

What could possibly go wrong? Troubleshooting spatially explicit capture–recapture models in **secr 3.1**

Murray Efford

2018-02-24

Contents

| | |
|---|----------|
| Introduction | 1 |
| When <code>secr.fit</code> fails | 1 |
| Bias limit exceeded | 1 |
| Initial log likelihood NA | 2 |
| Variance calculation failed | 3 |
| Log likelihood becomes NA or improbably large after a few evaluations | 4 |
| Possible maximization error: nlm code 3 | 4 |
| <code>secr.fit</code> requests more memory than is available | 5 |
| Covariate factor levels differ between sessions | 6 |
| Estimates from finite mixture models appear unstable | 6 |
| Faster is better | 6 |
| Mask tuning | 6 |
| Conditional likelihood | 6 |
| Parallel fitting of multiple models | 7 |
| Mashing | 7 |
| Reducing complexity (session- or group-specific models) | 8 |
| Collapsing occasions | 8 |
| Collapsing detectors | 9 |
| “multi” detector type instead of “proximity” | 9 |
| Some models are just slower than others | 9 |
| References | 9 |

Introduction

A lot can go wrong when fitting spatially explicit capture–recapture (SECR) models. This vignette assembles a list of known difficulties with the R package **secr 3.1**, with examples, and suggests some solutions. It largely supercedes the Troubleshooting help page in **secr 3.1**. Potential problems with density surface models are also discussed in `secr-densitysurfaces.pdf`.

Speed issues are also addressed.

When `secr.fit` fails

Bias limit exceeded

Suppose we omitted to specify the `buffer` argument for the first snowshoe hare model in `secr-tutorial.pdf`:

```
fit <- secr.fit (hareCH6, trace = FALSE)
```

Warning messages:

```
1: In secr.fit(hareCH6, trace = FALSE) : using default buffer width 100 m
2: In bufferbiascheck(output, buffer, biasLimit) :
predicted relative bias exceeds 0.01 with buffer = 100
```

The second warning message is clearly a consequence of the first: relying on the 100-m default buffer for an animal as mobile as the snowshoe hare is likely to cause mask truncation error. This is easily fixed by following advice on choosing the buffer width in `secr-tutorial.pdf` or `secr-habitatmasks.pdf`.

A check for mask truncation bias¹ is performed routinely by `secr.fit` after fitting a model that relies on the ‘buffer’ argument. It may be avoided by setting `biasLimit = NA` or providing a pre-computed habitat mask in the ‘mask’ argument.

Initial log likelihood NA

Maximization will fail completely if the likelihood cannot be evaluated at the starting values. This will be obvious with `trace = TRUE`. Otherwise, the first indication will be a premature end to fitting and a lot of NAs in the estimates.

As an example, try modelling the dataset `infraCH` (*Oligosoma infrapunctatum* skinks sampled with pitfall traps over two 3-occasion sessions labelled ‘6’ and ‘7’). Fitting fails with the default starting values:

```
fit <- secr.fit(infraCH, buffer = 25, trace = TRUE)
```

```
## Checking data
## Preparing detection design matrices
## Preparing density design matrix
## Finding initial parameter values...
## Initial values D = 241.50915, g0 = 0.18231, sigma = 2.39814
## Maximizing likelihood...
## Eval      Loglik      D      g0      sigma
##  1         NA    5.4869  -1.5007  0.8747
##  2         NA    5.4869  -1.5007  0.8747
##  3         NA    5.4869  -1.5007  0.8747
##  4         NA    5.4869  -1.5007  0.8747
##  5         NA    5.4869  -1.5007  0.8747
##  6         NA    5.4875  -1.5007  0.8747
##  7         NA    5.4869  -1.5006  0.8747
##  8         NA    5.4869  -1.5007  0.8748
##  9         NA    5.4880  -1.5007  0.8747
## 10         NA    5.4875  -1.5006  0.8747
## 11         NA    5.4875  -1.5007  0.8748
## 12         NA    5.4869  -1.5005  0.8747
## 13         NA    5.4869  -1.5006  0.8748
## 14         NA    5.4869  -1.5007  0.8749
## Completed in 19.39 seconds at 20:16:20 22 Feb 2018
```

