

Modelling population trend in **secr** 4.6

Murray Efford

2023-10-17

Contents

Introduction	1
Direct estimation of λ	1
The Dlambda parameterization	2
The summary function <code>predictDlambda</code>	2
Covariate and other models	2
Fixing coefficients	3
Technical notes and tips	4
Caveats	4
Spatial variation in density	4
Underestimation of sampling variance	4
References	4
Appendix 1. Older approach	5
Appendix 2. Simulated example with known λ	6
$\hat{\lambda}$ from sequence of density estimates	6
Old style $\hat{\lambda}$ using <code>contr.sdif</code>	6
Session-specific $\hat{\lambda}$	7
Using session covariate	7

Introduction

secr is primarily for estimating closed population density (density at one point in time), but multi-session data may also be modelled to describe population trend over time. The fundamentals of multi-session analysis are described separately in `secr-multisession.pdf`. This document describes methods specifically for population trend, defined as change in density between sessions and measured by the finite rate of increase $\lambda_t = D_{t+1}/D_t$.

Simple trend analysis may be achieved with a particular coding of the session factor as described in earlier versions of `secr-multisession.pdf` and repeated here in Appendix 1. The more flexible methods described here allow the direct estimation of λ_t , possibly including covariate effects. The newer methods require **secr** \geq 4.6.2 2023-09-30.

Direct estimation of λ

A satisfying and flexible approach is to parameterize the density model in terms of the initial density D_1 and the finite rates of increase λ_t for the remaining sessions (λ_1 refers to the density increase between Session 1 and Session 2, etc.).

The Dlambda parameterization

Reparameterization of the density model is achieved internally in `secr.fit` by manipulating the density design matrix to provide a new array of mask-cell- and session-specific densities at each evaluation of the full likelihood. This happens when the details argument ‘Dlambda’ is set to TRUE. The density model (D~) and the fitted beta parameters take on a new meaning determined by the internal function `Dfn2`. More explanation is given later.

Now, fitting the ovenbird model with D~1 results in two density parameters (density in session 1, constant finite rate of increase across remaining sessions):

```
library(secr)
msk <- make.mask(traps(ovenCH[[1]]), buffer = 300, nx = 32)
fit1 <- secr.fit(ovenCH, model = D~1, mask = msk, trace = FALSE,
                details = list(Dlambda = TRUE))
coef(fit1)
```

```
##           beta    SE.beta      lcl      ucl
## D.D1  0.03203014 0.19132423 -0.3429585  0.40701873
## D.D2 -0.06385933 0.07015147 -0.2013537  0.07363502
## g0    -3.56191645 0.15065337 -3.8571916 -3.26664127
## sigma  4.36396664 0.08105189  4.2051079  4.52282542
```

Density-relevant beta parameters have names starting with ‘D.’¹. The first is the log initial density; others relate to the λ parameters.

The summary function predictDlambda

To make the most of the reparameterization we need the special function `predictDlambda` to extract the lambda estimates (the simple `predict` method does not work).

```
predictDlambda (fit1)

##           estimate SE.estimate      lcl      ucl
## D1           1.032549  0.19937327 0.7096677  1.502332
## lambda1  0.938137   0.06589274 0.8176232  1.076414
## lambda2  0.938137   0.06589274 0.8176232  1.076414
## lambda3  0.938137   0.06589274 0.8176232  1.076414
## lambda4  0.938137   0.06589274 0.8176232  1.076414
```

This is an advance on the earlier approach using `sdif` contrasts, as we have constrained λ to a constant.

