

ipsecr 1.1 - spatially explicit capture–recapture by inverse prediction

Murray Efford

2022-06-22

Contents

Inverse prediction for capture–recapture estimation	1
Simple example	2
Proxy functions	3
Multi-session models	4
Fitting a density gradient	4
Non-target interference	6
Tuning the algorithm	7
Fractional designs	7
Relationship to package <code>secr</code>	9
Troubleshooting and limitations	9
References	10

This document provides an overview of **ipsecr** 1.1, an R package for spatially explicit capture–recapture analysis (SECR) that uses simulation and inverse prediction instead of maximum likelihood, as in **secr** (Efford 2022), or MCMC, as in various Bayesian approaches. The parent package **secr** included the function `ip.secr`; **ipsecr** extends the capability of that function in its central function `ipsecr.fit` whose arguments are closely modelled on `secr.fit` from **secr**.

Simulation and inverse prediction allows some models to be fitted that strictly cannot be fitted by other methods. Single-catch traps are a common example.

There are limitations: **ipsecr** 1.1 does not yet allow variation in detection parameters (e.g., λ_0 , σ) and NT may vary with session and detector (trap) but not occasion.

Inverse prediction for capture–recapture estimation

The method (Efford 2004; Efford, Dawson and Robbins 2004; see also Carothers 1979 and Pledger and Efford 1998) uses a vector-valued function of the raw data that provides a proxy for each coefficient (beta parameter) in the capture–recapture model. Each proxy is assumed to be monotonically related to the corresponding coefficient.

We use \mathbf{x}_p for the vector of known parameter values at which simulations are performed and \mathbf{y}_p for the vectors of proxies computed from these simulated data. The method fits a multivariate multiple regression over a set of points in parameter space ('box') and inverts that regression to estimate parameter values \mathbf{x} from the observed proxy vector \mathbf{y} .

The default proxy function `proxy.ms` works for simple models with two detection parameters. It uses a naive non-spatial estimate of population size (simply the number detected), the corresponding non-spatial detection probability p , and the 'root-pooled spatial variance' measure from the function `RPSV` in `secr` (argument `CC = TRUE`):

Parameter	Proxy	Proxy scale
Density D	number detected n	log
Detection intercept g_0 or λ_0^*	\hat{p}	cloglog
Detection spatial scale σ	RPSV	log

* depends on detection function

The monotonic relationship is demonstrated by simulation in the following figure.

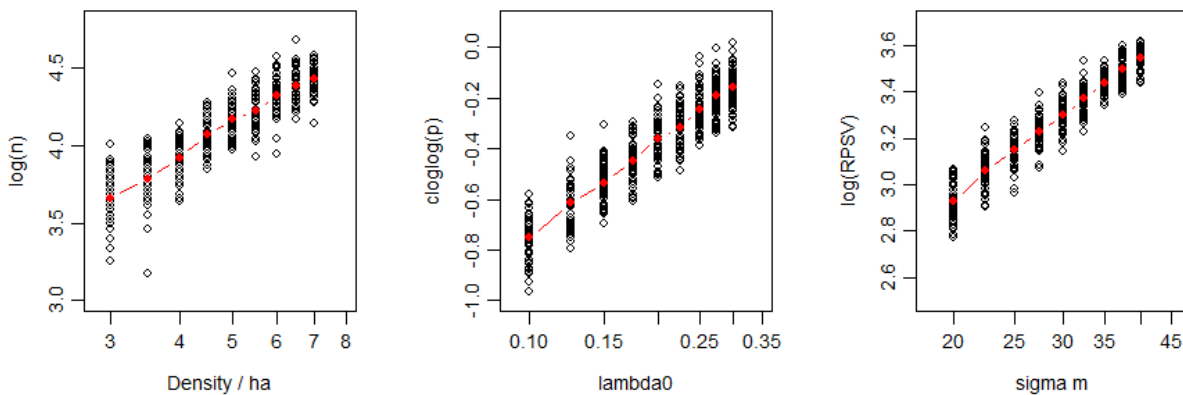


Fig. Monotonic relationships between parameters (x-axis, log scale) and their default proxies (y-axis). 100 single-catch traps at 20-m spacing on a square grid. 50 simulations for each level, with other parameters held at their central value. Red line follows mean.

