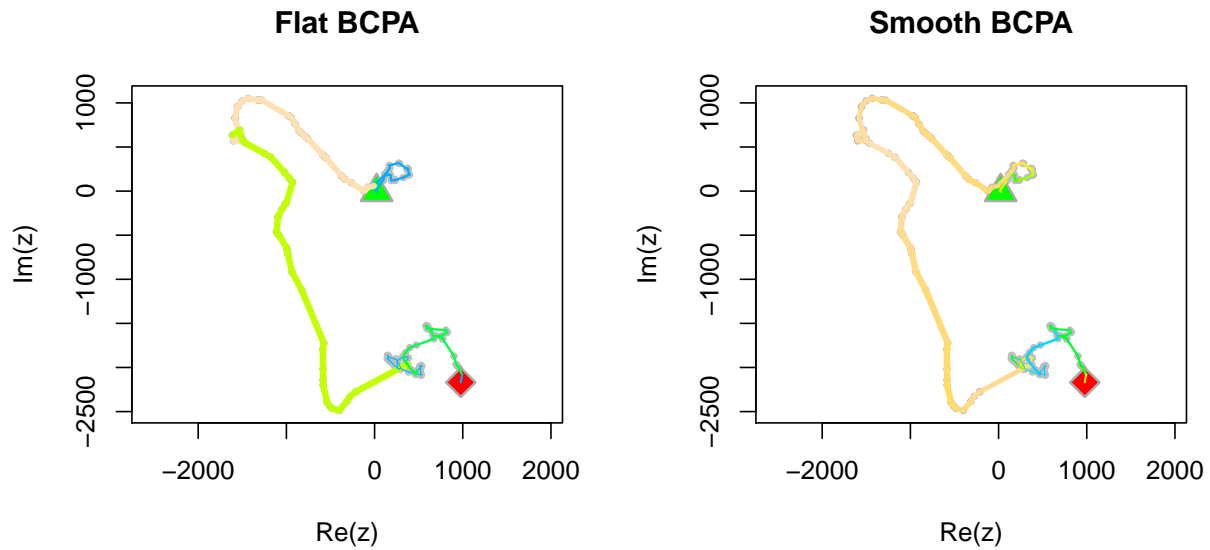
```
ChangePointSummary(Simp.ws, clusterwidth = 3)

## $breaks
##   middle size modelmode middle.POSIX
## 1  8.972  23         1         9.133
## 2 18.609  45         1        18.630
## 3 31.082  39         4        31.053
## 4 46.605  40         1        46.630
##
## $phases
##   t.cut mu.hat s.hat rho.hat t0 t1 interval
## 1 (-0.52,8.97] 118.24 107.7 0.3511 -0.520 8.972 9.492
## 2 (8.97,18.6] 306.38 180.6 1.2868 8.972 18.609 9.637
## 3 (18.6,31.1] 339.79 161.0 0.8083 18.609 31.082 12.473
## 4 (31.1,46.6] 59.51 59.9 0.3678 31.082 46.605 15.523
## 5 (46.6,59.4] 122.03 110.8 0.4408 46.605 59.410 12.805
```

This summary suggests five phases, with phases 2 and 3 consisting of a much higher velocity and longer time-scale movement than in the other phases.

The results of the BCPA can also be visualized with a so-called “path plot”:


```
PathPlot(Simp, Simp.ws, type = "flat", clusterwidth = 3, main = "Flat BCPA")
PathPlot(Simp, Simp.ws, type = "smooth", main = "Smooth BCPA")
```



The width of the line is proportional to the mean speed, while the colors correspond to the time-scales in the plots above. Both of these plots clearly separate the period of faster, more directed movement

An additional, potentially interesting visualization of the analysis is the “phase plot”, which illustrates how the three parameters change with respect to each other:


```
PhasePlot(Simp.ws, type = "smooth", clusterwidth = 3)
```



Finally, it is important to assess the assumptions of the BCPA using diagnostic plots. In particular, to assess

whether the standardized residuals of the final model are indeed distributed roughly as  $\mathcal{N}(0, 1)$  random variables

```
DiagPlot(Simp.ws)
```



This plot illustrates the qq-norm plot, the histogram and the auto-correlation function of the standardized data ( $Z_i = X_i - \hat{\mu}(T_i)/\hat{\sigma}(T_i)$ ). Overall the results seem to satisfy the assumptions of normality.

## 4 Conclusions

It is hoped that this package will facilitate analysis of complex behavioral movement data. However, it should be stressed that this is perhaps first and foremost an exploratory tool. Its strength is that it can distill complex information into some tabulated and visual summaries that outline underlying structures. It is also relatively fast - a long data set with tens of thousands of locations takes well under a minute to analyze on most machines.

That said, this tool merely *describes* and does not *explain* complex behavioral profiles. The BCPA can be used to propose some appropriate movement models or behavioral hypotheses for further testing. Alternatively, an explanatory or predictive analysis of behaviors with respect to covariates can be performed by extracting some biologically meaningful summary from the BCPA, and performing a post-hoc analysis to model the observed patterns with respect to covariates.

## 5 Acknowledgments

Gratitude is extended to F. Cagnacci, M. Panzacchi, B. van Moorter and colleagues at the Norwegian Institute of Nature Institute, who organized an animal movement analysis workshop and spearheaded a special issue of *J. Animal Ecology* (in progress), lighting the fuse to finally follow through on this package. C. Bracis and G. Passolt helped with the technical aspects of finalizing this package. Thanks, finally, to the many diverse colleagues and animal movement ecologists that tested earlier versions on actual data.

## 6 References

- Gurarie, E., R. Andrews and K. Laidre. 2009. A novel method for identifying behavioural changes in animal movement data. *Ecology Letters*. 12: 395-408.
- Gurarie, E., O. Ovaskainen. 2011. Characteristic spatial and temporal scales unify models of animal movement. *American Naturalist*. 178: 113-123.