

Density surfaces in **secr** 2.9

Murray Efford

2014-11-17

Contents

Introduction	1
Density surfaces - some background	1
Brushtail possum example	2
Using the ‘model’ argument in <code>secr.fit</code>	4
Link function	4
Built-in variables	5
User-provided variables	5
Covariates computed from coordinates	5
Pre-computed resource selection functions	6
Regression splines	7
Prediction and plotting	9
Scaling	9
Potential problems	10
“This is not a density surface”	10
References	12
Appendix: User-provided model functions	12

Introduction

The formulation of spatially explicit capture–recapture (SECR) by Borchers and Efford (2008) allows for population density to vary over space. Models of density may that include spatial covariates (e.g., habitat class), or spatial trend. This document describes the fitting of density models in **secr** 2.9 by means of a worked example.

Density surfaces - some background

SECR fits a state model and an observation model to data from incomplete spatial detections of individuals. The observation model is based on a distance-dependent detection function with parameters g_0 (or λ_0) and σ , as described in [secr-overview.pdf](#). The state model, with which we are concerned here, is a spatial Poisson process for animal range centres. The expected value of the process (centres per unit area) is either homogeneous (constant over space) or inhomogeneous (varying over space). The notation $D(\mathbf{x}; \phi)$ covers both possibilities: density (D) is a function of location (the vector \mathbf{x} represents a pair of x- and y-coordinates). If density is homogeneous its expected value is a flat surface and the parameter ϕ is one number, the density. In most other cases, ϕ is a vector of several parameters.

To model variation in a density surface we need to maximise the full likelihood. Maximising the conditional likelihood (conditional on n , the number of observed individuals) is a way to estimate the observation model; to go from there to a Horvitz-Thompson estimate of density we assume that density is homogeneous. Here we are concerned with inhomogeneous models that are all fitted with `CL = FALSE` in `secr.fit`.

Although $D(\mathbf{x}; \phi)$ may be a smooth function, in `secr` we evaluate it only at the fixed points of a habitat ‘mask’ (think of these as the cell centres of a pixellated or raster representation). A mask defines the region of habitat relevant to a particular study: in the simplest case it is a buffered zone inclusive of the detector locations, but it may exclude interior areas of non-habitat or have an irregular outline.

A density model $D(\mathbf{x}; \phi)$ is specified in the ‘model’ argument of `secr.fit`¹. Spatial covariates, if any, are needed for each mask point; they are stored in the ‘covariates’ attribute of the mask. Results from fitting the model (including the estimated coefficients ϕ) are saved in an object of class ‘secr’. To visualise a fitted density model we first evaluate it at each point on a mask with the function `predictDsurface` to create an object of class ‘Dsurface’. A Dsurface is a mask with added density data, and plotting a Dsurface is like plotting a mask covariate.

Brushtail possum example

For illustration we use a brushtail possum (*Trichosurus vulpecula*) dataset from the Orongorongo Valley, New Zealand. Possums were live-trapped in mixed evergreen forest near Wellington for nearly 40 years (Efford and Cowan 2004). Single-catch traps were set for 5 consecutive nights, three times a year. The dataset ‘OVpossumCH’ has data from the years 1996 and 1997. The study grid was bounded by a shingle riverbed to the north and west. See ?OVpossum in `secr` for more details.

First we import data for the habitat mask from a polygon shapefile included with the package:

```
library(secr)
library(maptools)
olddir <- setwd(system.file('extdata', package = 'secr'))
OVforest <- readShapeSpatial('OVforest')
leftbank <- read.table('leftbank.txt')[21:195,] ## drop some we don't need
setwd(olddir)
```