This can be addressed by manually providing a value for sigma:

```
fit1 <- secr.fit(infraCH, buffer = 25, start = list(sigma = 5), trace = FALSE)
```

¹The positive bias that results when the buffer width for a detector array in continuous habitat is too small to encompass the centres of all animals potentially detected.

The problem seems to be due to a discrepancy between the two sessions (try `RPSV(infraCH, CC = TRUE)`). The default sigma is suitable for the first session and not the second, whereas a larger sigma suits both. Rather than manually providing the starting value we can direct `secr.fit` to use the second session for determining starting values:

```
fit2 <- secr.fit(infraCH, buffer = 25, details = list(autoini = 2), trace = FALSE)
```

Variance calculation failed

If warnings had not been suppressed in the preceding example we would have seen

Warning message:

```
In secr.fit(infraCH, buffer = 25, details = list(autoini = 2), trace = FALSE) :  
at least one variance calculation failed
```

Examination of the output would reveal missing values for SE, lcl and ucl in both the coefficients and predicted values for `g0`.

Failure to compute the variance can be a sign that the model is inherently non-identifiable or that the particular dataset is inadequate (e.g., Gimenez et al. 2004). However, here it is due to a weakness in the default algorithm. Call `secr.fit` with `method = "none"` to recompute the variance-covariance matrix using a more sophisticated approach:

```
fit2r <- secr.fit(infraCH, buffer = 25, start = fit2, method = "none")
```

```
## Checking data  
## Preparing detection design matrices  
## Preparing density design matrix  
## Computing Hessian with fdHess in nlme  
## 1 -2260.348 5.5119 -2.2331 1.4911  
## 2 -2260.355 5.5174 -2.2331 1.4911  
## 3 -2260.349 5.5119 -2.2309 1.4911  
## 4 -2260.350 5.5119 -2.2331 1.4926  
## 5 -2260.354 5.5064 -2.2331 1.4911  
## 6 -2260.349 5.5119 -2.2354 1.4911  
## 7 -2260.350 5.5119 -2.2331 1.4896  
## 8 -2260.358 5.5174 -2.2309 1.4911  
## 9 -2260.360 5.5174 -2.2331 1.4926  
## 10 -2260.353 5.5119 -2.2309 1.4926  
## 11 -2260.348 5.5119 -2.2331 1.4911  
## Completed in 17.08 seconds at 20:21:03 22 Feb 2018
```

```
predict(fit2r)
```

```
## $`session = 6`  
## link estimate SE.estimate lcl ucl  
## D log 247.62577739 13.966392881 221.73036479 276.545460  
## g0 logit 0.09681491 0.008500931 0.08138541 0.114804  
## sigma log 4.44199303 0.164288001 4.13149062 4.775831  
##  
## $`session = 7`  
## link estimate SE.estimate lcl ucl  
## D log 247.62577739 13.966392881 221.73036479 276.545460  
## g0 logit 0.09681491 0.008500931 0.08138541 0.114804  
## sigma log 4.44199303 0.164288001 4.13149062 4.775831
```

The trapping sessions were only 4 weeks apart in spring 1995. We can further investigate session differences by fitting a session-specific model. The fastest way to fit fully session-specific models is to fit each session separately; `lapply` here applies `secr.fit` separately to each component of `infraCH`:

```
fits3 <- lapply(infraCH, secr.fit, buffer = 25, trace = FALSE)
class(fits3) <- "seclist" # ensure secr will recognise the fitted models
predict(fits3)
```

```
## $`6`
##      link      estimate SE.estimate      lcl      ucl
## D      log 255.7579522 28.39104034 205.8868408 317.7091350
## g0     logit  0.1712699  0.02998793  0.1201776  0.2382027
## sigma  log   2.5190414  0.18159019  2.1875310  2.9007908
##
## $`7`
##      link      estimate SE.estimate      lcl      ucl
## D      log 278.0159896 18.70906015 243.69854627 317.1659892
## g0     logit  0.0989329  0.01007041  0.08087273  0.1204975
## sigma  log   4.9514263  0.21284561  4.55152402  5.3864645
```

