

Multi-session models in **secr** 3.1

Murray Efford

2018-02-23

Contents

Input of session-specific data	1
Detections	1
Session-specific detector layouts	2
Manipulating multi-session data	3
Fitting multi-session models with secr.fit	3
Session-specific habitat masks	3
Session models	4
Session-stratified estimates	4
Trend across sessions	4
Session covariates	6
More on simulating multi-session data	6
Simulating multi-session populations	6
Multi-session sampling	6
Problems specific to multi-session data	6
Failure of autoini	6
Covariates with incompatible factor levels	7
Faster fitting of fully session-specific models	7
Caveats	8
Independence is a strong assumption	8
Parameters are assumed constant by default	9
Other notes	9
References	9

A ‘session’ in **secr** is a block of sampling that may be treated as independent from all other sessions. For example, sessions may correspond to trapping grids that are far enough apart that they sample non-overlapping sets of animals. Multi-session data and models combine data from several sessions. Sometimes this is merely a convenience, but it also enables the fitting of models with parameter values that apply across sessions – data are then effectively pooled with respect to those parameters.

Dealing with multiple sessions adds another layer of complexity, and raises some entirely new issues. This document tries for a coherent view of multi-session analyses, covering material that is otherwise scattered.

Input of session-specific data

A multi-session capthist object is essentially an R list of single-session capthist objects. We assume the functions `read.capthist` or `make.capthist` will be used for data input (simulated data are considered separately later on).

Detections

Entering session-specific detections is simple because all detection data are placed in one file or dataframe. Each session uses a character-valued code (the session identifier) in the first column. For demonstration let’s assume you have a file ‘msCHcapt.txt’ with data for 3 sessions, each sampled on 4 occasions.

```
## # Session ID Occasion Detector sex
## 1 28 1 A1 m
## 1 23 4 A1 f
## 1 6 3 A4 m
## 1 19 4 A4 f
## .
## .
## 2 12 4 A1 m
## 2 17 3 A5 f
## 2 3 1 A6 m
## 2 3 2 A6 m
## .
## .
## 3 28 1 A1 f
## 3 4 3 A2 f
## 3 6 3 A6 m
## 3 4 1 B1 f
## .
## .
```

(clipped lines are indicated by ‘.’).

Given a trap layout file ‘msCHtrap.txt’ with the coordinates of the detector sites (A1, A2 etc.), the following call of `read.caphist` will construct a single-session caphist object for each unique code value and combine these in a multi-session caphist:

```
msCH <- read.caphist('msCHcapt.txt', 'msCHtrap.txt', covnames = 'sex')
```

```
## No errors found :-)
```

Use the `summary` method or `str(msCH)` to examine `msCH`. Session-by-session output from `summary` can be excessive; the ‘terse’ option gives a more compact summary across sessions (columns).

```
summary(msCH, terse = TRUE)
```

```
##           1  2  3
## Occasions  4  4  4
## Detections 54 36 49
## Animals    31 20 29
## Detectors  64 64 64
```

Sessions are ordered in `msCH` according to their identifiers (‘1’ before ‘2’, ‘Albert’ before ‘Beatrice’ etc.). The order becomes important for matching with session-specific trap layouts and masks, as we see later. The vector of session names (identifiers) may be retrieved with `session(msCH)` or `names(msCH)`.

It is possible for there to be no detections in some sessions (but not all!). To create a session with no detections, include a dummy row with the value of the `noncapt` argument as the animal identifier; the default `noncapt` value is ‘NONE’. The dummy row should have some nonsense value (e.g. 0) in each of the other fields (occasion, trapID etc.).

Session-specific detector layouts

All sessions may share the same detector layout. Then the ‘trapfile’ argument of `read.caphist` is a single name, as in the example above. The trap layout is repeated as an attribute of each component (single-session) caphist.

Alternatively, each session may have its own detector layout. Unlike the detection data, each session-specific layout is specified in a separate input file or traps object. For `read.caphist` the ‘trapfile’ argument is then

a vector of file names, one for each session. For `make.caphist`, the ‘traps’ argument may be a list of traps objects, one per session. The first trap layout is used for the first session, the second for the second session, etc.