OVforest is a SpatialPolygonsDataFrame object defined in package `sp`. Now we can build a habitat mask object, selecting the first two polygons in OVforest and discarding the third that lies across the river. The attribute table of the shapefile (and hence OVforest) includes a categorical variable ‘forest’ that is either ‘beech’ or ‘nonbeech’; ‘addCovariates’ attaches these data to each cell in the mask.


```
ovtrap <- traps(OVpossumCH[[1]])
ovmask <- make.mask(ovtrap, buffer = 120, type = 'trapbuffer',
  poly = OVforest[1:2,], spacing = 7.5, keep.poly = FALSE)
ovmask <- addCovariates(ovmask, OVforest[1:2,])
```

Plotting is easy:

```
par(mar = c(1,6,2,8))
forestcol <- terrain.colors(6)[c(4,2)]
plot(ovmask, cov='forest', dots = FALSE, col = forestcol)
plot(ovtrap, add = TRUE)
par(cex = 0.8)
library(GISTools)
```

¹Technically, it may also be specified in a user-written function supplied to `secr.fit` (see [Appendix](#)), but you are unlikely to need this.

```
map.scale(2674670, 5982930, 200, "metres", 4, 50)
north.arrow(2674670, 5982830, 8, lab = 'N')
```



We fit some simple models to data from February 1996 (session 49).

```
fit.0 <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ 1)
fit.Dxy <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ x + y)
fit.Dxy2 <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ x + y + x2 + y2 + xy)
fit.Dforest <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ forest)
```


```
AIC(fit.0, fit.Dxy, fit.Dxy2, fit.Dforest, criterion = 'AIC')[,-2]
```

##		model	npar	logLik	AIC	AICc	dAIC	AICwt
##	fit.Dxy2	D~x + y + x2 + y2 + xy g0~1 sigma~1	8	-1549.33	3114.66	3115.33	0.000	0.4466
##	fit.Dxy	D~x + y g0~1 sigma~1	5	-1552.89	3115.78	3116.06	1.126	0.2543
##	fit.Dforest	D~forest g0~1 sigma~1	4	-1554.17	3116.34	3116.52	1.676	0.1932
##	fit.0	D~1 g0~1 sigma~1	3	-1555.77	3117.54	3117.65	2.877	0.1060

Each of the inhomogeneous models seems marginally better than the null model, but there is little to choose among them.

To visualise the entire surface we compute predicted density at each mask point. For example, we can plot the quadratic surface like this:

```
par(mar = c(1,6,1,8))
surfaceDxy2 <- predictDsurface(fit.Dxy2)
plot(surfaceDxy2, plottype = 'shaded', poly = FALSE, breaks = seq(0,22,2),
      title = 'Density / ha', text.cex = 1)
## Various graphical elements may be added, including contours of the Dsurface:
plot(ovtrap, add = TRUE)
plot(surfaceDxy2, plottype = 'contour', poly = FALSE, breaks = seq(0,22,2), add = TRUE)
lines(leftbank)
```



Following sections expand on the options for specifying and displaying density models.

Using the ‘model’ argument in `secr.fit`

The model argument of `secr.fit` is a list of formulae, one for each ‘real’ parameter² in both the state model (usually just D for density) and the observation model (typically g_0 or λ_{0i} , and σ). A model formula defines variation in each parameter as a function of covariates (including geographic coordinates and their polynomial terms) that is linear on the ‘link’ scale, as in a generalized linear model.

The options differ between the state and observation models. D may vary with respect to group, session or point in space; g_0 , λ_{0i} , and σ may vary by group, session, occasion or latent class (finite mixture), but not with respect to continuous space. This was a choice made in the software design, aiming to tame the complexity that would result if g_0 and σ were allowed to vary continuously in space.

The predictors ‘group’ and ‘session’ behave for D as they do for other real parameters. They determine variation in the expected density for each group or session that is (by default) uniform across space, leading to a homogeneous Poisson model and a flat surface. No further explanation is therefore needed.

