

# An Introduction to stockR

Scott D. Foster

CSIRO, Hobart, Tasmania, Australia

---

## Abstract

The stockR package is useful for finding stock structure within species (combined) populations. It does so by using a (potentially large) number of co-dominant genetic markers, for example SNPs. The package implements the methods described in Foster *et al.* (2018), which is a variation on the models and computational methods described elsewhere (see reference list in Foster *et al.* 2018). In this tutorial, we will go through the steps of:

1. Simulate (from a mixture model) SNP data for a population of animals with a known number of stocks (sup-populations); and
2. Analyse the simulated data (and retrieve the simulated populations).

We shall perform this twice, once in the situation where none of the individuals are assumed to share the same stock and once where some individuals do. These individuals may have been sampled from a breeding aggregation (as in the yellowfin tuna data in Foster *et al.* 2018). This introduction is, largely because there is (at present) limited functionality in this package. However, we have concentrated on making the routines available to be *good* (we hope that the interface is fine too).

*Keywords:* Stock Structure, Sub-Population, SNP Marker, Mixture Model, R.

---

## First Things First (setting up R for using stockR)

Before starting with this introduction to `stockR`, we need to make sure that everything is set up properly. Much of this will vary from computer to computer, but you must have a working version of R installed (preferably the latest one). This introduction was written and tested using R-3.4.0. It does not matter whether you prefer to use R through a development environment (such as RStudio) or through the command line – the results should be the same. So, start R and then:

```
install.packages("stockR")
```

You will be asked which repository you want to use. Just use one that is geographically close to where you are (or where your computer is). Next load the package.

```
library(stockR)
```

For illustration it is also good to fix the random number seed, so that this document is reproducible *exactly*.

```
set.seed( 747)  #a 747 is a big plane
```

Now, we are good to go with the rest of the introduction.

## Simulating a Co-Dominant Genetic Marker Data set

Let's simulate a data set, using the mixture model in [Foster \*et al.\* \(2018\)](#), where the mathematical/technical treatment is also given). We shall simulate data from  $K = 3$  stocks, without any sampling groups. It is important to note that the current implementation does not allow for specified stock sizes – the expectation of a stock size is the number of sampling groups divided by the number of stocks.

```
#number of individuals
N <- 100
#number of markers to measure on each of the individuals
M <- 5000
#number of sampling groups (same as number of individuals)
S <- N
#number of stocks
K <- 3
#simulate the data
myData <- sim.stock.data( nAnimal=N, nSNP=M, nSampleGrps=N, K=K)
```

This produces a data matrix of dimensions ( $M \times N$ ) so that the rows index the markers and the columns index the individuals. This data format is common to all functions within `stockR`. Although not highlighted here, it is possible to simulate datasets that have a limited number of informative markers – markers that are segregating between the stocks. Also, it is possible to alter the amount of separation between the stocks, but the default is demonstrated here (see [Foster \*et al.\* 2018](#), for details and parameterisation).

We have now simulated data, so let's have a look at the contents of the data. Basically, the object `myData` is a matrix with the number of columns equal to the number of individuals and the number of rows equal to the number of loci (in this case it is a  $5000 \times 100$  matrix). This object has a number of attributes though. These correspond to the stocks that each individual belongs to, as well as their sampling group. Ordering is the same as in the matrix of data too. Let's have a look at the data with code.

```
#the dimensions of the data
dim( myData)

[1] 5000  100

#the number of fish
ncol( myData)

[1] 100
```



function. For any real data set, there may have to be slightly more manipulation of the data. The key features of the input data is that: 1) it has as many rows as there are markers; 2) there are as many columns as individuals; 3) it is numeric (not a factor); 4) NAs are OK (they are ignored in grouping); and 5) the coding of data is by number of allele copies. The attributes of the simulated data (sample groups and stock membership) are not required to be present – although sample groups do need to be specified as their own function argument. Let's have a look at some code to find stocks:

```
#find the stocks in the data
stocks <- stockSTRUCTURE( myData, K=3) #K-EM estimation is default
#the (posterior) membership probability to each stock
stocks$postProb #shows high discrimination, which could be erroneous (bootstrap soon)
```