Manipulating multi-session data

The standard extraction and manipulation functions of `secr` (`summary`, `verify`, `covariates`, `subset`, `reduce` etc.) mostly allow for multi-session input, applying the manipulation to each component session in turn.

Plotting a multi-session caphist object (e.g., `plot(msCH)`) will create one new plot for each session unless you specify `add = TRUE`.

Methods that extract attributes from multi-session caphist object will generally return a list in which each component is the result from one session. Thus for the ovenbird mistnetting data `traps(ovenCH)` extracts a list of 5 traps objects, one for each annual session 2005–2009.

The `subset` method for caphist objects has a ‘sessions’ argument for selecting particular session(s) of a multi-session object.

Function	Purpose	Input	Output
<code>join</code>	collapse sessions	multi-session caphist	single-session caphist
<code>MS.caphist</code>	build multi-session caphist	several single-session caphist	multi-session caphist
<code>split</code>	subdivide caphist	single-session caphist	multi-session caphist

The `split` method for caphist objects (`?split.caphist`) may be used to break a single-session caphist object into a multi-session object, segregating detections by some attribute of the individuals, or by occasion or detector groupings.

Fitting multi-session models with `secr.fit`

Given multi-session caphist input, `secr.fit` automatically fits a multi-session model by maximising the product of session-specific likelihoods (Efford et al. 2009). For fitting a model separately to each session see the later section on Faster fitting. . . .

Session-specific habitat masks

The default mechanism for constructing a habitat mask in `secr.fit` is to buffer around the trap layout. This extends to multi-session data; buffering is applied to each trap layout in turn.

Override the default buffering mechanism by specifying the ‘mask’ argument of `secr.fit`. This is necessary if you want to –

1. reduce or increase mask spacing (pixel size; default 1/64 x-range)
2. clip the mask to exclude non-habitat
3. include mask covariates (predictors of local density)
4. define non-Euclidean distances (`secr-noneuclidean.pdf`)
5. specify a rectangular mask (type = “traprect” vs type = “trapbuffer”)

For any of these you are likely to use the `make.mask` function (the manual alternative is usually too painful to contemplate). If `make.mask` is provided with a list of traps objects as its ‘traps’ argument then the result is a list of mask objects - effectively, a multi-session mask.


If `addCovariates` receives a list of masks and a single spatial data source¹ then it will add the requested covariate(s) to each mask and return a new list of masks. The single spatial data source is expected to span all the regions; mask points that are not covered receive NA covariate values. As an alternative to a single spatial data source, the `spatialdata` argument may be a list of spatial data sources, one per mask, in the order of the sessions in the corresponding `capthist` object.

To eliminate any doubt about the matching of session-specific masks to session-specific detector arrays it is always worth plotting one over the other. We don't have an interesting example, but

```

masks <- make.mask(traps(msCH), buffer = 80, nx = 32, type = 'trapbuffer')
par (mfrow = c(1,3), mar = c(1,1,3,1))
for (sess in 1:length(msCH)) {
  plot(masks[[sess]])
  plot(traps(msCH)[[sess]], add = TRUE)
  mtext(side=3, paste('session', sess))
}

```



Session models

The default in `secr.fit` is to treat all parameters as constant across sessions. For detection functions parameterized in terms of cumulative hazard (e.g., 'HHN' or 'HEX') this is equivalent to `model = list(D ~ 1, lambda0 ~ 1, sigma ~ 1)`. Two automatic predictors are provided specifically for multi-session models: 'session' and 'Session'.

Session-stratified estimates

A model with lowercase 'session' fits a distinct value of the parameter (D, g0, lambda0, sigma) for each level of `factor(session(msCH))`.

Trend across sessions

A model with initial uppercase 'Session' fits a *trend* across sessions using the session number as the predictor. A trend model for density may be interesting if the sessions fall in some natural sequence, such as a series of annual samples (as in the ovenbird dataset `ovenCH`). The fitted trend is linear on the link scale; using the default link function for density ('log') this corresponds to exponential growth or decline if samples are equally spaced in time.