Link function The default link for D is ‘log’. It is equally feasible in most cases to choose ‘identity’ as the link (see the `secr.fit` argument ‘link’), and for the null model $D \sim 1$ the estimate will be the same to numerical accuracy, as will estimates involving only categorical variables (e.g., session). However, with an ‘identity’ link the usual (asymptotic) confidence limits will be symmetrical (unless truncated at zero) rather than asymmetrical. In models with continuous predictors, including spatial trend surfaces, the link function will affect the result, although the difference may be small when the amplitude of variation on the surface is small. Otherwise, serious thought is needed regarding which model is biologically more appropriate: logarithmic or linear.

You may wonder why `secr.fit` is ambivalent: link functions have seemed a necessary part of the machinery for capture–recapture modelling since Lebreton et al. (1992). Their key role is to keep the ‘real’ parameter within feasible bounds (e.g., 0–1 for probabilities). In `secr.fit` any modelled value of D that falls below zero is truncated at zero (of course this condition will not arise with a log link).

²Null formulae such as $D \sim 1$ may be omitted, and if a single formula is used, it may be presented on its own rather than in `list()` form.

Built-in variables `secr.fit` automatically recognises the spatial variables `x`, `y`, `x2`, `y2` and `xy` if they appear in the formula for `D`. These refer to the x-coordinate, y-coordinate, x-coordinate² etc. for each mask point, and will be constructed automatically as needed.

The formula for `D` may also include the non-spatial variables `g` (group), `session` (categorical), and `Session` (continuous), defined as for modelling `g0` and `sigma` (see [secr-overview.pdf](#)).

The built-in variables offer limited model possibilities:

Formula	Interpretation
$D \sim 1$	flat surface (default)
$D \sim x + y$	linear trend surface (planar)
$D \sim x + x^2$	quadratic trend in east-west direction only
$D \sim x + y + x^2 + y^2 + xy$	quadratic trend surface

User-provided variables More interesting models can be made with variables provided by the user. These are stored in a data frame as the ‘covariates’ attribute of a mask object. Covariates must be defined for every point on a mask.

Variables may be categorical (a factor or character value that can be coerced to a factor) or continuous (a numeric vector). The habitat variable ‘habclass’ constructed in the Examples section of the `skink` help is an example of a two-class categorical covariate. Remember that categorical variables entail one additional parameter for each extra level.

There are several ways to create or input mask covariates.

1. Read columns of covariates along with the x- and y-coordinates when creating a mask from a dataframe or external file (`read.mask`)
2. Read the covariates dataframe separately from an external file (`read.table`)
3. Infer covariate values by computation on in existing mask (see below).
4. Infer values for points on an existing mask from a GIS data source, such as a polygon shapefile or a `SpatialPolygonsDataFrame` (see `sp` (Pebesma and Bivand 2005), `maptools` (Lewin-Koh, Bivand et al. 2011), and related R packages)

Use the function `addCovariates` for the third and fourth options.

Covariates computed from coordinates Higher-order polynomial terms may be added as covariates if required. For example,

```
covariates(ovmask)[,'x3'] <- covariates(ovmask)$x^3
```


allows a model like $D \sim x + x^2 + x^3$.

If you have a strong prior reason to suspect a particular ‘grain’ to the landscape then this may be also be computed as a new, artificial covariate. This code gives a covariate representing a northwest – southeast trend:

```
covariates(ovmask)[,'NWSE'] <- ovmask$y - ovmask$x - mean(ovmask$y - ovmask$x)
```

Another trick is to compute distances to a mapped landscape feature. For example, possum density in our Orongorongo example may relate to distance from the river; this corresponds roughly to elevation, which we do not have to hand. The `distancetotrap` function of `secr` computes distances from mask cells to the nearest vertex on the riverbank, which are precise enough for our purpose.

```
par(mar = c(1,6,1,8))
covariates(ovmask)[,'DTR'] <- distancetotrap(ovmask, leftbank)
plot(ovmask, covariate = 'DTR', breaks = seq(0,500,50),
     title = 'Distance to river m', dots = FALSE, inset= 0.07)
```



Pre-computed resource selection functions A resource selection function (RSF) was defined by Boyce et al. (2002) as “any model that yields values proportional to the probability of use of a resource unit”. An RSF combines habitat information from multiple sources in a single variable. Typically the function is estimated from telemetry data on marked individuals, and primarily describes individual-level behaviour (3rd-order habitat selection of Johnson (1980)).