The algorithm includes these steps

1. Compute proxy function from data
2. Simulate data for parameters at the vertices* of a box around plausible estimates
3. Compute proxy function for each simulation to generate predictor matrix
4. Fit a multivariate multiple regression model with proxies at each vertex as the dependent variables
5. Invert regression model to estimate vector of parameters from the data proxies (1)
6. If the estimated parameters do not all lie inside box, adjust the box and repeat from (2)
7. Simulate at the estimates to obtain variance-covariance matrix

* including some centre points

Simple example

This example uses a simulated single-catch trap dataset in `secr` that is loaded automatically when `ipsecr` is loaded. See here for instructions on reading actual data.

```
library(ipsecr)
setNumThreads(2) # adjust to number of available cores
```

```
ip.single <- ipsecr.fit(captdata, buffer = 100, detectfn = 'HHN')
```

The fitted model is of class 'ipsecr' for which there are the usual methods matching those in `secr` (`print`, `coef`, `predict`, `summary` etc. as detailed below).

```
predict(ip.single)
```

```
##          link  estimate SE.estimate      lcl      ucl
## D          log  5.6056852  0.65615122  4.459985  7.0456969
## lambda0    log  0.4420398  0.07161501  0.322441  0.6059999
## sigma      log 28.2761160  1.27720726 25.881612 30.8921539
```

Proxy functions

The proxy function takes an `secr` capthist object as its first argument. It returns a vector of surrogate values, one for each coefficient (beta parameter) in the model. The default is `proxy.ms`. For example,

```
proxy.ms(captdata)
```

```
##          logn  cloglogp  logRPSV
## 4.33073334 -0.03724765  3.24371994
```

`proxy.ms` works for models with variation in parameters D and NT. Users are free to define their own proxyfn.

Changing the proxy function may have little effect on the estimates. Here we demonstrate an older builtin proxy function `proxyfn0`:

```
ip.single.0 <- ipsecr.fit(captdata, buffer = 100, detectfn = 'HHN', proxyfn = proxyfn0,
  verbose = FALSE)
```

```
# secr function 'collate' works for both secr and ipsecr fits
collate(ip.single, ip.single.0)[1,,]
```

```
## , , D
##
##          estimate SE.estimate      lcl      ucl
## ip.single  5.605685  0.6561512  4.459985  7.045697
## ip.single.0 5.638251  0.7491143  4.350566  7.307068
##
## , , lambda0
##
##          estimate SE.estimate      lcl      ucl
## ip.single  0.4420398  0.07161501  0.3224410  0.6059999
## ip.single.0 0.4328522  0.06449507  0.3237472  0.5787264
##
## , , sigma
##
##          estimate SE.estimate      lcl      ucl
## ip.single  28.27612  1.277207  25.88161  30.89215
## ip.single.0 28.16993  1.366206  25.61697  30.97731
```

Multi-session models

Data may take the form of independent samples. Models are constructed as in `secr` (`secr-multisession.pdf`). In `ipsecr 1.1` detection parameters are common to all sessions, but density `D` and non-target `NT` parameters may depend on session or session covariates as in `secr`.

Fitting a density gradient

Distiller and Borchers (2015) simulated an example with a gradient in population density to demonstrate their method for data from single-catch traps when the time of each capture is known. Inverse prediction may be used to estimate density from these data, but assuming a constant density can result in bias. Here we illustrate a strategy for fitting the density gradient with inverse prediction.

This requires a proxy function that includes a proxy for the density gradient. As the function is a function of the capture histories alone we must use a gradient over detectors to stand for a gradient over points on the habitat mask. The default function `proxy.ms` does this automatically. A model in terms of mask coordinates and covariates is re-cast as a model in terms of detector coordinates and covariates. Proxies are coefficients of a glm for detector-specific counts as a function of the predictors, using a log link.

First simulate some data with an east-west gradient in density.

```
tr <- traps(captdata)
mask <- make.mask(tr)
covariates(mask) <- data.frame(D = (mask$x-265)/20) # for sim.pop
set.seed(1237)
pop <- sim.popn(D = 'D', core = mask, model2D = 'IHP', buffer = 100)
ch <- sim.caphist(tr, popn = pop, detectfn = 'HHN', nooccasions = 5,
  detectpar = list(lambda0 = 0.2, sigma = 25))
# show east-west trend
table(tr[trap(ch), 'x'])
```

```
##
## 365 395 425 455 485 515 545 575 605 635
## 12 17 21 20 28 24 22 23 31 36
```

