

# Adjustment for varying effort in **secr**

Murray Efford

June 7, 2013

When sampling effort varies between detectors or over time in a capture–recapture study we expect a commensurate change in the number of detections. Allowing for known variation in effort when modelling detections has these benefits: (i) detection parameters are related to a consistent unit of effort (e.g. one trap set for one day), (ii) the fit of the detection model is improved, and (iii) trends in the estimates may be modelled without confounding.

Borchers and Efford (2008) allowed the duration of exposure to vary between sampling occasions in their competing-hazard model for multi-catch traps. Efford, Borchers and Mowat (2013) generalised the method to allow joint variation in effort over detectors and over time (occasions), and considered other detector types. This document reiterates some of their material and describes the implementation in ‘secr’.

## Contents

<b>1</b>	<b>Background</b>	<b>2</b>
1.1	Previous methods . . . . .	2
1.2	Linear hazard models . . . . .	3
<b>2</b>	<b>Implementation in secr</b>	<b>3</b>
2.1	Data entry . . . . .	4
2.2	Model fitting . . . . .	4
2.3	Data manipulation and checking . . . . .	7
2.4	Polygons and transects . . . . .	8
2.5	Miscellaneous . . . . .	8
<b>3</b>	<b>References</b>	<b>8</b>

# 1 Background

We use  $T_{sk}$  for the effort on occasion  $s$  at detector  $k$ . It is expected that for small  $T_{sk}$  the number of detections increases linearly with  $T_{sk}$  (saturation may occur at higher effort, depending on the detector type) and that there are no detections when  $T_{sk} = 0$ . Examples of possible effort variables are the number of days that each automatic camera was operated in a tiger study, or the number of rub trees sampled for DNA in each grid cell of a grizzly bear study.

The observations to be modelled are either binary (represented by  $\delta_{sk}$ , an indicator variable for the presence of an animal on occasion  $s$  at binary detector  $k$ ), or integer (represented by  $y_{sk}$ , the number of detections on occasion  $s$  at count detector  $k$ ). We assume the probability of detecting an individual declines with the distance  $d_k(X)$  between a detector  $k$  and the animal's range centre at coordinates  $X = (x, y)$ . The binary relationship is described by a spatial detection function  $g(d_k(X); \theta)$ , where  $\theta$  is a vector of parameters. We define  $g(\cdot)$  so that its intercept when  $d_k = 0$  is a non-spatial scale parameter  $g_0$  ( $0 < g_0 \leq 1$ ). For a concrete example, the half-normal detection function uses  $g(d_k(X); g_0, \sigma) = g_0 \exp(-d_k(X)^2 / (2\sigma^2))$ .

If the data are counts rather than binary observations, we may choose to define the spatial detection function as the decline in expected count with distance  $\lambda(d_k(X); \theta')$ . We use the symbol  $\lambda_0$  for the intercept ( $\lambda_0 > 0$ ). For a particular distribution of the counts we can switch back and forth between the binary and expected-count representations (e.g.,  $g(X) = 1 - e^{-\lambda(X)}$  when the counts are Poisson-distributed). The transformation is non-linear so, for example, a half-normal form for  $g(\cdot)$  does not correspond to half-normal form for  $\lambda(\cdot)$ .

## 1.1 Previous methods

Two options were provided for effort adjustment in earlier versions of **secr**:

i. If only a subset of detectors is used on any occasion  $s$ , and there is no other variation in effort,  $T_{sk}$  is a binary indicator taking the values 0 (detector not used) or 1 (detector used). This case is handled simply by setting log-likelihood components for occasion  $s$  and detector  $k$  to 0 whenever  $T_{sk} = 0$  (via the 'usage' attribute of 'traps' objects in **secr**).

ii. The parameter  $g_0$  or  $\lambda_0$  may be modelled on an appropriate link scale (logit or log) as a linear function of  $T_{sk}$  or other time-varying detector-level covariates.