|    | Pop_1 | Pop_2 | Pop_3 |
|----|-------|-------|-------|
| 1  | 1     | 0     | 0     |
| 2  | 1     | 0     | 0     |
| 3  | 1     | 0     | 0     |
| 4  | 1     | 0     | 0     |
| 5  | 1     | 0     | 0     |
| 6  | 1     | 0     | 0     |
| 7  | 1     | 0     | 0     |
| 8  | 1     | 0     | 0     |
| 9  | 1     | 0     | 0     |
| 10 | 1     | 0     | 0     |
| 11 | 1     | 0     | 0     |
| 12 | 1     | 0     | 0     |
| 13 | 1     | 0     | 0     |
| 14 | 1     | 0     | 0     |
| 15 | 1     | 0     | 0     |
| 16 | 1     | 0     | 0     |
| 17 | 1     | 0     | 0     |
| 18 | 1     | 0     | 0     |
| 19 | 1     | 0     | 0     |
| 20 | 1     | 0     | 0     |
| 21 | 0     | 0     | 1     |
| 22 | 0     | 0     | 1     |
| 23 | 0     | 0     | 1     |
| 24 | 0     | 0     | 1     |
| 25 | 0     | 0     | 1     |
| 26 | 0     | 0     | 1     |
| 27 | 0     | 0     | 1     |
| 28 | 0     | 0     | 1     |
| 29 | 0     | 0     | 1     |
| 30 | 0     | 0     | 1     |
| 31 | 0     | 0     | 1     |
| 32 | 0     | 0     | 1     |
| 33 | 0     | 0     | 1     |

|    |   |   |   |
|----|---|---|---|
| 34 | 0 | 0 | 1 |
| 35 | 0 | 0 | 1 |
| 36 | 0 | 0 | 1 |
| 37 | 0 | 0 | 1 |
| 38 | 0 | 0 | 1 |
| 39 | 0 | 0 | 1 |
| 40 | 0 | 0 | 1 |
| 41 | 0 | 0 | 1 |
| 42 | 0 | 0 | 1 |
| 43 | 0 | 0 | 1 |
| 44 | 0 | 0 | 1 |
| 45 | 0 | 0 | 1 |
| 46 | 0 | 0 | 1 |
| 47 | 0 | 0 | 1 |
| 48 | 0 | 0 | 1 |
| 49 | 0 | 0 | 1 |
| 50 | 0 | 0 | 1 |
| 51 | 0 | 0 | 1 |
| 52 | 0 | 0 | 1 |
| 53 | 0 | 0 | 1 |
| 54 | 0 | 0 | 1 |
| 55 | 0 | 0 | 1 |
| 56 | 0 | 0 | 1 |
| 57 | 0 | 0 | 1 |
| 58 | 0 | 0 | 1 |
| 59 | 0 | 0 | 1 |
| 60 | 0 | 0 | 1 |
| 61 | 0 | 0 | 1 |
| 62 | 0 | 0 | 1 |
| 63 | 0 | 0 | 1 |
| 64 | 0 | 0 | 1 |
| 65 | 0 | 0 | 1 |
| 66 | 0 | 0 | 1 |
| 67 | 0 | 0 | 1 |
| 68 | 0 | 0 | 1 |
| 69 | 0 | 0 | 1 |
| 70 | 0 | 0 | 1 |
| 71 | 0 | 0 | 1 |
| 72 | 0 | 0 | 1 |
| 73 | 0 | 0 | 1 |
| 74 | 0 | 0 | 1 |
| 75 | 0 | 0 | 1 |
| 76 | 0 | 0 | 1 |
| 77 | 0 | 0 | 1 |
| 78 | 0 | 0 | 1 |
| 79 | 0 | 0 | 1 |

```

80      0      0      1
81      0      0      1
82      0      0      1
83      0      0      1
84      0      0      1
85      0      0      1
86      0      0      1
87      0      0      1
88      0      0      1
89      0      0      1
90      0      0      1
91      0      0      1
92      0      0      1
93      0      0      1
94      0      0      1
95      0      0      1
96      0      0      1
97      0      0      1
98      0      1      0
99      0      1      0
100     0      1      0