Note that the `x`- and `y`-coordinates of traps and mask are scaled internally and independently to zero mean and unit variance.

```
ipx <- ipsecr.fit(ch, mask = mask, detectfn = 'HHN', model = list(D=x))
```


```
coef(ipx)
```

```
##          beta    SE.beta      lcl      ucl
## D          2.3877675 0.09743882 2.1967909 2.5787441
## D.x         0.5052538 0.14472281 0.2216023 0.7889053
## lambda0    -1.7474722 0.12401415 -1.9905354 -1.5044089
## sigma       3.2569141 0.05239229 3.1542271 3.3596011
```

```
predict(ipx)
```


```
##      link  estimate SE.estimate      lcl      ucl
## D      log 10.8891567 1.06354999 8.9960979 13.1805739
## lambda0 log 0.1742138 0.02168831 0.1366223 0.2221486
## sigma   log 25.9692748 1.36152396 23.4349173 28.7777091
```

```
plot(predictDsurface(ipx))
plot(tr, add = TRUE)
plotMaskEdge(ipx$mask, add=T)
```



The fitted relationship is linear on the log scale, whereas it was simulated as linear on the natural scale. To better match the original it is possible to use an identity link for D.

```
oldpar <- par(mar = c(4,6,4,4))
# model refers to scaled values of x, so repeat here
m <- mean(traps(ch)$x); s <- sd(traps(ch)$x)
xval <- seq(270,730,10)
xvals <- scale(xval, m, s)
pred <- predict(ipx, newdata = data.frame(x = xvals))
plot(0,0, type='n', xlim= c(270,730), ylim = c(0,40), xlab = 'x', ylab = 'Density')
lines(xval, sapply(pred, '[', 'D', 'estimate'), col = 'red', lwd=2)
abline(-265/20,0.05) # true linear trend
rug(unique(tr$x))    # trap locations
```



```
par(oldpar)
```

Figure: Estimated trend in density from log-linear model (red). Actual trend shown as black line.

Non-target interference

The capthist data object may include an attribute ‘nontarget’ that is a binary $K \times S$ matrix (K detectors, S occasions) indicating whether detector k was disturbed or occupied by non-target species on each occasion s .

Disturbance at single-catch and capped detectors is usually taken to be exclusive of captures: an occupied trap cannot later catch a non-target animal or be disturbed, and vice versa. Multi-catch traps or binary proximity detectors may be disturbed after having registered detections, and both detections and interference are recorded (disturbance is independent of prior detector status, and has the effect of truncating the sampling interval). Variations may be specified with the details argument ‘nontargettype’ that takes one of the values ‘exclusive’, ‘truncated’, ‘erased’ or ‘independent’.

Type	Allowed detectors	Explanation
‘exclusive’	single, capped	detector may be occupied or interfered with, but not both
‘truncated’	all	sampling interval ends at time of interference (default for all detectors except single and capped)
‘erased’	all	interference erases any previous detections at a site in the particular interval
‘independent’	all	interference has no effect on detections, but is recorded and modelled

We illustrate the fitting of a model with non-target captures by simulating disturbance at about 50% of the 235 unoccupied traps in the ‘captdata’ dataset. (This is an ad hoc method, used only for generating demonstration data).

```
set.seed(123)
ch <- captdata
attr(ch, 'nontarget') <- (1-t(apply(ch,2:3,sum))) * (runif(500)>0.5)
summary(ch)$nontarget
```

```
##              1  2  3  4  5 Total
## detectors nontarget 25 27 29 27 28   136
```

The attribute ‘nontarget’ is a matrix with one row per trap and one column per occasion. Entries are either 1 (nontarget or disturbed) or 0 (target species or undisturbed). For a dataset including non-target data `ipsecr.fit` automatically adds the parameter ‘NT’ for the hazard of a single ubiquitous disturbance process competing (in the ‘exclusive’ case) for detectors with each individual-specific capture process. The non-target model may be suppressed by setting the details argument ‘ignorenontarget = TRUE’.