Notice that there is no issue with starting values when the sessions are treated separately. The skinks appeared to enlarge their home ranges as the weather warmed; they may also have become more active overall². It is plausible that density did not change: the estimate increased slightly, but there is substantial overlap of confidence intervals.

Log likelihood becomes NA or improbably large after a few evaluations

The default maximization method (Newton-Raphson in function `nlm`) takes a large step away from the initial values at evaluation `np + 3` where `np` is the number of estimated coefficients. This often results in a very negative or NA log likelihood, from which the algorithm quickly returns to a reasonable part of the parameter space. However, for some problems it does not return and estimation fails³. Two solutions are suggested:

- change to the more robust Nelder-Mead maximization algorithm

```
secr.fit(CH, method = "Nelder-Mead", ...)
```

- vary the scaling of each parameter in `nlm` by passing the `typsize` (typical size) argument. The default is `typsize = rep(1, np)`. Suppose your model has four coefficients and it is the second one that appears to be behaving wildly:

```
secr.fit(CH, typsize = c(1, 0.1, 1, 1), ...)
```

In these examples `CH` is your capthist object and `...` indicates other arguments of `secr.fit`.

Possible maximization error: `nlm` code 3

The default algorithm for numerical maximization of the likelihood is `nlm` in the `stats` package. That uses a Newton-Raphson algorithm and numerical derivatives. It was chosen because it is significantly faster than the alternatives. However, it sometimes returns estimates with the ambiguous result code 3, which means that the optimization process terminated because “[the] last global step failed to locate a point lower than estimate. Either estimate is an approximate local minimum of the function or `stoptol` is too small.”

Here is an example:

²Home-range area increased about 4-fold; `g0` showed some compensatory decrease, but compensation was incomplete, implying increased total activity (treating `g0` as an estimate of λ_0 ; see Efford and Mowat 2014).

³You may also get weird messages about infinite densities; these messages are due to be removed in the next release.

```
fit3 <- secr.fit(infraCH, buffer = 25, model = list(g0~session, sigma~session),
               details = list(autoini = 2), trace = FALSE)
```

Warning message:

```
In secr.fit(infraCH, buffer = 25, model = list(g0 session, sigma :
possible maximization error: nlm returned code 3. See ?nlm
```

The results seem usually to be reliable even when this warning is given. If you are nervous, you can try a different algorithm in `secr.fit` – “Nelder-Mead” is recommended. We can derive starting values from the earlier fit:

```
fit3nm <- secr.fit(infraCH, buffer = 25, model = list(g0~session, sigma~session),
                 method = "Nelder-Mead", start = fit3, trace = FALSE)
```

There is no warning. Comparing the density estimates we see a trivial difference in the SE and confidence limits and none at all in the estimates themselves:

```
collate(fit3 = fit3, fit3nm = fit3nm)[1,,'D']
```

```
##      estimate SE.estimate      lcl      ucl
## fit3  271.9454    15.83433 242.6395 304.7909
## fit3nm 271.9454    15.80942 242.6830 304.7363
```

This suggests that only the variance-covariance estimates were in doubt, and it would have been much faster merely to check them with `method = "none"` as in the previous section.