However, the individual-level RSF is also a plausible hypothesis for 2nd-order habitat selection i.e. for modelling the relationship between habitat and population density. Then we interpret the RSF as a single variable that is believed to be proportional to the expected population density in each cell of a habitat mask.

Suppose, for example, we have a polygon shapefile (RSF.shp, RSF.dbf etc.) with the attribute “rsf” defined for each polygon. Given mask and capthist objects “habmask” and “myCH”, this code fits a SECR model that calibrates the RSF in terms of population density:

```
library(maptools)
rsfshape <- readShapePoly("RSF")
habmask <- addCovariates(habmask, rsfshape, columns = "rsf")
secr.fit(myCH, mask = habmask, model = D ~ rsf - 1)
```

- “rsf” must be known for every pixel in the habitat mask
- Usually it make sense to fit the density model through the origin (rsf = 0 implies D = 0). This is not true of habitat suitability indices in general.

This is a quite different approach to fitting multiple habitat covariates within **secr**, and one that should be considered. There are usually too few individuals in a SECR study to usefully fit models with multiple covariates of density, even given a large dataset such as our possum example. However, 3rd-order and 2nd-order habitat selection are conceptually distinct, and their relationship is an interesting research topic.

Regression splines Regression splines are a flexible alternative to polynomials for spatial trend analysis. Regression splines are familiar as the smooth terms in ‘generalized additive models’ (gams) implemented (differently) in the base R package **gam** and in Simon Wood’s package **mgcv**.

Some of the possible smooth terms from **mgcv** can be used in model formulae for **secr.fit** – see the help page for ‘smooths’ in **secr**. Smooths are specified with terms that look like calls to the functions **s** and **te**. Smoothness is determined by the number of knots which is set by the user via the argument ‘k’. The number of knots cannot be determined automatically by the penalty algorithms of **mgcv**.

Here we fit a regression spline with the same number of parameters as a quadratic polynomial, a linear effect of the ‘distance to river’ covariate on log(D), and a nonlinear smooth.


```
fit.D6 <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ s(x,y, k = 6))
fit.DDTR <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ DTR)
fit.DDTR3 <- secr.fit(OVpossumCH[[1]], mask = ovmask, model = D ~ s(DTR, k = 3))
```

Now add these to the AIC table and plot the ‘AIC-best’ model:

```
AIC(fit.0, fit.Dxy, fit.Dxy2, fit.Dforest, fit.DDTR, fit.D6, fit.DDTR3, criterion = 'AIC')[,-2]
```


##		model	npar	logLik	AIC	AICc	dAIC	AICwt
## fit.DDTR3	D~s(DTR, k = 3)	g0~1 sigma~1	5	-1552.03	3114.07	3114.34	0.000	0.2661
## fit.Dxy2	D~x + y + x2 + y2 + xy	g0~1 sigma~1	8	-1549.33	3114.66	3115.33	0.592	0.1979
## fit.DDTR	D~DTR	g0~1 sigma~1	4	-1553.39	3114.77	3114.95	0.705	0.1870
## fit.Dxy	D~x + y	g0~1 sigma~1	5	-1552.89	3115.78	3116.06	1.718	0.1127
## fit.D6	D~s(x, y, k = 6)	g0~1 sigma~1	8	-1549.98	3115.95	3116.63	1.886	0.1036
## fit.Dforest	D~forest	g0~1 sigma~1	4	-1554.17	3116.34	3116.52	2.268	0.0856
## fit.0	D~1	g0~1 sigma~1	3	-1555.77	3117.54	3117.65	3.469	0.0470

```
tmp <- predict(fit.DDTR3, newdata = data.frame(DTR = seq(0,400,5)))
par(mar=c(5,8,2,4), pty = 's')
plot(seq(0,400,5), sapply(tmp, '[', 'D','estimate'), ylim = c(0,20),
     xlab = 'Distance from river (m)', ylab = 'Density / ha', type = 'l')
```



