

Advanced topics in distance sampling

Workshop, 26-30 August 2019

Centre for Research into Ecological and Environmental Modelling

Solution 5. Advanced density surface models

Here is a “solution” for exercise 5. As with any data analysis, there is no correct answer, but this shows how I would approach this analysis. The analysis here is conditional on selecting a detection function in the previous exercises; I’ve chosen `df_hn` and `df_hr_ss_size` that we saved previously (see detection function practical).

Much of the text below is as in the exercise itself, so it should be relatively easy to navigate.

Additional text and code is highlighted using boxes like this.

Load data and packages

```
library(Distance)

## Loading required package: mrds

## This is mrds 2.2.1
## Built: R 3.5.2; ; 2019-08-19 13:07:30 UTC; unix

##
## Attaching package: 'Distance'

## The following object is masked from 'package:mrds':
## 
##     create.bins

library(dsm)

## Loading required package: mgcv

## Loading required package: nlme

## This is mgcv 1.8-28. For overview type 'help("mgcv-package")'.

## Loading required package: numDeriv

## This is dsm 2.2.17
## Built: R 3.5.2; ; 2019-01-20 08:22:08 UTC; unix

library(ggplot2)
library(knitr)
library(plyr)
library(reshape2)
```

Load the data processed from GIS and the fitted detection function objects from the previous exercises:

```
load("sperm-data.RData")
load("df-models.RData")
```

Exploratory analysis

We can do some exploratory analysis by aggregating the counts to each cell and plotting what's going on.

Don't worry about understanding what this code is doing at the moment.

```
# join the observations onto the segments
join_dat <- join(segs, obs, by="Sample.Label", type="full")
# sum up the observations per segment
n <- ddply(join_dat, .(Sample.Label), summarise, n=sum(size, na.rm=TRUE), .drop = FALSE)
# sort the segments by their labels
segs_eda <- segs[sort(segs$Sample.Label),]
# make a new column for the counts
segs_eda$n <- n$n

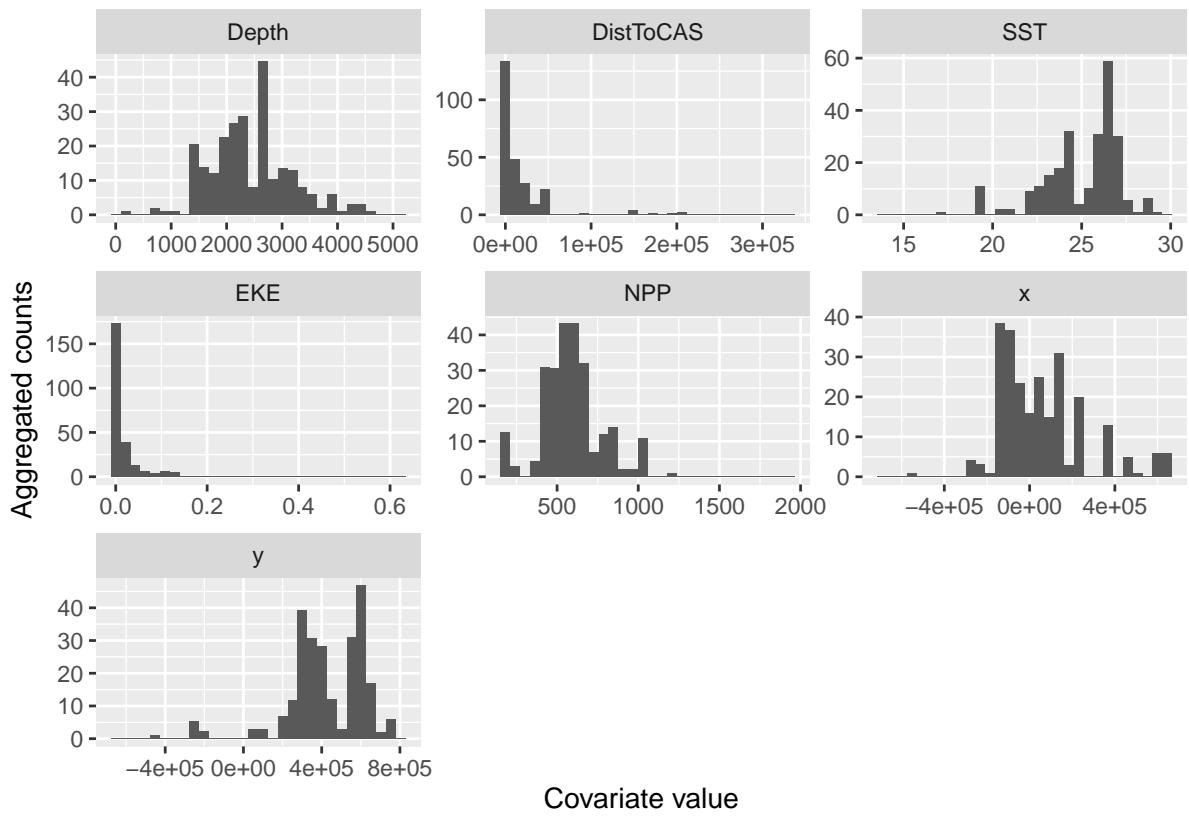
# remove the columns we don't need,
segs_eda$CentreTime <- NULL
segs_eda$POINT_X <- NULL
segs_eda$POINT_Y <- NULL
segs_eda$segment.area <- NULL
segs_eda$off.set <- NULL
segs_eda$CenterTime <- NULL
segs_eda$Effort <- NULL
segs_eda$Length <- NULL
segs_eda$SegmentID <- NULL
segs_eda$coords.x1 <- NULL
segs_eda$coords.x2 <- NULL

# "melt" the data so we have four columns:
#   Sample.Label, n (number of observations),
#   variable (which variable), value (its value)
segs_eda <- melt(segs_eda, id.vars=c("Sample.Label", "n"))
# try head(segs_eda)
```