The pre-fitted model `ovenbird.model.D` provides an example. The coefficient 'D.Session' is the rate of change in $\log(D)$:

¹See the help for `addCovariates` regarding possible spatial data sources.

```
coef(ovenbird.model.D)
```

```
##           beta    SE.beta      lcl      ucl
## D           0.03169708 0.19146033 -0.3435583  0.40695243
## D.Session -0.06385425 0.07015155 -0.2013488  0.07364026
## g0          -3.56192706 0.15062257 -3.8571419 -3.26671225
## sigma       4.36411019 0.08114870  4.2050617  4.52315872
```

The overall finite rate of increase (equivalent to Pradel's lambda) is given by

```
beta <- coef(ovenbird.model.D)['D.Session','beta']
sebeta <- coef(ovenbird.model.D)['D.Session','SE.beta']
exp(beta)
```

```
## [1] 0.9381417
```

Confidence intervals may also be back-transformed with `exp`. To back-transform the SE use the delta-method approximation $\exp(\beta) * \sqrt{\exp(\text{sebeta}^2)-1} = 0.0658931$.

This is fine for a single overall lambda. However, if you are interested in successive estimates (session 1 to session 2, session 2 to session 3 etc.) the solution is slightly more complicated and uses a new feature of `secr` 3.1.1. The trick is to use 'backward difference' coding of the levels of the factor `session`, specified with the new details argument 'contrasts'. This coding is provided by the function `contr.sdif` in the **MASS** package (e.g., Venables and Ripley 1999 Section 6.2).

```
fit <- secr.fit(ovenCH, model = D~session, buffer = 300, trace = FALSE,
               details = list(contrasts = list(session = MASS::contr.sdif)))
coef(fit)
```

```
##           beta    SE.beta      lcl      ucl
## D          -0.1060440 0.13808471 -0.3766851  0.1645970
## D.session2-1  0.0458424 0.30898503 -0.5597571  0.6514419
## D.session3-2  0.1670526 0.28968649 -0.4007225  0.7348277
## D.session4-3 -0.3136594 0.30181438 -0.9052047  0.2778859
## D.session5-4 -0.1718478 0.33930017 -0.8368639  0.4931684
## g0          -3.5609571 0.15067158 -3.8562680 -3.2656463
## sigma       4.3639971 0.08109701  4.2050498  4.5229443
```

The beauty of this method is that it estimates the session-specific $\log(\lambda)$ directly - just back-transform as before:

```
beta <- coef(fit)[2:5, 'beta']
sebeta <- coef(fit)[2:5, 'SE.beta']
exp(beta) # lambda-hat
```

```
## [1] 1.0469094 1.1818164 0.7307679 0.8421074
```

```
exp(beta) * sqrt(exp(sebeta^2)-1) # SE(lambda-hat)
```

```
## [1] 0.3313559 0.3496659 0.2256756 0.2941514
```

The lambda estimates match the density ratios computed directly (a warning is suppressed):

```
D <- collate(fit, realnames='D')[,1,1,]
D[2:5]/D[1:4]
```

```
## session=2006 session=2007 session=2008 session=2009
## 1.0469094 1.1818164 0.7307679 0.8421074
```

The ovenbird population appeared to increase in density for two years and then decline for two years, but the effects are far from significant.

Session covariates

Other variation among sessions may be modelled with session-specific covariates. These are provided to `secr.fit` on-the-fly in the argument ‘`sessioncov`’ (they cannot be embedded in the `capthist` object like detector or individual covariates). The value for ‘`sessioncov`’ should be a dataframe with one row per session. Each column is a potential predictor in a model formula; other columns are ignored.

Session covariates are extremely flexible. The linear trend of the ‘`Session`’ predictor may be emulated by defining a covariate `sessnum = 0:(R-1)` where `R` is the number of sessions. Sessions of different types may be distinguished by a factor-valued covariate. Supposing for the ovenbird dataset we wished to distinguish years 2005 and 2006 from 2007, 2008 and 2009, we could use `earlylate = factor(c('early', 'early', 'late', 'late', 'late'))`. Quantitative habitat attributes might also be coded as session covariates.