Confidence intervals are computed in `predictDsurface` by back-transforming $\pm 2\text{SE}$ from the link (log) scale:

```
par(mar = c(1,1,1,1), mfrow = c(1,2), xpd = FALSE)
surfaceDDTR3 <- predictDsurface(fit.DDTR3, cl.D = TRUE)
plot(surfaceDDTR3, covariate= 'lcl', breaks = seq(0,22,2), legend = FALSE)
mtext(side=3,line=-1.5,'Lower 95% confidence limit of D (possums / ha)')
plot(surfaceDDTR3, plottype = 'contour', breaks = seq(0,22,2), add = TRUE)
lines(leftbank)
plot(surfaceDDTR3, covariate= 'ucl', breaks = seq(0,22,2), legend = FALSE)
mtext(side=3,line=-1.5,'Upper 95% confidence limit of D (possums / ha)')
plot(surfaceDDTR3, covariate= 'ucl', plottype = 'contour', breaks = seq(0,22,2), add = TRUE)
lines(leftbank)
mtext(side=3, line=-1, outer=TRUE, 's(DTR, k = 3) model')
```



Multiple predictors may be included in one ‘s’ smooth term, implying interaction. This assumes isotropy – equality of scales on the different predictors – which is appropriate for geographic coordinates such as x and y in this example. In other cases, predictors may be measured on different scales (e.g., structural complexity of vegetation and elevation) and isotropy cannot be assumed. In these cases a tensor-product smooth (`te`) is appropriate because it is scale-invariant. For `te`, ‘k’ represents the order of the smooth on each axis, and we must fix the number of knots with ‘`fx = TRUE`’ to override automatic selection.

For more on the use of regression splines see the documentation for **mgcv**, the **secr** help page ‘?smooths’, Wood (2006), and Borchers and Kidney (submitted).

Prediction and plotting

Fitting a model provides estimates of its coefficients or ‘beta parameters’; use the **coef** method to extract these from an **secr** object. The coefficients are usually of little use in themselves, but we can use them to make predictions. In order to plot a fitted model we first predict the height of the density surface at each point on a mask. As we have seen, this is done with **predictDsurface**, which has arguments (**object**, **mask** = **NULL**, **se.D** = **FALSE**, **cl.D** = **FALSE**, **alpha** = 0.05). By default, prediction is at the mask points used when fitting the model (i.e. **object\$mask**); specify the **mask** argument to extrapolate the model to a different area.

The output from **predictDsurface** is a specialised mask object called a **Dsurface** (class “c(‘Dsurface’, ‘mask’, ‘data.frame’)”). The covariate dataframe of a **Dsurface** has columns for the predicted density of each group (D.0 if there is only one). Usually when you print a mask you see only the x- and y-coordinates. The **print** method for **Dsurface** objects displays both the coordinates and the density values as one dataframe, as also do the **head** and **tail** methods.

Use the arguments ‘se.D’ and ‘cl.D’ to request computation of the estimated standard error and/or upper and lower confidence limits for each mask point³. If requested, values are saved as additional covariates of the output **Dsurface** (**SE.0**, **lcl.0**, and **ucl.0** if there is only one group).

The plot method for a **Dsurface** object has arguments (**x**, **covariate** = ‘D’, **group** = **NULL**, **plottype** = ‘shaded’, **scale** = 1, ...). **covariate** may either be a prefix (one of ‘D’, ‘SE’, ‘lcl’, ‘ucl’) or any full covariate name. ‘plottype’ may be one of “shaded”, “dots”, “persp”, or “contour”. A coloured legend is displayed centre-right (see ?plot.mask and ?strip.legend for options).