`secr.fit` requests more memory than is available

In `secr` 3.1 the memory required by the external C code is at least $32 \times C \times M \times K$ bytes, where C is the number of distinct sets of values for the detection parameters (across all individuals, occasions, detectors and finite mixture classes), M is the number of points in the habitat mask and K is the number of detectors. Each distinct set of values appears as a row in a lookup table⁴ whose columns correspond to real parameters; a separate parameter index array (PIA) has entries that point to rows in the lookup table. Four arrays with dimension $C \times M \times K$ are pre-filled with, for example, the double-precision (8-byte) probability an animal in mask cell m is caught in detector k under parameter values c .

The number of distinct parameter sets C can become large when any real parameter (g_0 , λ_0 , σ) is modelled as a function of continuous covariates, because each unit (individual, detector, occasion) potentially has a unique level of the parameter. A rough calculation may be made of the maximum size of C for a given amount of available RAM. Given say 6GB of available RAM, $K = 200$ traps, and $M = 4000$ mask cells, C should be substantially less than $6e9 / 200 / 4000 / 32 \approx 234$. Allowance must be made for other memory allocations; this is simply the largest.

There is a different lookup table for each session; the limiting C is for the most complex session. The memory constraint concerns detection parameters only.

Most analyses can be re-configured to bring the memory request down to a reasonable number.

1. C may be reduced by replacing each continuous covariate with one using a small number of discrete levels (e.g. the mid-points of weight classes). For example, `weightclass <- 10 * trunc(weight/10) + 5` for midpoints of 10-g classes.
2. M can be reduced by building a habitat mask with an appropriate spacing (see `secr-habitatmasks.pdf`).
3. K might seem to be fixed by the design, but in extreme cases it may be appropriate to combine data from adjacent detectors (see Collapsing detectors).

⁴If you are really keen you can see this table by running `secr.fit` with `details = list(debug = 3)` and typing `Xrealparval` at the browser prompt (type Q to exit).

The `mesh` function (see Mashing) may be used to reduce the number of detectors when the design uses many identical and independent clusters. Otherwise, apply your ingenuity to simplify your model, e.g., by casting ‘groups’ as ‘sessions’. Memory is less often an issue on 64-bit systems (see also `?"Memory-limits"`).

Covariate factor levels differ between sessions

This is fairly explicit; `secr.fit` will stop if you include in a model any covariate whose factor levels vary among sessions, and `verify` will warn if it finds any covariate like this. This commonly occurs in multi-session datasets with ‘sex’ as an individual covariate when only males or only females are detected in one session. Use the function `shareFactorLevels` (new in `secr` 3.0) to force covariates to use the same superset of levels in all sessions.

Estimates from finite mixture models appear unstable

These models have known problems due to multimodality of the likelihood. See `secr-finitemixtures.pdf`.

Faster is better

There is nothing virtuous about waiting days for a model to fit if there is a faster alternative. Here are some things you can do.

Mask tuning

Consider carefully the necessary extent of your habitat mask and the acceptable cell size (`secr-habitatmasks.pdf` has detailed advice). If your detectors are clustered then your mask may have gaps between the clusters. Masks with more than 2000 points are generally excessive (and the default is about 4000!).

Conditional likelihood

The default in `secr.fit` is to maximize the full likelihood (i.e., to jointly fit both the state process and the observation process). If you do not need to model spatial, temporal or group-specific variation in density (the sole real parameter of the state model) then you can save time by first fitting only the observation process⁵. This is achieved by maximizing only the likelihood conditional on n , the number of detected individuals (Borchers and Efford 2008). Conditional likelihood maximization is selected in `secr.fit` by setting `CL = TRUE`.

```
fit <- secr.fit(hareCH6, buffer = 250, trace = FALSE) # default CL = FALSE
fitCL <- secr.fit(hareCH6, buffer = 250, CL = TRUE, trace = FALSE)
```

Fitting time is reduced by 32% because maximization is over two parameters (g_0 , σ) instead of three. The relative reduction will be less for more complex detection models, but still worthwhile.

