

# Introduction to R for distance sampling

Workshop, 26 August 2018

*Centre for Research into Ecological and Environmental Modelling*

## 1.2 Using real data

In this practical, we use some real data to fit different detection function models and estimate density and abundance. The data were collected during a line transect survey of duck nests in Monte Vista National Wildlife Refuge, Colorado, USA. Transects were of length 128.75km and a distance out to 2.4m was searched. Distances of detected nests have been provided in a ‘csv’ text file in a basic format required by ‘Distance’. The columns in the the file are:

- Study.Area - name of the study region (Monte Vista NWR)
- Region.Label - identifier of regions (in this case there is only one region and it is set to ‘Default’)
- Sample.Label - line transect identifier
- distance - perpendicular distances (m).

The distances allow different key functions/adjustments to be fitted in the detection function model and, by including the transect lengths and area of the region, density and abundance can be estimated.

## Objectives of the practical

1. Import a text file
2. Understand the structure of a data frame
3. Fit different key functions/adjustments in the detection function model
4. Explore the model object i.e. `ddf` and `dht`
5. Create a data frame
6. Estimate density and abundance using `ds`.

## Importing the data

To let R know where to look for files (and also where to save the R workspace and ‘.Rmd’ files) we can set the ‘working directory’; from the menu along the top of the RStudio window click

on ‘Session > Set Working Directory > Choose Directory’ and select your chosen directory, for example ‘C:/workshop’.

Load the data into R with the following command:

```
nestts <- read.csv(file="ducknestts.csv", header=TRUE)
```

This command is made up of several components:

- `read.csv` is an in-built R function which reads in a data file of type ‘csv’,
- the function has two arguments specified; `file` specifies the name of the data file and `header=TRUE` specifies that the first row of the data file contains the names of the data columns. If the columns had not been named in the data file, then `header=FALSE` should be specified and the columns would have been called X1, X2, etc. automatically. They can be subsequently renamed.
- the `<-` symbol has assigned the data set to an object called `nestts`. Note that there is now an object called `nestts` listed on the ‘Environment’ tab.

To check that the data file has been read into R correctly, use the `head` and `tail` ‘functions’ to look at the top and bottom rows of the data, respectively. To look at the first few rows of `nestts` type the following command.

```
head(nestts)
```

The `head` function as used above displays the first 6 records of the named object. By default, `head` and `tail` display 6 rows of data but this can be changed by specifying a value for the function argument which controls this action. To display the last 2 records in the data, type the command:

```
tail(nestts, n=2)
```

In this function, `n` is the argument which controls the number of rows to display.

The object `nestts` is a dataframe object made up of rows and columns. Use the function `dim` to find out the dimensions of the data set (i.e. the total number of rows and columns):

```
dim(nestts)
```

How many rows and columns are there in `nestts`?

Another way to look at a data frame is to move to the ‘Environment tab’ and click on the rectangle (with the grid); this opens a new tab showing the data.

## Summarising the perpendicular distances

To access an individual column within a data frame we use the `$` symbol, for example to summarise the distances, then the following command is used:

```
summary(nests$distance)
```

Similarly to plot the histogram of distances, the command is:

```
hist(nests$distance, xlab="Distance (m)")
```

## Fitting different models

To use the `ds` function, first ensure that the `Distance` package (Miller 2017) has been loaded - see practical 1.1 for help on this.

The function `ds` requires a data frame to have a column called `distance` - since we have this in our `nests` data, we can simply specify the name of the data frame as follows:

```
nest.model1 <- ds(nests, key="hn", adjustment=NULL)
```

In this command, a half-normal key function is selected with no adjustment terms. Summarise the model:

```
summary(nest.model1)
```

Plot the detection function with the histogram having 12 bins:

```
plot(nest.model1, nc=12)
```

To fit different detection function shapes, we can change the `key` and `adjustment` arguments. For example to fit a half-normal key function with cosine adjustment terms, then use the command:

```
nest.model2 <- ds(nests, key="hn", adjustment="cos")
```

By default, AIC selection will be used to fit adjustment terms of up to order 5. Have any adjustment terms been selected?