Information on the disturbance process is given by the frequency of nontarget events in the capthist input (attribute ‘nontarget’). The last element of the vector returned by the proxy function `proxy.ms` transforms this to a hazard on the log scale (i.e. complementary log-log of proportion):

```
proxy.ms(ch)

##      logn      cloglogp      logRPSV      cloglogNT
## 4.33073334 -0.03724765  3.24371994 -1.14742163

ip.single.nontarget <- ipsecr.fit(ch, detectfn = 'HHN')
```

The estimate of lambda0 has risen compared to ip.single:

```
predict(ip.single.nontarget)

##      link      estimate SE.estimate      lc1      uc1
## D      log  5.6560584  0.65541402  4.5103152  7.0928518
## lambda0 log  0.6275791  0.11170379  0.4439577  0.8871465
## sigma   log 28.1720779  1.36661240 25.6183886 30.9803236
```

```
## NT      log  0.5501976  0.04808575  0.4637325  0.6527846
```

Note that ‘NT’ cannot be estimated if all traps are occupied.

Tuning the algorithm

The inverse prediction method assumes

1. linear relationships between proxies and parameters (jointly, a hyper-plane), and
2. negligible error in the predicted (mean) proxy for a given set of parameter values.

Performance of the algorithm with respect to these assumptions depends on several components of the ‘details’ argument of `ipsecr.fit` that may be seen as tuning parameters. The first assumption is more plausible for small regions of parameter space (hence a small ‘boxsize’). The second assumption is eased by increasing the number of simulations. Actual settings are a compromise between these requirements and execution time. Users may choose a different compromise.

By default, the size of the box in parameter space is set to \pm ‘boxsize’ units on the link scale. This may be changed to a multiple of the central value with `boxtype = 'relative'`.

It is usual to start with a wide box and to use a narrower box for subsequent simulations, on the assumption that the first box has selected a region of parameter space very close to the solution.

The stopping criterion ‘dev.max’ is used to exit the simulation loop early when sufficient precision has been achieved for all parameters. If `boxtype = 'absolute'` then the criterion is the standard error on the link scale. If `boxtype = 'relative'` then the criterion is the relative standard error (RSE or CV) on the link scale.

Tuning parameter	Default	Description
<code>boxtype</code>	‘absolute’	‘absolute’ or ‘relative’
<code>boxsize1</code>	0.2	size of first box
<code>boxsize2</code>	0.05	boxsize for boxes after the first
<code>centre</code>	3	number of centre points
<code>dev.max</code>	0.002	stopping criterion
<code>min.nsim</code>	20	minimum number of simulations per vertex
<code>max.nsim</code>	2000	maximum number of simulations (excluding for var-covar)
<code>max.nbox</code>	5	maximum number of boxes
<code>max.ntries</code>	2	maximum number of attempts to achieve valid simulation
<code>var.nsim</code>	2000	number of simulations for variance-covariance matrix

Fractional designs

By default, `ipsecr` fits a full factorial design in the parameter space. For NP parameters, simulations are performed at 2^{NP} points, the corners of a hyperrectangle (box), plus possible centre points. The total number of simulations grows rapidly for large NP. Fractional factorial designs may omit some parameter combinations while retaining balance and other desirable characteristics.

The package `FrF2` may be used for fractional factorial designs (Groemping, 2014). Fractional designs are selected by setting `details = list(factorial = 'fractional')` in `ipsecr.fit`.

```
ip.Fr <- ipsecr.fit(captdata, detectfn = 'HHN', details = list(factorial = 'fractional'))  
# ignore message "Registered S3 method overwritten by 'DoE.base'" etc.
```

```
## Registered S3 method overwritten by 'DoE.base':  
##   method      from  
##   factorize.factor conf.design
```

```
collate(ip.single, ip.Fr)[1,,]
```

```
## , , D
##
##          estimate SE.estimate      lcl      ucl
## ip.single 5.605685   0.6561512 4.459985 7.045697
## ip.Fr      5.594845   0.7293871 4.337965 7.215892
##
## , , lambda0
##
##          estimate SE.estimate      lcl      ucl
## ip.single 0.4420398 0.07161501 0.3224410 0.6059999
## ip.Fr      0.4400630 0.07982957 0.3092794 0.6261506
##
## , , sigma
##
##          estimate SE.estimate      lcl      ucl
## ip.single 28.27612   1.277207 25.88161 30.89215
## ip.Fr      28.23530   1.409061 25.60591 31.13469
```

```
ip.single$proctime
```

```
## [1] 49.23
```

```
ip.Fr$proctime
```

```
## [1] 65.02
```

In this example the fractional design was actually slower than the full design because (i) there is little difference in the number of design points when $NP = 3$ (7 vs 11 with 3 centre points) and (ii) the fractional fit went to a third box. Conditions when fractional designs are faster have not been determined - they are probably useful only when parameters are numerous.