```

```

#hard classificaiton of individuals to stocks
stocks$hardClass <- apply( stocks$postProb, 1, which.max)
#Bootstrap to see about uncertainty in assignment
#for serious applications many more than B=25 will be needed.
#only 2 cores used to pass CRAN's arbitrary checks
stocks$boot <- stockBOOT( stocks, B=25, mc.cores=2)
#Assignment accounting for uncertainty.
#Could also look at other model quantities (e.g. allele frequencities)
stocks$uncertClass <- apply( stocks$boot$postProbs, 1:2,
                           quantile, probs=c(0.05,0.5,0.95))
#careful inspection of this object will give the lower and
#upper 90% confidence intervals and the median
#e.g. for the 99th individual
print( round( stocks$uncertClass[,99,], 3))

```

|     | Pop_1 | Pop_2 | Pop_3 |
|-----|-------|-------|-------|
| 5%  | 0     | 0     | 0.000 |
| 50% | 0     | 1     | 0.000 |
| 95% | 0     | 1     | 0.996 |

```

#showing that there is some uncertainty about which
#stock this indivudal belongs to.

```

There are a bunch of other output data in the `stocks` object. These either relate to the input

data, the arguments used, or to potentially useful quantities (such as mean allele frequencies within each stock).

For these data, the stocks were reliably found. Care should be exerted though as there is no effort given to trying to account for label-switching of output (so estimated stock 1 may actually correspond to simulated stock 3, for example). The last lines of code assess the uncertainty in the assignment of the stock assignment. This is done by bootstrap methods (see Foster *et al.* 2018), and in this case we find 90% confidence intervals as well as the median. Results for only the 99<sup>th</sup> individual are presented, which shows considerable variation – it is uncertain if this individual should be in the second or third stock.

The `stockSTRUCTURE` function has a number of different options for how the estimation of stock structure is performed. In most cases, the defaults should be reasonable but it is upon the user to make sure. Have a look at `?stockSTRUCTURE` and at Foster *et al.* (2018) for the different options and what they implement.

## Simulating and Finding Stocks with Sampling Groups

Code is now presented for simulating and analysing data with sampling groups, such as might arise from sampling a breeding aggregation. In this case, these samples are *a priori* known to come from the same stock and the analysis should assign all these fish to the same stock. See Foster *et al.* (2018) for more details.

```
#number of sampling groups (same as number of individuals)
S <- 15 #there are now 15 sampling groups (100 individuals will be
#distributed between them).
#number of stocks
K <- 3
#simulate the data
myData <- sim.stock.data( nAnimal=N, nSNP=M, nSampleGrps=S, K=K)
#find stocks in data
stocks <- stockSTRUCTURE( myData, K=3,
                           sample.grps=attributes(myData)$sampleGrps)
#once again, the stocks have been found (up to labelling)
#the simulated values
attributes( myData)$grps

[1] 1 1 1 2 2 2 2 2 2 2 3 3 3 3 3

#the data-derived values (as hard clustered)
apply( stocks$postProbs, 1, which.max)

samp1 samp10 samp11 samp12 samp13 samp14 samp15 samp2 samp3 samp4 samp5
  1     1     1     3     3     3     3     3     3     3     2
samp6 samp7 samp8 samp9
  2     2     2     2

#Bootstrap to see about uncertainty in assignment
```

```

#for serious applications many more than B=25 will be needed.
#only 2 cores used to pass CRAN's arbitrary checks
stocks$boot <- stockBOOT( stocks, B=25, mc.cores=2)
#Assignment accounting for uncertainty.
#Could also look at other model quantities (e.g. allele frequencies)
stocks$uncertClass <- apply( stocks$boot$postProbs, 1:2,
                             quantile, probs=c(0.05,0.5,0.95))
#careful inspection of this object will give the lower and
#upper 90% confidence intervals and the median
print( round( stocks$uncertClass, 3))

, , Pop_1

      samp1 samp10 samp11 samp12 samp13 samp14 samp15 samp2 samp3 samp4 samp5
5%      1      1      1      0      0      0      0      0      0      0      0
50%     1      1      1      0      0      0      0      0      0      0      0
95%     1      1      1      0      0      0      0      0      0      0      0
      samp6 samp7 samp8 samp9
5%      0      0      0      0
50%     0      0      0      0
95%     0      0      0      0

, , Pop_2

      samp1 samp10 samp11 samp12 samp13 samp14 samp15 samp2 samp3 samp4 samp5
5%      0      0      0      0      0      0      0      0      0      0      1
50%     0      0      0      0      0      0      0      0      0      0      1
95%     0      0      0      0      0      0      0      0      0      0      1
      samp6 samp7 samp8 samp9
5%      1      1      1      1
50%     1      1      1      1
95%     1      1      1      1

, , Pop_3

      samp1 samp10 samp11 samp12 samp13 samp14 samp15 samp2 samp3 samp4 samp5
5%      0      0      0      1      1      1      1      1      1      1      0
50%     0      0      0      1      1      1      1      1      1      1      0
95%     0      0      0      1      1      1      1      1      1      1      0
      samp6 samp7 samp8 samp9
5%      0      0      0      0
50%     0      0      0      0
95%     0      0      0      0

#Showing that there is almost no uncertainty about which
#stock individuals belong to.