The first is effective for binary use vs non-use of detectors, but does not encompass other gradations of effort. The second is suboptimal because varying effort is not expected to have a linear additive effect on either of the default link

scales, and the estimation of additional parameters is an unnecessary burden.

## 1.2 Linear hazard models

A more comprehensive approach to effort adjustment follows from the hazard model of Borchers and Efford (2008). We assume detections are independent of each other except as allowed by the competing hazard model for multi-catch traps. The variables to be modelled are  $\delta_{sk}$ , an indicator variable for the presence of an animal on occasion  $s$  at binary detector  $k$ , and  $y_{sk}$ , the number of detections on occasion  $s$  at detector  $k$ , if  $k$  is a ‘count’ detector (i.e., one that can record multiple independent occurrences of an animal).

The properties of various detector types and the expressions for detection as a function of effort  $T_{sk}$  are given in Table 1. Only in the Poisson case is the expected number of detections linear on effort. For binomial count detectors we use a formulation not based directly on instantaneous hazard, as explained in Efford, Borchers and Mowat (2013).

Table 1: Including effort in SECR models for various detector types.  $\delta_{sk}$  is an indicator variable for detection of an individual on occasion  $s$  at detector  $k$ , and  $y_{sk}$  is a count of detections of an individual on occasion  $s$  at detector  $k$ .

Detector type	Model	
Multi-catch trap	$\delta_{sk} \sim \text{Bernoulli}(p_{sk})$	$p_{sk}(X) = [1 - e^{-h_{.s}(X)}]h_{sk}(X)/h_{.s}(X)$
Binary proximity	$\delta_{sk} \sim \text{Bernoulli}(p_{sk})$	$p_{sk}(X) = 1 - [1 - g(d_k(X))]^{T_{sk}}$
Poisson count	$y_{sk} \sim \text{Poisson}(\lambda_{sk})$	$\lambda_{sk}(X) = \lambda_0 T_{sk} g(d_k(X))$
Binomial count	$y_{sk} \sim \text{Binomial}(N_{sk}, p_{sk})$	$N_{sk} = T_{sk}, p_{sk} = g(d_k(X))$

## 2 Implementation in secr

The ‘usage’ attribute of a ‘traps’ object in **secr** is a  $K \times S$  matrix recording the effort ( $T_{sk}$ ) at each detector  $k$  ( $k = 1 \dots K$ ) and occasion  $s$  ( $s = 1 \dots S$ ). If the attribute is missing (NULL) it will be treated as all ones. Extraction and replacement functions are provided (usage() and usage<-()), as demonstrated below). All detector types accept usage data in the same format, except

- polygon                    usage matrix has one row for each polygon
- transect                    usage matrix has one row for each transect
- signal strength            usage is not considered when fitting acoustic models
- binomial counts             $N_{sk}$  is determined by **secr.fit** from usage, rounded to an integer, when **binomN = 1**, or equivalently **binomN = ‘usage’**

## 2.1 Data entry

Usage data may be read as extra columns in the text file of detector coordinates (see `read.traps` and `secr-datainput.pdf`). When only binary (0/1) codes are used, and the `read.traps` argument `binary.usage = TRUE`, separation with white space is optional. This means that ‘01000’ and ‘0 1 0 0 0’ are equivalent. For non-binary values always set `binary.usage = FALSE` and separate with spaces.

The input file for polygons and transects has multiple rows for each unit (one row for each vertex). Usage data are taken from the first vertex for each polygon or transect.