Covariate and other models

The method allows many covariate models for λ . We can fit a time trend in λ using:

```
fit2 <- secr.fit(ovenCH, model = D~Session, mask = msk, trace = FALSE,
                details = list(Dlambda = TRUE))
predictDlambda (fit2)
```

```
##           estimate SE.estimate      lcl      ucl
## D1           0.9013177 0.21201775 0.5719246  1.420421
## lambda1  1.1473737   0.22123516 0.7889924  1.668541
## lambda2  0.9997690   0.09247382 0.8343249  1.198020
## lambda3  0.8711530   0.08598820 0.7182558  1.056598
## lambda4  0.7590829   0.15341581 0.5128313  1.123580
```

¹Their indices are listed in component ‘D’ of the ‘parindx’ component of the fitted model (e.g. `fit1$parindx$D`), but you are unlikely to need this.

Session-specific λ (lower-case ‘session’) provide a direct comparison with the original analysis:

```
fit3 <- secr.fit(ovenCH, model = D~session, mask = msk, trace = FALSE,
                details = list(Dlambda = TRUE))
predictDlambda (fit3)
```

```
##          estimate SE.estimate      lcl      ucl
## D1          0.9202639   0.2276259 0.5708037 1.483672
## lambda1     1.0468967   0.3313237 0.5713690 1.918187
## lambda2     1.1818222   0.3496452 0.6698626 2.085060
## lambda3     0.7307712   0.2256664 0.4044713 1.320308
## lambda4     0.8420998   0.2941328 0.4330774 1.637426
```

Model selection procedures apply as usual:

```
AIC(fit1, fit2, fit3, criterion = 'AIC')[,-6]
```

```
##          model detectfn npar   logLik      AIC  dAIC  AICwt
## fit1      D~1 g0~1 sigma~1 halfnormal    4 -930.7267 1869.453 0.000 0.5692
## fit2 D~Session g0~1 sigma~1 halfnormal    5 -930.0667 1870.133 0.680 0.4051
## fit3 D~session g0~1 sigma~1 halfnormal    8 -929.8226 1875.645 6.192 0.0257
```

Session covariates are readily applied. The covariate for the second session predicts $\lambda_1 = D_2/D_1$, for the third session predicts $\lambda_2 = D_3/D_2$, etc. The covariate for the first session is discarded (remember D_1 is constant). This all may be confusing, but you can work it out, and it saves extra coding.

```
covs <- data.frame(acov = c(0,2,1,1,2)) # a fabricated covariate
fit4 <- secr.fit(ovenCH, model = D~acov, mask = msk, trace = FALSE,
                details = list(Dlambda = TRUE), sessioncov = covs)
predictDlambda (fit4)
```

```
##          estimate SE.estimate      lcl      ucl
## D1          1.0280556   0.2124441 0.6885799 1.534896
## lambda1     0.9500980   0.2120186 0.6167683 1.463574
## lambda2     0.9303015   0.1452797 0.6862724 1.261104
## lambda3     0.9303015   0.1452797 0.6862724 1.261104
## lambda4     0.9500980   0.2120186 0.6167683 1.463574
```

Fixing coefficients

Another possibility is to fit the model with fixed trend (the second beta coefficient corresponds to lambda, before).

```
fit5 <- secr.fit(ovenCH, model = D~1, mask = msk, trace = FALSE,
                details = list(Dlambda = TRUE, fixedbeta = c(NA, log(0.9), NA, NA)))
predictDlambda(fit5)
```

```
##          estimate SE.estimate      lcl      ucl
## D1          1.115395   0.1538136 0.8523148 1.45968
## lambda1     0.900000   0.0000000 0.9000000 0.90000
## lambda2     0.900000   0.0000000 0.9000000 0.90000
## lambda3     0.900000   0.0000000 0.9000000 0.90000
## lambda4     0.900000   0.0000000 0.9000000 0.90000
```