For details on how to specify colours, levels etc. read the help pages for **plot.mask**, **contour** and **persp** (these functions may be controlled by extra arguments to **plot.Dsurface**, using the ‘dots’ convention).

A plot may be enhanced by the addition of contours. This is a challenge, because the **contour** function in R requires a rectangular matrix of values, and our mask is not rectangular. We could make it so with the **secr** function **rectangularmask**, which makes a rectangular **Dsurface** with missing (NA) values of density at all the external points. **plot.Dsurface** recognises an irregular mask and attempts to fix this with an internal call to **rectangularMask**.

Scaling

So far we have ignored the scaling of covariates, including geographic coordinates.

secr.fit scales the x- and y-coordinates of mask points to mean = 0, SD = 1 before using the coordinates in a model. Remember this when you come to use the coefficients of a density model. Functions such as **predictDsurface** take care of scaling automatically. **predict.secr** uses the scaled values (‘newdata’ x = 0, y = 0), which provides the predicted density at the mask centroid. The mean and SD used in scaling are those saved as the ‘meanSD’ attribute of a mask (dataframe with columns ‘x’ and ‘y’ and rows ‘mean’ and ‘SD’).

Scaling of covariates other than x and y is up to the user. It is not usually needed.

The numerical algorithms for maximising the likelihood work best when the absolute expected values are roughly similar for all parameters on their respective ‘link’ scales (i.e. all beta parameters) rather than varying by orders of magnitude. The default link function for D and sigma (log) places the values of these parameters on a scale that is not wildly different to the variation in g0 or lambda0, so this is seldom an issue. In extreme cases you may want to make allowance by setting the **typsize** argument of **nlm** or the **parscale** control argument of **optim** (via the ... argument of **secr.fit**).

³Option available only for models specified in generalized linear model form with the ‘model’ argument of **secr.fit**, not for user-defined functions.

Scaling is not performed routinely by `secr.fit` for distance calculations. Sometimes, large numeric values in coordinates can cause loss of precision in distance calculations (there are a lot of them at each likelihood evaluation). The problem is serious in datasets that combine large coordinates with small detector spacing, such as the Lake Station `skink` dataset. Set `details = list(centred = TRUE)` to force scaling; this may become the default setting in a future version of `secr`.

Potential problems

Modelling density surfaces can be tricky. Recognise when model fitting has failed. If there is no asymptotic variance-covariance matrix, the estimates cannot be trusted. Some forensic work may be needed. If in doubt, try repeating the fit, perhaps starting from the previously fitted values (you can use `secr.fit(..., start = last.model)` where `last.model` is a previously fitted `secr` object) or from new arbitrary values. Problems may result when the discretization is too coarse, so try with smaller mask cells.

You can try another optimization method; `method = 'Nelder-Mead'` is generally more robust than the default gradient-based method. Any method may fail to find the true maximum from a given starting point. We have no experience with simulated annealing ('SANN' in `optim`); it is reputedly effective, but slow. In the `optim` help it is stated ominously that "the 'SANN' method depends critically on the settings of the control parameters. It is not a general-purpose method".

Avoid using '[' to extract subsets from `mask`, `capthist` and other `secr` objects. Use the provided `subset` methods. It is generally safe to use '[' to extract one session from a multi-session object, as in our possum example, but this is not guaranteed. With care, it is also possible to replace selected elements in situ, but note that any change in coordinates will require the attribute 'meanSD' to be recalculated (see `?getMeanSD`).


Do you really want to model density on the log scale? If not, change the link.

"This is not a density surface"⁴

The surfaces we have fitted involve inhomogeneous Poisson models for the distribution of animal home range centres. The models have parameters that determine the relationship of expected density to location or to habitat covariates.