Finally, we can plot histograms of counts aggregated for different values of the covariates:

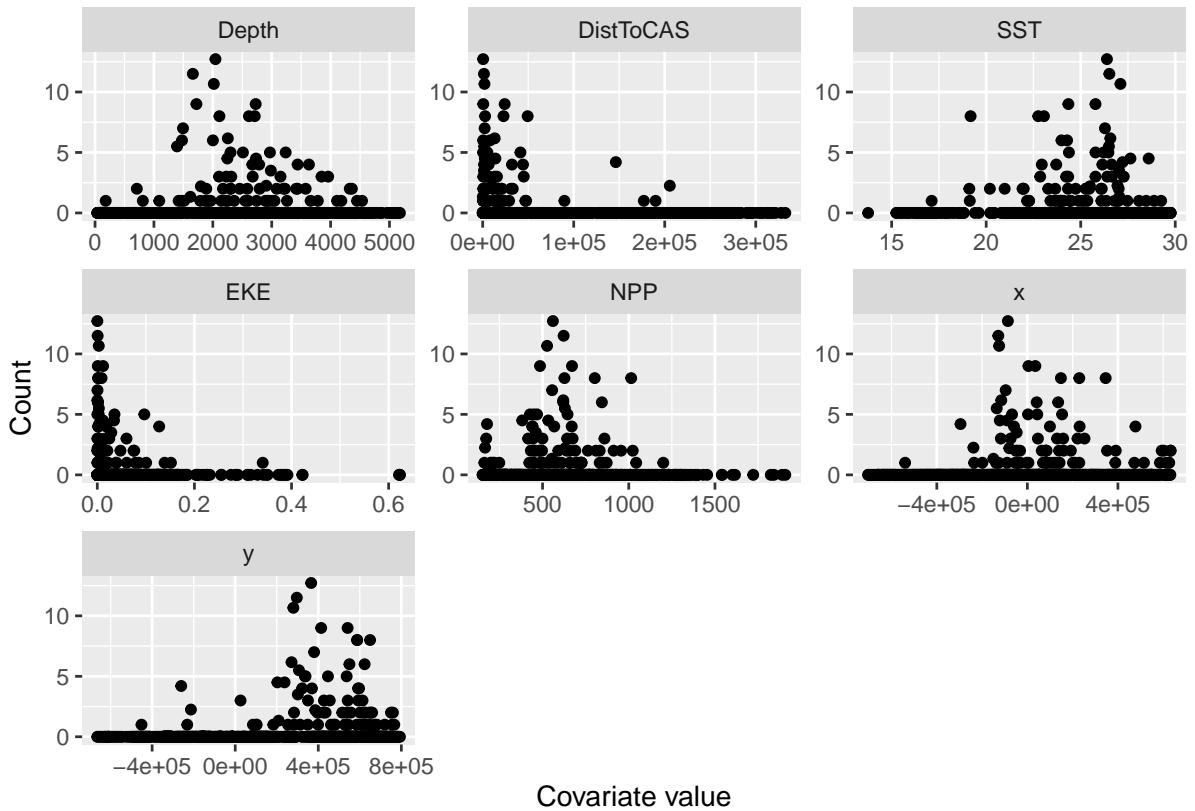
```
p <- ggplot(segs_eda) +
  geom_histogram(aes(value, weight=n)) +
  facet_wrap(~variable, scale="free") +
  xlab("Covariate value") +
  ylab("Aggregated counts")
print(p)
```

```
## `stat_bin()` using `bins = 30`. Pick better value with `binwidth`.
```