To fit a hazard rate key function with hermite polynomial adjustment terms, then use the command:

```
nest.model3 <- ds(nests, key="hr", adjustment="herm")
```

```
summary(nest.model3)
```

Use the `help` command to find out what other key functions and adjustment terms are available.

## The `ds` object

The objects created with `ds` (e.g. `nest.model1`) are made up of several parts. We can list them using the `names` function as below:

```
names(nest.model1)
```

The detection function information is in the `ddf` part and the density and abundance estimates would be stored in the `dht` part. To access each part, then the `$` can be used (as with columns in a data frame). For example to see what information is stored in the `ddf` part, we can use the `names` function again:

```
names(nest.model1$ddf)
```

The `dht` part is essentially empty because we haven't included the necessary information to calculate density. We will do this later.

```
nest.model1$dht
```

## Goodness of fit

The usual tools for checking goodness of fit are available, for example:

```
ddf.gof(nest.model1$ddf, nc=12)
```

Note that the `ddf` part of the object is passed to the `ddf.gof` function. In this command, the number of bins has been set to 12 for the chi-squared goodness of fit test.

## Estimating density and abundance

So far, we have concentrated on the detection function but, with more information such as transect lengths and the area of the region, we can estimate density and abundance. This information can be included in the text file along with the distances but this would involve repeating some information and so `Distance` also uses a series of linked tables. Even though it is a bit more work, we will use the latter approach and illustrate a few more R commands along the way.

The tables are linked with columns that have reserved names: the column called `distance` is the reserved name for the perpendicular distances. The other columns (e.g. `Region.Label`, `Sample.Label`) are also reserved names and these are used to link to other tables.

`Sample.Label` contains the transect identifier, the values can be listed using the `unique` function:

```
unique(nests$Sample.Label)
```

We can create a data frame containing the names of the transects (sample) and search effort and also the name of the region. In this example, there was only one region and the effort was the same for all transects (128.75km) and so this information gets repeated.

```
 nests.sample <- data.frame(Region.Label=rep("Default", times=20),  
                           Sample.Label=1:20, Effort=rep(128.75, times=20))
```

Let's see what has this done:

```
head(nests.sample)
```

There is only one region and so the region table is trivial - for the purposes of this practical, we assume that the area of the region was 50km<sup>2</sup>.

```
nests.region <- data.frame(Region.Label="Default", Area=50)
```

Let's look at this data frame:

```
nests.region
```

The final table required is for the observations: some of the necessary information is in `nests` and so we can select these columns from `nests` and add a unique identifier for each detected nest as follows:

```
# Select necessary columns
nests.obs <- nests[,2:3]
# Create unique identifier
nests.obs$object <- 1:534
```

Check what this has done. Why does `object` go from 1 to 534?

```
head(nests.obs)
```

Now we have all the necessary information to estimate density and abundance, therefore, we can include these tables in the `ds` function:

```
nest.model4 <- ds(nests, key="hn", adjustment=NULL, region.table=nests.region,
                 sample.table=nests.sample, obs.table=nests.obs,
                 convert.units=0.001)
```

The `convert.units` argument ensures that the correct units are specified - in this example, distances are in metres, lengths in km and the area in km<sup>2</sup>.

Having run the function, the estimates are stored in the `dht` part of the object:

```
# Print out estimates
nest.model4$dht
```

## References

Buckland ST, Rexstad EA, Marques TA and Oedekoven CS (2015) Distance Sampling: Methods and Applications. Springer 277 pp. ISBN: 978-3-319-19218-5 (Print) 978-3-319-19219-2 (Online)

Miller DL (2017). Distance: Distance Sampling Detection Function and Abundance Estimation. R package version 0.9.7. <https://CRAN.R-project.org/package=Distance>

R Core Team (2017) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>