Having selected a suitable observation model with `CL = TRUE` you can then either resort to a full-likelihood fit to estimate density, or compute a Horvitz-Thompson-like (HT) estimate in function `derived`. In models without individual covariates the HT estimate is $n/a(\hat{\theta})$ where n is the number of detected individuals, θ represents the parameters of the observation model, and a is the effective sampling area as a function of the estimated detection parameters.

⁵Selecting an observation model with `CL = TRUE` (and first focussing on detection parameters) is a good strategy even if you intend to model density later. It may occasionally be desirable to re-visit the selection if covariates can affect both density and detection parameters.

Compare

```
predict(fit) # CL = FALSE
```

```
##      link      estimate SE.estimate      lcl      ucl
## D      log  1.46598633 0.191312391  1.13636035  1.89122749
## g0     logit  0.06158393 0.009300473  0.04568546  0.08253654
## sigma   log 68.34577994 4.469318854 60.13241671 77.68098957
```

```
derived(fitCL) # CL = TRUE
```

```
##      estimate SE.estimate      lcl      ucl      CVn      CVa      CVD
## esa 46.385112          NA      NA      NA      NA      NA      NA
## D   1.465988    0.1905122  1.137563  1.889232  0.1212678  0.04671589  0.1299548
```

Estimated density is exactly the same, to 6 significant figures (1.46599). This is expected when n is Poisson; slight differences arise when n is binomial (because the number of animals N in the masked area is considered fixed rather than random). The estimated variance differs slightly - that from `derived` follows an unpublished and slightly *ad hoc* procedure (Borchers and Efford 2007).

Parallel fitting of multiple models

Your computer almost certainly has multiple cores, allowing computations to be run in parallel. Use the function `par.secr.fit` to take advantage of them when you are fitting multiple models. Parallel processing is available in some other functions (look for the ‘ncores’ argument), but it is of considerable benefit for only a few functions (see ‘?Parallel’).

See the brushtail possum example in `secr-densitysurfaces.pdf` for a good example of the use of `par.secr.fit`.

Mashing

Mashing is a very effective way of speeding up estimation when the design uses many replicate clusters of detectors, each with the same geometry, and far enough apart that animals are not detected on more than one. The approach for M clusters is to overlay all data as if from a single cluster; the estimated density will be M times the per-cluster estimate, and SE etc. will be inflated by the same factor. This relies on individuals being detected independently of each other, which is a standard assumption in any case. The present implementation assumes density is uniform.

We describe in general terms an actual example in which 18 separated clusters of 12 traps were operated on 6 occasions. Each cluster had the same geometry (two parallel rows of traps separated by 200 m along and between rows). Trap numbering was consistently up one row and down the other. The capthist object CH included detections of 150 animals in the 216 traps. To mash these data we first assign attributes for the cluster number (`clusterID`) and the sequence number of each trap within its cluster (`clustertrap`). The function `mash` then collapses the data as if all detections were made on one cluster. A mask based on this single notional cluster has many fewer points than a mask with the same spacing spanning all clusters.

```
clustertrap(traps(CH)) <- rep(1:12,18)
clusterID(traps(CH)) <- rep(1:18, each = 12)
mashedCH <- mash(CH)
mashedmask <- make.mask(traps(mashedCH), buffer = 900, spacing = 100, type = "trapbuffer")
fitmash <- secr.fit(mashedCH, mask = mashedmask)
```

The model for the mashed data fitted in 4% of the time required for the original. The mashed estimate of density shrank by 2% in this case, which is probably due to slight variation among clusters in the actual spacing of traps (one cluster was arbitrarily chosen to represent all clusters). Mashing prevents the inclusion

of cluster-specific detail in the model (such as discontinuous habitat near the traps). For further details see `?mash`.

Reducing complexity (session- or group-specific models)

Simultaneous estimation of many parameters can be painfully slow, but it can be completely avoided. If your model is fully session- or group-specific then it is much faster to analyse each group separately. For sessions this can be done simply with `lapply` above and in `secr-multisession.pdf`). For groups you may need to construct new `capthist` objects using `subset` to extract groups corresponding to the levels of one or more individual covariates.