```



## Summary

In this short introduction the two functions in `stockR` have been introduced. These functions simulate data and analyse it for stock structure. They are easy to use and the default estimation method is fast (much faster than many other methods and slower than only one or two). See [Foster \*et al.\* \(2018\)](#) for the comparison and for assessment of reliability.

If you have any queries regarding the package, then please contact the developer/maintainer. Your comments will be welcomed.

## Last Things Last

The only remaining thing to do is to tidy up our workspace. This is just removing all objects for this analysis from your workspace. I like to do this, in tutorial situations, but you may not. It is entirely up to you whether you clean or not.

```
#THIS IS SUGGESTIVE ONLY.  
#You may wish to tidy your workspace.  
#rm( list=ls())
```

## References

Foster SD, Feutry P, Grewe PM, Berry O, Hui FKC, Davies CR (2018). “Reliably discriminating stock structure with genetic markers: Mixture models with robust and fast computation.” *Molecular Ecology Resources*, **18**(6), 1310–1325. doi:10.1111/1755-0998.12920.

## Appendix

### Computational details

This vignette was created using the following R and add-on package versions

- R version 4.4.2 (2024-10-31), x86\_64-pc-linux-gnu
- Locale: LC\_CTYPE=en\_US.UTF-8, LC\_NUMERIC=C, LC\_TIME=en\_US.UTF-8, LC\_COLLATE=C, LC\_MONETARY=en\_US.UTF-8, LC\_MESSAGES=en\_US.UTF-8, LC\_PAPER=en\_US.UTF-8, LC\_NAME=C, LC\_ADDRESS=C, LC\_TELEPHONE=C, LC\_MEASUREMENT=en\_US.UTF-8, LC\_IDENTIFICATION=C
- Time zone: Etc/UTC
- TZcode source: `system (glibc)`
- Running under: Ubuntu 24.04.1 LTS
- Matrix products: default

- BLAS: `/usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3`
- LAPACK:  
`/usr/lib/x86_64-linux-gnu/openblas-pthread/libopenblas-p-r0.3.26.so ;`  
LAPACK version 3.12.0
- Base packages: base, datasets, grDevices, graphics, methods, stats, utils
- Other packages: knitr 1.48, stockR 1.0.76
- Loaded via a namespace (and not attached): RColorBrewer 1.1-3, Rcpp 1.0.13-1, buildtools 1.0.0, codetools 0.2-20, compiler 4.4.2, digest 0.6.37, evaluate 1.0.1, gtools 3.9.5, highr 0.11, maketools 1.3.1, parallel 4.4.2, sys 3.4.3, tools 4.4.2, xfun 0.49

**Affiliation:**

Scott D. Foster  
CSIRO  
Marine Laboratories  
GPObox 1538  
Hobart 7001  
Australia E-mail: [scott.foster@csiro.au](mailto:scott.foster@csiro.au)