The default fractional design is a $1/2$ factorial, illustrated by this code:

```
if (require('FrF2')) {
  NP <- 3
  boxsize <- rep(0.2,3)
  design <- FrF2(2^(NP-1),NP, factor.names = c('D','lambda0','sigma'), ncenter = 2)
  # recast factors as numeric
  design <- sapply(design, function(x) as.numeric(as.character(x)))
  design <- sweep(design, MAR=2, STATS = boxsize, FUN='*')
  # apply to beta
  beta <- log(c(5,0.2,25))
  designbeta <- sweep(design, MAR=2, STATS=beta, FUN='+')
  round(designbeta,3)
}
```

```
      D lambda0 sigma
[1,] 1.609  -1.609 3.219
[2,] 1.809  -1.409 3.419
[3,] 1.809  -1.809 3.019
[4,] 1.409  -1.809 3.419
[5,] 1.409  -1.409 3.019
[6,] 1.609  -1.609 3.219
```

The first and last rows are centre points.

For other designs you may specify the desired arguments of FrF2 as a list e.g., `details = list(factorial = 'fractional', FrF2args = list(nruns = 4, nfactors = 3, ncenter = 3))`.

Relationship to package `secr`

Some functionality of `ipsecr` is shared with `secr`.

Methods for ‘ipsecr’ of generic functions defined in `secr`

Function	Description
<code>makeNewData</code>	generate dataframe suitable for <code>predict.ipsecr</code>
<code>region.N</code>	estimate number of animals in mask or polygon

Methods for ‘ipsecr’ of generic functions from base R and stats

Function	Description
<code>coef</code>	coefficients (beta parameters) of fitted model
<code>plot</code>	plot detection function from fitted model
<code>print</code>	display ‘ipsecr’ object
<code>predict</code>	predict real parameters
<code>summary</code>	summarise fitted model
<code>vcov</code>	variance-covariance matrix of coefficients (beta parameters)

Functions from `secr` that work on fitted ‘ipsecr’ objects as-is

Function	Description
<code>ellipse.secr</code>	plot confidence ellipse for 2 parameters
<code>predictDsurface</code>	predict density for each cell of mask
<code>makeStart</code>	generate starting values from various inputs
<code>collate</code>	tabulate results from multiple models

Troubleshooting and limitations

`ipsecr 1.1` is not intended for models with many sessions or many covariates.

This version does not allow for any variation in detection parameters (`g0`, `lambda0`, `sigma`).

Sparse data

Simulation becomes unreliable with very sparse populations, or sparse sampling, because some simulated datasets will have no recaptures or even no captures. The code allows a failed simulation to be repeated (set the ‘`max.ntries`’ `details` argument > 1), but results probably should not be relied upon when there are warning messages regarding failed simulations.

“simulations for box 1 did not reach target for proxy SE 0.002”

This message may not be terminal. The target precision is arbitrary. Review the ‘Variance bootstrap’ table of verbose output.

References

- Borchers, D. L. and Efford, M. G. (2008) Spatially explicit maximum likelihood methods for capture–recapture studies. *Biometrics* **64**, 377–385.
- Carothers, A. D. (1979) The effects of unequal catchability on Jolly–Seber estimates. *Biometrics* **29**, 79–100.
- Distiller, G. and Borchers, D. L. (2015) A spatially explicit capture–recapture estimator for single-catch traps. *Ecology and Evolution* **5**, 5075–5087.
- Efford, M. G. (2004) Density estimation in live-trapping studies. *Oikos* **106**, 598–610.
- Efford, M. G. (2022). `secr`: Spatially explicit capture–recapture models. R package version 4.5.4. <https://CRAN.R-project.org/package=secr/>
- Efford, M. G., Dawson, D. K. and Robbins C. S. (2004) DENSITY: software for analysing capture-recapture data from passive detector arrays. *Animal Biodiversity and Conservation* **27**, 217–228.
- Groemping, U. (2014). R Package FrF2 for Creating and Analyzing Fractional Factorial 2-Level Designs. *Journal of Statistical Software*, **56**, 1–56. <https://www.jstatsoft.org/article/view/v056i01>.
- Pledger, S. and Efford, M. G. (1998) Correction of bias due to heterogeneous capture probability in capture–recapture studies of open populations. *Biometrics* **54**, 888–898.