Collapsing occasions

If there is no temporal aspect to the model you want to fit (such as a behavioural response) then it is attractive to collapse all sampling occasions to one with `reduce`. This can usually be done without loss of data by choosing 'outputdetector' carefully. For example,

```
FTHL.fit <- secr.fit (hornedlizardCH, buffer = 80, trace = FALSE)
```

```
## Warning in secr.fit(hornedlizardCH, buffer = 80, trace = FALSE): using
## default starting values
```

```
collapsedCH <- reduce(hornedlizardCH, by = "ALL", outputdetector = "polygon")
usage(traps(collapsedCH))
```

```
##      [,1]
## [1,]  14
```

```
FTHL.fit14 <- secr.fit (collapsedCH, buffer = 80, trace = FALSE, binomN = 1)
```

```
## Warning in secr.fit(collapsedCH, buffer = 80, trace = FALSE, binomN = 1):
## using default starting values
```

The collapsed `capthist` object receives a `usage` attribute equal to the number of collapsed occasions (14). Unfortunately, in this instance the collapsed model takes longer to fit, on account of having to compute a Binomial probability with size = 14 rather than a Bernoulli probability. A Poisson model (not shown) is slightly faster than the original, but the estimates differ by about 5%.

```
collate(FTHL.fit = FTHL.fit, FTHL.fit14 = FTHL.fit14)[1,, 'D']
```

```
##           estimate SE.estimate      lcl      ucl
## FTHL.fit   8.013068   1.061701  6.187400 10.37742
## FTHL.fit14 8.013062   1.061700  6.187395 10.37741
```

```
FTHL.fit$proctime
```

```
## elapsed
##   64.46
```

```
FTHL.fit14$proctime
```

```
## elapsed
##   94.78
```


Collapsing detectors

The theoretical (but often unrealised) benefit from collapsing occasions has a spatial analogue: if there are many detectors and they are closely spaced relative to animal movements σ then nearby detectors may be aggregated into new notional detectors located at the centroid. The `reduce.traps` method has an argument ‘span’ explained as follows in the help –

If `span` is specified a clustering of detector sites will be performed with `hclust` and detectors will be assigned to groups with `cutree`. The default algorithm in `hclust` is complete linkage, which tends to yield compact, circular clusters; each will have diameter less than or equal to `span`.

“multi” detector type instead of “proximity”

The type of the detectors is usually determined by the sampling reality. For example, if individuals can physically be detected at several sites on one occasion then the “proximity” detector type is preferred over “multi”. However, if data are very sparse, so that individuals in practice are almost never observed at multiple sites on one occasion, then the detectors may as well be of type “multi”, in the sense that there is no observable difference between the two types of detection process. This is significant because “multi” models fit much more quickly than “proximity” models.

Some models are just slower than others

Detector covariates are a particular problem. Models with learned responses take slightly longer to fit.

References

- Borchers, D. L. and Efford, M. G. (2007) Supplements to Biometrics paper. Available online at <http://www.otago.ac.nz/density>.
- Borchers, D. L. and Efford, M. G. (2008) Spatially explicit maximum likelihood methods for capture–recapture studies. *Biometrics* **64**, 377–385.
- Gimenez, O., Viallefont, A., Catchpole, E. A., Choquet, R. and Morgan, B. J. T. (2004) Methods for investigating parameter redundancy. *Animal Biodiversity and Conservation* **27**, 561–572.
- Otis, D. L., Burnham, K. P., White, G. C. and Anderson, D. R. (1978) Statistical inference from capture data on closed animal populations. *Wildlife Monographs* No. **62**.
- Royle, J. A. and Young, K. V. (2008) A hierarchical model for spatial capture–recapture data. *Ecology* **89**, 2281–2289.