Usage codes may be added to an existing traps object, even after it has been included in a capthist object. For example, the traps object in the demonstration dataset `captdata` starts with no usage attribute:

```
> library(secr, quietly = TRUE)
> usage(traps(captdata))
```

NULL

Suppose that we knew that traps 14 and 15 caught no animals on occasions 1–3 because they were not set. We could construct and assign a binary usage matrix to indicate this:

```
> mat <- matrix(1, nrow = 100, ncol = 5)
> mat[14:15,1:3] <- 0
> usage(traps(captdata)) <- mat
```

## 2.2 Model fitting

Following on from the preceding example, we can confirm our assignment and fit a new model:

```
> summary(traps(captdata))
```

```
Object class      traps
Detector type     single
Detector number   100
Average spacing   30 m
x-range           365 635 m
y-range           365 635 m
Usage range by occasion
```

```

      1 2 3 4 5
min 0 0 0 1 1
max 1 1 1 1 1

> fit <- secr.fit(captdata, trace = FALSE)
> predict(fit)

      link estimate SE.estimate      lcl      ucl
D      log  5.4674928  0.64521578  4.3419569  6.8847937
g0     logit 0.2766941  0.02734596  0.2264047  0.3333407
sigma   log 29.3937605  1.30807419 26.9397720 32.0712869

```

The result in this case is only subtly different from the model with uniform usage (compare `predict(secrdemo.0)`).

Usage is 'hardwired' into the traps object, and will be applied (in the sense of Table 1) when a model is fitted with `secr.fit`. There are two ways to suppress this. The first is to remove or replace the usage attribute. For example,

```
> usage(traps(captdata)) <- NULL
```

returns our demonstration dataset to its original state (this would happen in any case when we started a new R session). The second is to bypass the attribute for a single model fit by calling `secr.fit` with `details = list(ignoreusage = TRUE)`.

For a more interesting example, we simulate data from an array of proximity detectors (such as automatic cameras) operated over 5 occasions, using the default density (5/ha) and detection parameters ( $g_0 = 0.2$ ,  $\sigma = 25$  m) of `sim.caphist`. We choose to expose all detectors twice as long on occasions 2 and 3 as on occasion 1, and three times as long on occasions 4 and 5:

```

> simgrid <- make.grid(nx = 10, ny = 10, detector = 'proximity')
> usage(simgrid) <- matrix(c(1,2,2,3,3), byrow = TRUE, nrow = 100, ncol = 5)
> simCH <- sim.caphist(simgrid)
> summary(simCH)

```

```

Object class      caphist
Detector type     proximity
Detector number   100
Average spacing   20 m
x-range           0 180 m
y-range           0 180 m
Usage range by occasion

```

```

      1 2 3 4 5
min 1 2 2 3 3
max 1 2 2 3 3
Counts by occasion
      1 2 3 4 5 Total
n      17 24 25 28 27 121
u      17 10 4 1 0 32
f       2 4 6 7 13 32
M(t+1) 17 27 31 32 32 32
losses  0 0 0 0 0 0
detections 28 54 61 101 91 335
detectors visited 25 39 47 60 66 237
detectors used 100 100 100 100 100 500

```

Now we fit three models with a half-normal detection function. The first implicitly adjusts for effort. The second has no adjustment because we wipe the usage information. The third allows for occasion-to-occasion variation by fitting a separate  $g_0$  each time. We use `trace = FALSE` to suppress output from each likelihood evaluation, and drop columns 1 and 2 (model and `detectfn`) from the AIC table to save space.

```

> fit.usage <- secr.fit(simCH, trace = FALSE)
> usage(traps(simCH)) <- NULL
> fit.null <- secr.fit(simCH, trace = FALSE)
> fit.t <- secr.fit(simCH, model = g0 ~ t, trace = FALSE)
> AIC(fit.usage, fit.null, fit.t)[,-(1:2)]

```

	npar	logLik	AIC	AICc	dAICc	AICwt
fit.usage	3	-1019.575	2045.149	2046.006	0.000	0.9868
fit.t	7	-1017.982	2049.964	2054.631	8.625	0.0132
fit.null	3	-1050.288	2106.577	2107.434	61.428	0.0000