More on simulating multi-session data

Back at the start of this document we used `sim.capthist` to generate `msCH`, a simple multi-session `capthist`. Here we look at various extensions. Generating SECR data is a 2-stage process. The first stage simulates the locations of animals to create an object of class ‘`popn`’; the second stage generates samples from that population according to a particular sampling regime (detector array, number of occasions etc.).

Simulating multi-session populations

By default `sim.capthist` uses `sim.popn` to generate a new population independently for each session. Centres are placed within a rectangular region obtained by buffering around a ‘`core`’ (the traps object passed to `sim.capthist`).

The session-specific populations may also be prepared in advance as a list of ‘`popn`’ objects (use `nsessions > 1` in `sim.popn`). This allows greater control. In particular, the population density may be varied among sessions by making argument `D` a vector of session-specific densities. Other arguments of `sim.popn` do not yet accept multi-session input – it might be useful for ‘`core`’ to accept a list of traps objects (or a list of mask objects if `model2D = “IHP”`).

We can also put aside the basic assumption of independence among sessions and simulate a single population open to births, deaths and movement between sessions. This does not correspond to any model that can be fitted in `secr`, but it allows the effects of non-independence to be examined. See `?turnover` for further explanation.

Multi-session sampling

A multi-session population prepared in advance is passed as the `popn` argument of `sim.capthist`, replacing the usual list (`D`, `buffer` etc.).

The argument ‘`traps`’ may be a list of length equal to `nsessions`. Each component potentially differs not just in detector locations, but also with respect to detector type (‘`detector`’) and resighting regime (‘`markocc`’). The argument ‘`noccasions`’ may also be a vector with a different number of occasions in each session.

Problems specific to multi-session data

Failure of `autoini`

Numerical maximisation of the likelihood requires a starting set of parameter values. This is either computed internally with the function `autoini` or provided by the user. Given multi-session data, the default procedure

is for `secr.fit` to apply `autoini` to the first session only. If the data for that session are inadequate or result in parameter estimates that are extreme with respect to the remaining sessions then model fitting may abort. One solution is to provide start values manually, but that can be laborious. A quick fix is often to switch the session used to compute starting values by changing the details option ‘`autoini`’. For example

```
fit0 <- secr.fit(ovenCH, buffer = 300, details = list(autoini = 2), trace = FALSE)
```

Covariates with incompatible factor levels

Individual or detector covariates used in a multi-session model obviously must appear in each of the component sessions. It is less obvious, and sometimes annoying, that a factor (categorical) covariate should have exactly the same levels in the same order in each component session. The `verify` methods for capthist objects checks that this is in fact the case (remember that `verify` is called by `secr.fit` unless you suppress it).

A common example might be an individual covariate ‘sex’ with the levels “f” and “m”. If by chance only males are detected in one of the sessions, and as a result the factor has a single level “m” in that session, then `verify` will give a warning.

The solution is to force all sessions to use the same factor levels. The function `shareFactorLevels` is provided for this purpose in `secr 3.1`.