Another type of plot is sometimes presented and described as a 'density surface' – the summed posterior distribution of estimated range centres from a Bayesian fit of a homogeneous Poisson model. A directly analogous plot may be obtained from the `secr` function `fx.total` (see also Borchers and Efford 2008 Section 4.3). The contours associated with the home range centre of each detected individual essentially represent 2-D confidence intervals for its home range centre, given the fitted observation model. Summing these gives a summed probability density surface for the centres of the observed individuals ('D.fx'), and to this we can add an equivalent scaled probability density surface for the individuals that escaped detection ('D.nc'). Both components are reported by `fx.total`, along with their sum ('D.sum') which we plot here for the flat possum model:

⁴http://en.wikipedia.org/wiki/The_Treachery_of_Images



The plot concerns only one realisation from the underlying Poisson model. It visually invites us to interpret patterns in that realisation that we have not modelled. There are serious problems with the interpretation of such plots as ‘density surfaces’:

- attention is focussed on the individuals that were detected; others that were present but not detected are represented by a smoothly varying base level that dominates in the outer region of the plot (contrast this figure with the previous quadratic and DTR3 models).
- the surface depends on sampling intensity, and as more data are added it will change shape systematically. Ultimately, the surface near the centre of a detector array becomes a set of spikes on a barren plain
- the ‘summed confidence interval’ plot is easily confused with the 2-D surface obtained by summing utilisation distributions across animals
- confidence intervals are not available for the height of the probability density surface.

The plots are also prone to artefacts. In some examples we see concentric clustering of estimated centres around the trapping grids, apparently ‘repelled’ from the traps themselves (e.g., plot below for a null model of the Waitarere ‘possumCH’ dataset in **secr**). This phenomenon appears to relate to lack of model fit (unpubl. results).

```
fxsurfaceW <- fx.total(possum.model.0)
par(mar = c(1,5,1,8))
plot(fxsurfaceW, covariate = 'D.sum', breaks = seq(0,5,0.5), poly = FALSE)
plot(traps(possumCH), add = TRUE)
```



References

- Borchers, D. L. and Efford, M. G. (2008) Spatially explicit maximum likelihood methods for capture–recapture studies. *Biometrics* **64**, 377–385.
- Borchers, D. L. and Kidney, D. J. (2014) Flexible density surface estimation using regression splines with spatially explicit capture-recapture data. In prep.
- Boyce, M. S., Vernier, P. R., Nielsen, S. E. and Schmiegelow, F. K. A. (2002) Evaluating resource selection functions. *Ecological modelling* **157**, 281–300.
- Efford, M. G. and Mowat, G. (2014) Compensatory heterogeneity in spatially explicit capture–recapture data. *Ecology* **95**, 1341–1348.
- Johnson, D. H. (1980) The comparison of usage and availability measurements for evaluating resource preference. *Ecology* **61**, 65–71.
- Lebreton, J.-D., Burnham, K. P., Clobert, J. and Anderson, D. R. (1992) Modeling survival and testing biological hypotheses using marked animals: a unified approach with case studies. *Ecological Monographs* **62**, 67–118.
- Lewin-Koh, N. J., Bivand, R., et al. (2011). maptools: Tools for reading and handling spatial objects. R package version 0.8-10. <http://CRAN.R-project.org/package=maptools>
- Pebesma, E.J. and Bivand, R. S. (2005) Classes and methods for spatial data in R. *R News* 5(2). <http://cran.r-project.org/doc/Rnews/>
- Pledger, S. (2000) Unified maximum likelihood estimates for closed capture–recapture models using mixtures. *Biometrics* **56**, 434–442.
- Wood, S. N. (2006) *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC.

Appendix: User-provided model functions

Some density models cannot be coded in the generalized linear model form of the model argument. To alleviate this problem, a model may be specified as an R function that is passed to `secr.fit`, specifically as the component ‘userDfn’ of the list argument ‘details’. We document this feature here, although you may never use it.