From the likelihoods we can see that failure to allow for effort (model `fit.null`) dramatically reduces model fit. The fully time-varying model (`fit.t`) captures the variation in detection probability, but at the cost of fitting  $S - 1$  additional parameters. The model with built-in adjustment for effort (`fit.usage`) has the lowest AIC, but how do the estimates compare? This is a task for the `secr` function `collate`.

```

> collate(fit.usage, fit.null, fit.t, newdata = data.frame(t =
  factor(1:5)))[, 'estimate', 'g0']

```

```

          secr1      secr2      secr3
t=1 0.1654736 0.3300741 0.1389921
t=2 0.1654736 0.3300741 0.2668054
t=3 0.1654736 0.3300741 0.3002824
t=4 0.1654736 0.3300741 0.4915173
t=5 0.1654736 0.3300741 0.4512911

```

The null model fits a single ‘average’  $g_0$  across all occasions that is approximately twice the true rate on occasion 1 (0.2). The estimates of  $g_0$  from `fit.t` mirror the variation in effort. The effort-adjusted model estimates the fundamental rate for one unit of effort (0.2).

```
> collate(fit.usage, fit.null, fit.t)[,,, 'D']
```

```

      estimate SE.estimate      lcl      ucl
secr1 3.997740  0.7180849 2.819237 5.668883
secr2 3.990727  0.7170070 2.814046 5.659432
secr3 3.986982  0.7162925 2.811461 5.654008

```

The density estimates themselves are almost entirely unaffected by the choice of model for  $g_0$ . This is not unusual. Nevertheless, the example shows how unbalanced data may be analysed with a minimum of fuss.

Adjustment for varying usage will be more critical in analyses where (i) the variation is confounded with temporal (between-session) or spatial variation in density, and (ii) it is important to estimate the temporal or spatial pattern. For example, if detector usage was consistently high in one part of a landscape, while true density was constant, failure to allow for varying usage might produce a spurious density pattern.

### 2.3 Data manipulation and checking

The various functions in `secr` for manipulating traps and capthist objects (`subset`, `split.traps`, `rbind.capthist`, `MS.capthist`, `join` etc.) attempt to deal with usage intelligently.

When occasions are collapsed or detectors are lumped with the `reduce` method for capthist objects, usage is summed for each aggregated unit.

The function `usagePlot` displays a bubble plot of spatially varying detector usage on one occasion. The arguments `markused` and `markvarying` of `plot.traps` may also be useful.

## 2.4 Polygons and transects

Binary or count data from searches of polygons or transects (Efford 2011) do not raise any new issues for including effort, at least when effort is homogeneous across each polygon or transect. Effects of varying polygon or transect size are automatically accommodated in the models of Efford (2011). Models for varying effort within polygons or transects have not been needed for problems encountered to date. Such variation might in any case be accommodated by splitting the searched areas or transects into smaller units that were more nearly homogeneous (see the `snip()` function for splitting transects).

## 2.5 Miscellaneous

The units of usage determine the units of  $g_0$  or  $\lambda_0$  in the fitted model. This must be considered when choosing starting values for likelihood maximisation. Ordinarily one relies on `secr.fit` to determine starting values automatically (via `autoini`), and a simple linear adjustment for usage, averaged across non-zero detectors and occasions, is applied to the value of  $g_0$  from `autoini`.

Usage values other than 0 and 1 require significant additional computation because the adjustment is re-computed for each combination of detector x occasion x mask point x detection history x finite mixture. Execution speed may be improved in future versions.

Absolute duration does not always equate with effort. Animal activity may be concentrated in part of the day, or older DNA samples from hair snares may fail to amplify (Efford et al. 2013).

## 3 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. (2011) Estimation of population density by spatially explicit capture–recapture analysis of data from area searches. *Ecology* **92**, 2202–2207.
- Efford, M. G., Borchers D. L. and Mowat, G. (2013) Varying effort in capture–recapture studies. *Methods in Ecology and Evolution*. DOI: 10.1111/2041-210X.12049.