For comparison, this can be achieved more conventionally by fixing the beta coefficient in a model with log-linear trend over sessions (D~Session):

```
fit6 <- secr.fit(ovenCH, model = D~Session, mask = msk, trace = FALSE,
                details = list(Dlambda = FALSE, fixedbeta = c(NA, log(0.9), NA, NA)))
```

```
t(sapply(predict(fit6), '[', 'D', ))
```

```
##                link estimate SE.estimate lcl      ucl
## session = 2005, Session = 0 "log" 1.115394 0.1538135 0.8523138 1.459679
## session = 2006, Session = 1 "log" 1.003855 0.1384322 0.7670824 1.313711
## session = 2007, Session = 2 "log" 0.9034693 0.124589 0.6903742 1.18234
## session = 2008, Session = 3 "log" 0.8131224 0.1121301 0.6213368 1.064106
## session = 2009, Session = 4 "log" 0.7318101 0.1009171 0.5592031 0.9576951
```

Technical notes and tips

`Dfn2` performs some tricky manipulations. You can see the code by typing `secre:::Dfn2`. A column is pre-pended to the density design matrix specifically to model the initial density; this takes the value one in Session 1 and is otherwise zero. Other columns in the design matrix are set to zero for the first session. Session-specific density on the link (log) scale is computed as the cumulative sum across sessions of the initial log density and the modelled log-lambda values.

Note –

- The model allows detector locations and habitat masks to vary between sessions.
- The coding of `Dfn2` relies on a log link function for density.
- `Dlambda` is ignored for single-session data and conditional-likelihood (CL) models.
- The method is not (yet) suitable for group models.
- The default start values for `D` in `secre.fit` work well: all lambda are initially 1.0 ($\log(\lambda_t) = 0$ for all t).
- If session covariates are used in any model, `AICcompatible()` expects the argument ‘sessioncov’ to be included in all models.

Caveats

Spatial variation in density

`D` for session 1 is constant over space. It is not possible in the present version of `secre` to model simultaneous spatial variation in density or λ , and using `Dlambda` with a density model that includes spatial covariates will cause an error.

Underestimation of sampling variance

Underestimation of sampling variance is expected 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.

The effect is often small. Possible solutions are to fit an open population model (e.g., in `openCR` Efford and Schofield 2020) or to apply some form of bootstrapping.

References

Efford, M. G. and Schofield, M. R. (2020) A spatial open-population capture–recapture model. *Biometrics* **76**, 392–402.

Appendix 1. Older approach

From an earlier version of `secr-multisession.pdf`:

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 $\log(D)$:

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

##           beta    SE.beta      lcl      ucl
## D           0.03171000 0.19145970 -0.3435441  0.40696412
## D.Session -0.06385893 0.07015152 -0.2013534  0.07363552
## g0        -3.56192637 0.15062251 -3.8571411 -3.26671168
## sigma      4.36410837 0.08114847  4.2050603  4.52315645
```

The overall finite rate of increase λ (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] 0.9381373
```

Confidence intervals may also be back-transformed with `exp`. To back-transform the SE use the delta-method approximation $\exp(\beta) * \sqrt{\exp(\text{sebeta}^2)-1} = 0.0658928$.

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. One 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).

```
msh <- make.mask(traps(ovenCH[[1]]), buffer = 300, nx = 32)
fit <- secr.fit(ovenCH, model = D-session, mask = msh, trace = FALSE,
               details = list(contrasts = list(session = MASS::contr.sdif)))
coef(fit)
```

```
##           beta    SE.beta      lcl      ucl
## D           -0.10603039 0.13809600 -0.3766936  0.1646328
## D.session2006-2005  0.04583922 0.30898492 -0.5597601  0.6514385
## D.session2007-2006  0.16705323 0.28968581 -0.4007205  0.7348270
## D.session2008-2007 -0.31365873 0.30181343 -0.9052022  0.2778847
## D.session2009-2008 -0.17184818 0.33930001 -0.8368640  0.4931676
## g0          -3.56086344 0.15063317 -3.8560990 -3.2656278
## sigma       4.36392352 0.08104733  4.2050737  4.5227734
```

This estimates each 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] 1.0469061 1.1818172 0.7307684 0.8421070
```

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