The userDfn function must follow some rules.

- It should accept four arguments, the first a vector of parameter values or a character value (below), and the second a ‘mask’ object, a data frame of x and y coordinates for points at which density must be predicted.

Argument	Description
Dbeta	coefficients of density model, or one of c(‘name’, ‘parameters’)
mask	habitat mask object
ngroup	number of groups
nsession	number of sessions

- When called with `Dbeta = ‘name’`, the function should return a character string to identify the density model in output. (This should not depend on the values of other arguments).
- When called with `Dbeta = ‘parameters’`, the function should return a character vector naming each parameter. (When used this way, the call always includes the `mask` argument, so information regarding the model may be retrieved from any attributes of `mask` that have been set by the user).
- Otherwise, the function should return a numeric array with `dim = c(nmask, ngroup, nsession)` where `nmask` is the number of points (rows in `mask`). Each element in the array is the predicted density (natural scale, in animals / hectare) for each point, group and session. This is simpler than it sounds, as usually there will be a single session and single group.

The coefficients form the density part of the full vector of beta coefficients used by the likelihood maximization function (`nlm` or `optim`). Ideally, the first one should correspond to an intercept or overall density, as this is what appears in the output of `predict.secr`. If transformation of density to the ‘link’ scale is required then it should be hard-coded in `userDfn`.

Covariates are available to user-provided functions, but within the function they must be extracted ‘manually’ (e.g., `covariates(mask)$habclass` rather than just ‘`habclass`’). To pass other arguments (e.g., a basis for splines), add attribute(s) to the `mask`.

It will usually be necessary to specify starting values for optimisation manually with the `start` argument of `secr.fit`.

If the parameter values in `Dbeta` are invalid the function should return an array of all zero values.

Here is a ‘null’ `userDfn` that emulates $D \sim 1$ with log link

```
userDfn0 <- function (Dbeta, mask, ngroup, nsession) {
  if (Dbeta[1] == 'name') return ('0')
  if (Dbeta[1] == 'parameters') return ('intercept')
  D <- exp(Dbeta[1]) ## constant for all points
  tempD <- array(D, dim = c(nrow(mask), ngroup, nsession))
  return(tempD)
}
```

We can compare the result using `userDfn0` to a fit of the same model using the ‘model’ argument. Note how the model description combines ‘user.’ and the name ‘0’.

```
library(secr)
model.0 <- secr.fit(captdata, model = D ~ 1, trace = FALSE)
```

```
## Warning: multi-catch likelihood used for single-catch traps
## Warning: using default buffer width 100 m
```

```
userDfn.0 <- secr.fit(captdata, details = list(userDfn = userDfn0), trace = FALSE)
```

```
## Warning: multi-catch likelihood used for single-catch traps
## Warning: using default buffer width 100 m
```

```
AIC(model.0, userDfn.0)
```

```
##              model  detectfn npar  logLik      AIC      AICc dAICc AICcwt
## model.0      D~1 g0~1 sigma~1 halfnormal    3 -759.026 1524.05 1524.38    0    0.5
## userDfn.0 D~userD.0 g0~1 sigma~1 halfnormal    3 -759.026 1524.05 1524.38    0    0.5
```

```
predict(model.0)
```

```
##      link  estimate SE.estimate      lcl      ucl
## D      log  5.479803  0.6467406  4.351622  6.900471
## g0     logit  0.273191  0.0270513  0.223477  0.329274
## sigma   log 29.365832  1.3049377 26.917572 32.036771
```

```
predict(userDfn.0)
```

```
##      link  estimate SE.estimate      lcl      ucl
## D      log  5.479806  0.6467407  4.351624  6.900474
## g0     logit  0.273191  0.0270513  0.223477  0.329274
## sigma   log 29.365833  1.3049380 26.917572 32.036773
```

Not very exciting, maybe, but reassuring!