Faster fitting of fully session-specific models

Fitting a multi-session model with each parameter stratified by session is unnecessarily slow. In this case no data are pooled across sessions and it is better to fit each session separately. If your data are already in a multi-session capthist object then the speedy solution is

```
fits <- lapply(ovenCH, secr.fit, buffer = 300, trace = FALSE)
class(fits) <- 'secrlist'
predict(fits)
```

```
## $`2005`
##      link      estimate SE.estimate      lcl      ucl
## D      log  0.84652539  0.305472248  0.42641323  1.68054175
## g0     logit  0.02324585  0.008200893  0.01158927  0.04608015
## sigma   log  88.12699702 19.479143588 57.43846822 135.21195541
##
## $`2006`
##      link      estimate SE.estimate      lcl      ucl
## D      log  1.02693138  0.293667997  0.59278023  1.77905403
## g0     logit  0.03012598  0.009279114  0.01639547  0.05471558
## sigma   log  73.02094898 11.712720765 53.42924809 99.79663162
##
## $`2007`
##      link      estimate SE.estimate      lcl      ucl
## D      log  1.07824795  0.287911704  0.64466547  1.80344491
## g0     logit  0.03476923  0.009326454  0.02046603  0.05847188
## sigma   log  76.07522785 11.497176801 56.66665453 102.13132115
##
## $`2008`
##      link      estimate SE.estimate      lcl      ucl
## D      log  1.46809715  0.48868938  0.7777201  2.77131733
## g0     logit  0.02829823  0.01125305  0.0128898  0.06098784
## sigma   log  52.22351449 10.47851706 35.3799434 77.08591943
```

```
##
## $`2009`
##      link estimate SE.estimate      lcl      ucl
## D      log  0.5310134  0.201044860  0.2591587  1.08804095
## g0     logit  0.0244262  0.008809249  0.0119864  0.04913431
## sigma   log  98.9732581 22.950467248 63.1997398 154.99598340
```

The first line (`lapply`) creates a list of ‘secr’ objects. The `predict` method works once we set the class attribute to ‘seclist’ (or you could `lapply(fits, predict)`).

```
fits2 <- secr.fit(ovenCH, model=list(D~session, g0~session, sigma~session), buffer = 300,
                 trace = FALSE)
```

What if we wish to compare this model with a less general one (i.e. with some parameter values shared across sessions)? For that we need the number of parameters, log likelihood and AIC summed across sessions:

```
apply(AIC(fits)[,3:5], 2, sum)
```

```
##      npar    logLik      AIC
## 15.0000 -925.8885 1881.7770
```

```
AIC(fits2)[,3:5]
```

```
##      npar    logLik      AIC
## fits2  15 -925.8885 1881.777
```

AICc is not a simple sum of session-specific AICc and should be calculated manually (hint: use `sapply(ovenCH, nrow)` for session-specific sample sizes).

The unified model fit and separate model fits with `lapply` give essentially the same answers, and the latter approach is faster by a factor of 36.

Using `lapply` does not work if some arguments of `secr.fit` other than ‘capthist’ themselves differ among sessions (as when ‘mask’ is a list of session-specific masks). Then we can use either a ‘for’ loop or the slightly more demanding function `mapply`, with the same gain in speed.

```
# one mask per session
masks <- make.mask(traps(ovenCH), buffer = 300, type = 'trapbuffer')
fits3 <- mapply(secr.fit, ovenCH, mask = masks, MoreArgs = list(buffer = 300,
                    trace = FALSE), SIMPLIFY = FALSE)
class(fits3) <- 'seclist'
```

Caveats

Independence is a strong assumption

If sessions are not truly independent then expect confidence intervals to be too short. This is especially likely when a trend model is fitted to temporal samples with incomplete population turnover between sessions. The product likelihood assumes a new realisation of the underlying population process for each session. If in actuality much of the sampled population remains the same (the same individuals in the same home ranges) then the precision of the trend coefficient will be overstated. Either an open population model is needed (at present that means a Bayesian model) or extra work will be needed to obtain credible confidence limits for the trend (probably some form of bootstrapping).

Parameters are assumed constant by default

Output from `predict.secr` for a multi-session model is automatically stratified by session even when the model does not include ‘session’, ‘Session’ or any session covariate as a predictor (the output simply repeats the constant estimates for each session).

Other notes

The function `ms` returns TRUE if its argument is a multi-session object and FALSE otherwise.

References

- Borchers, D. L. and Efford, M. G. (2008) Spatially explicit maximum likelihood methods for capture–recapture studies. *Biometrics* **64**, 377–385.
- Efford, M. G., Borchers D. L. and Byrom, A. E. (2009) Density estimation by spatially explicit capture–recapture: likelihood-based methods. In: D. L. Thomson, E. G. Cooch and M. J. Conroy (eds) *Modeling Demographic Processes in Marked Populations*. Springer. Pp 255–269.
- Venables, W. N. and Ripley, B. D. (1999) *Modern Applied Statistics with S-Plus*. 3rd edition. Springer.