```
## [1] 0.3313547 0.3496653 0.2256750 0.2941511
```

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  
## 1.0469061 1.1818172 0.7307684 0.8421070
```

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

Appendix 2. Simulated example with known λ

In this example the true λ_t are (1, 2, 1, 1, 1.5) for sessions 1–5.

```
# traps and mask  
grid144 <- make.grid(12,12, detector='proximity', spacing = 2)  
grid144mask <- make.mask(grid144, spacing = 0.5, buffer = 4)  
# simulate some data  
set.seed(123)  
# some warnings are suppressed  
pop <- sim.popn(D = 3000, core = grid144, buffer = 4, nsessions = 6,  
               details = list(lambda = c(1, 2, 1, 1, 1.5)))  
ch <- sim.caphist(grid144, popn = pop, detectfn = 'HHN', renumber = FALSE,  
                 detectpar = list(lambda0 = 0.5, sigma = 1))  
summary(ch, terse = TRUE)
```

```
##           1  2  3  4  5  6  
## Occasions  5  5  5  5  5  5  
## Detections 627 671 1228 1188 1163 1833  
## Animals   192 211 388 377 364 569  
## Detectors 144 144 144 144 144 144
```

$\hat{\lambda}$ from sequence of density estimates

```
setNumThreads(18) # adjust as required
```

```
## [1] 18
```

```
fit0 <- secr.fit(ch, mask = grid144mask, detectfn = 'HHN', model = D ~ session,  
                trace = FALSE)  
D <- sapply(predict(fit0), '[[', 'D', 'estimate')  
D[2:6]/D[1:5]
```

```
## session = 2 session = 3 session = 4 session = 5 session = 6  
## 1.0989973 1.8388559 0.9716593 0.9655128 1.5631850
```

Old style $\hat{\lambda}$ using `contr.sdif`

```
fit1 <- secr.fit(ch, mask = grid144mask, detectfn = 'HHN', model = D ~ session,  
                trace = FALSE, details = list(contrasts = list(session = MASS::contr.sdif)))
```

```
# lambda estimates
exp(coef(fit1)[2:6,1])
```

```
## [1] 1.0989573 1.8388750 0.9716453 0.9655213 1.5631831
```

Session-specific $\hat{\lambda}$

```
fit2 <- secr.fit(ch, mask = grid144mask, detectfn = 'HHN', model = D-session,
  trace = FALSE, details = list(Dlambda = TRUE))
predictDlambda(fit2)
```

```
##           estimate SE.estimate          lcl          ucl
## D1          3024.3167722 218.59153305 2625.3322097 3483.936968
## lambda1      1.0989925   0.10987997   0.9038598   1.336252
## lambda2      1.8388655   0.15759569   1.5550098   2.174537
## lambda3      0.9716466   0.07041738   0.8431418   1.119737
## lambda4      0.9655201   0.07117930   0.8357853   1.115393
## lambda5      1.5631800   0.10545086   1.3697862   1.783878
```

Using session covariate

The covariate used here is a perfect match to the original λ - completely unrealistic, but a test of the idea.

```
covs <- data.frame(sesscov = log(c(1, 1, 2, 1, 1, 1.5)))
fit3 <- secr.fit(ch, mask = grid144mask, detectfn = 'HHN', model = D ~ sesscov,
  sessioncov = covs, trace = FALSE, details = list(Dlambda = TRUE))
predictDlambda(fit3)
```

```
##           estimate SE.estimate          lcl          ucl
## D1          3138.7003752 147.14963649 2863.2906572 3440.600772
## lambda1      0.9901158   0.03255250   0.9283424   1.056000
## lambda2      1.9563585   0.14237976   1.6966081   2.255877
## lambda3      0.9901158   0.03255250   0.9283424   1.056000
## lambda4      0.9901158   0.03255250   0.9283424   1.056000
## lambda5      1.4746737   0.04756057   1.3843645   1.570874
```