We can also just plot the counts against the covariates, note the high number of zeros (but still some interesting patterns):

```
p <- ggplot(segs_eda) +
  geom_point(aes(value, n)) +
  facet_wrap(~variable, scale="free") +
  xlab("Covariate value") +
  ylab("Count")
print(p)
```



These plots give a very rough idea of the relationships we can expect in the model. Notably these plots don't take into account interactions between the variables and potential correlations between the terms, as well as detectability.

Pre-model fitting

As we did in the previous exercise we must remove the observations from the spatial data that we excluded when we fitted the detection function, i.e. those observations at distances greater than the truncation.

```
obs <- obs[obs$distance <= df_hn$ddf$meta.data$width, ]
```

Here we used the value of the truncation stored in the detection function object (`df_hn$ddf`), but we could also use the numeric value (which we can also find by checking the model's `summary()`).

Again note that if you want to fit DSMs using detection functions with different truncation distances, then you'll need to reload the `sperm-data.RData` and do the truncation again for that detection function.

Our new friend +

We can build a really big model using `+` to include all the terms that we want in the model. We can check what covariates are available to us by using `head()` to look at the segment table:

```
head(segs)
```

```
##          CenterTime SegmentID    Length  POINT_X  POINT_Y      Depth
## 1 2004/06/24 07:27:04           1 10288.91 214544.0 689074.3 118.5027
## 2 2004/06/24 08:08:04           2 10288.91 222654.3 682781.0 119.4853
## 3 2004/06/24 09:03:18           3 10288.91 230279.9 675473.3 177.2779
## 4 2004/06/24 09:51:27           4 10288.91 239328.9 666646.3 527.9562
## 5 2004/06/24 10:25:39           5 10288.91 246686.5 659459.2 602.6378
## 6 2004/06/24 11:00:22           6 10288.91 254307.0 652547.2 1094.4402
##   DistToCAS      SST       EKE      NPP coords.x1 coords.x2      x
## 1 14468.1533 15.54390 0.0014442616 1908.129 214544.0 689074.3 214544.0
## 2 10262.9648 15.88358 0.0014198086 1889.540 222654.3 682781.0 222654.3
## 3 6900.9829 16.21920 0.0011704842 1842.057 230279.9 675473.3 230279.9
## 4 1055.4124 16.45468 0.0004101589 1823.942 239328.9 666646.3 239328.9
## 5 1112.6293 16.62554 0.0002553244 1721.949 246686.5 659459.2 246686.5
## 6 707.5795 16.83725 0.0006556266 1400.281 254307.0 652547.2 254307.0
##   y   Effort Sample.Label
## 1 689074.3 10288.91        1
## 2 682781.0 10288.91        2
## 3 675473.3 10288.91        3
## 4 666646.3 10288.91        4
## 5 659459.2 10288.91        5
## 6 652547.2 10288.91        6
```

We can then fit a model with the available covariates in it, each as an `s()` term.

```
dsm_nb_xy_ms <- dsm(count~s(x,y, bs="ts") +
                         s(Depth, bs="ts") +
                         s(DistToCAS, bs="ts") +
                         s(SST, bs="ts") +
                         s(EKE, bs="ts") +
                         s(NPP, bs="ts"),
                         df_hn, segs, obs,
                         family=nb())
summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.114)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(DistToCAS,
##     bs = "ts") + s(SST, bs = "ts") + s(EKE, bs = "ts") + s(NPP,
##     bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7732     0.2295   -90.5   <2e-16 ***
## ---

##
```

```

## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##          edf Ref.df Chi.sq p-value
## s(x,y)      1.8636924    29 19.141 2.90e-05 ***
## s(Depth)     3.4176460     9 46.263 1.65e-11 ***
## s(DistToCAS) 0.0000801     9  0.000  0.9053
## s(SST)       0.0002076     9  0.000  0.5402
## s(EKE)       0.8563344     9  5.172  0.0134 *
## s(NPP)       0.0001018     9  0.000  0.7820
## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) = 0.0947  Deviance explained = 39.3%
## -REML = 382.76  Scale est. = 1           n = 949

```

Notes:

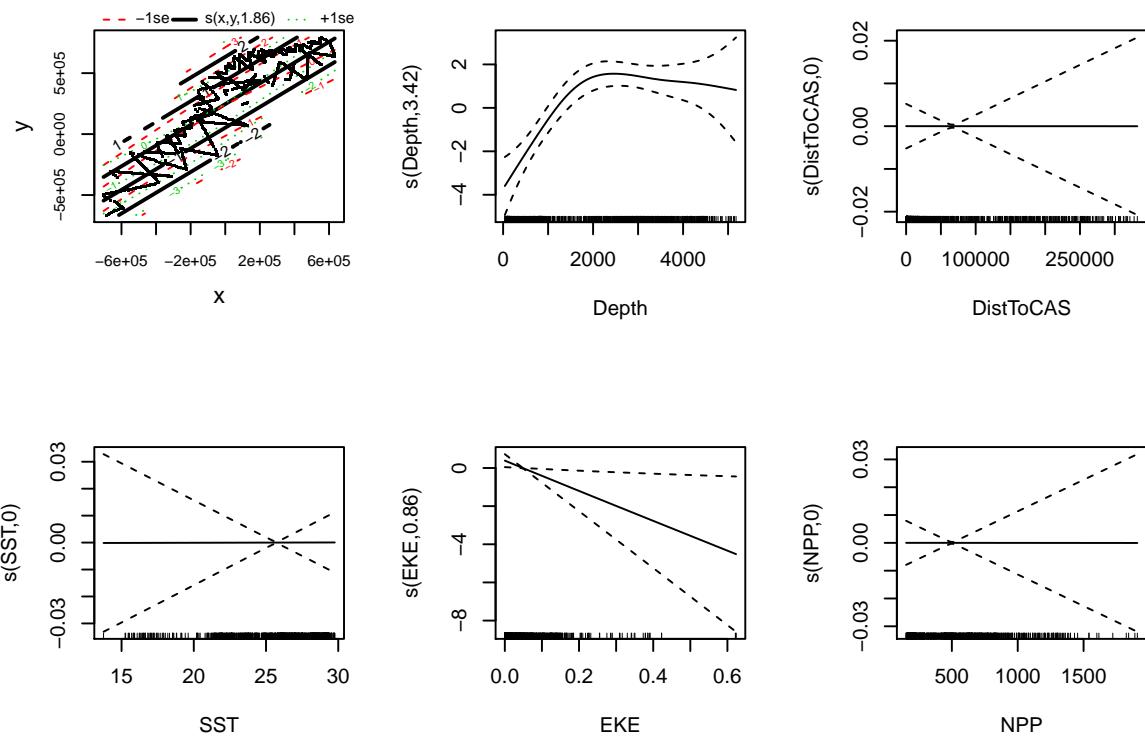
1. We're using `bs="ts"` to use the shrinkage thin plate regression spline. More technical detail on these smooths can be found on their manual page `?smooth.construct.ts.smooth.spec`.
2. We've not specified basis complexity (`k`) at the moment. Note that if you want to specify the same complexity for multiple terms, it's often easier to create a variable that can then be provided to `k` (for example, specify `k1 <- 15` and then set `k=k1` in the required `s()` terms).

Looking at the `summary()` here we can see that many of the terms have been shrunk out of the model. We could simply use this model to make predictions at this point since those shrunken terms will have little effect on the results (since their EDFs are so small), but it makes sense to remove them (because predictions will be faster and we might avoid building many prediction grids for those covariates – more on that later).

Plot

Let's plot the smooths from this model:

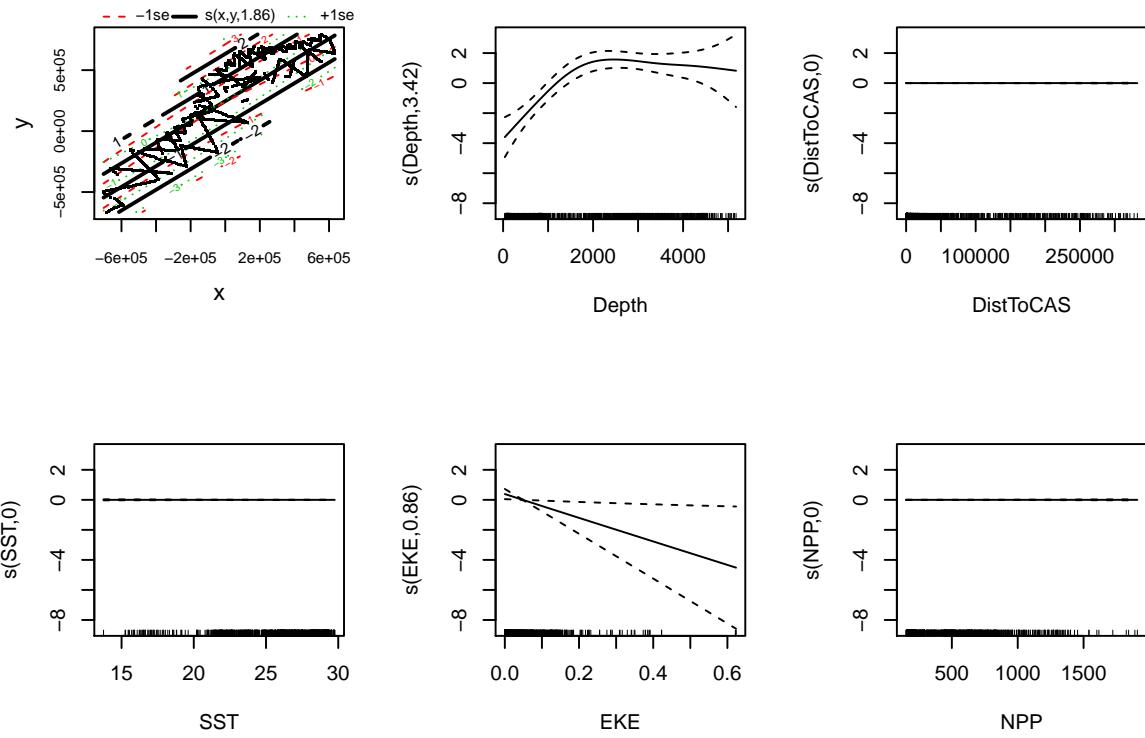
```
plot(dsm_nb_xy_ms, pages=1, scale=0)
```



Notes:

1. Setting `shade=TRUE` gives prettier confidence bands (by default `shade=FALSE`).
2. As with `vis.gam()` the response is on the link scale.
3. `scale=0` puts each plot on a different *y*-axis scale, making it easier to see the effects. Setting `scale=-1` will put the plots on a common *y*-axis scale so you can see the relative influence of the effects, as shown below.

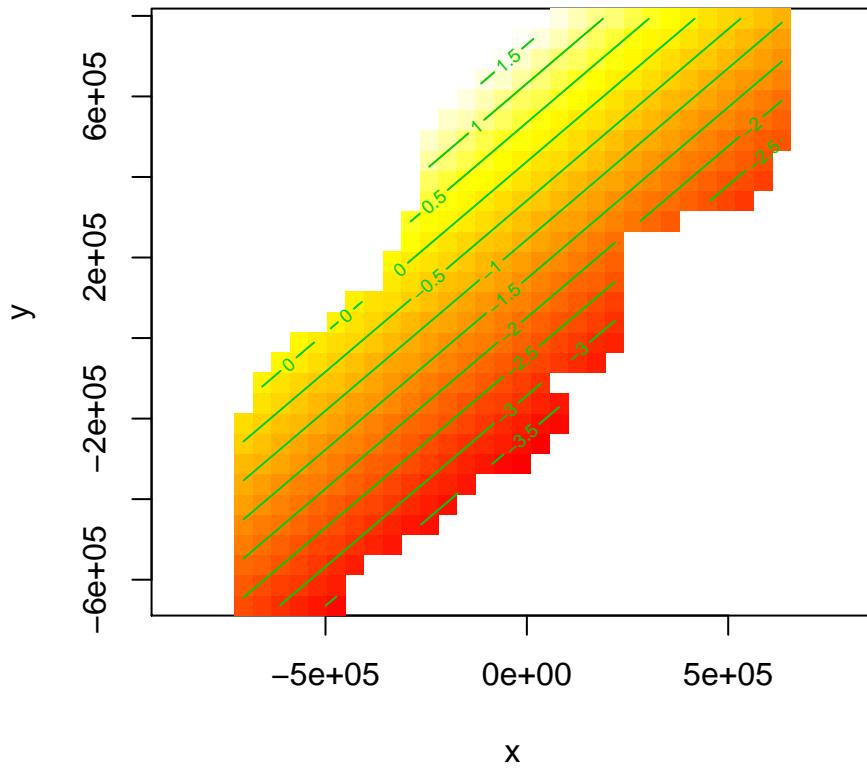
```
plot(dsm_nb_xy_ms, pages=1, scale=-1)
```



We can also plot the bivariate smooth of x and y as we did before, using `vis.gam()`:

```
vis.gam(dsm_nb_xy_ms, view=c("x","y"), plot.type="contour", too.far=0.1,
        main="s(x,y) (link scale)", asp=1)
```

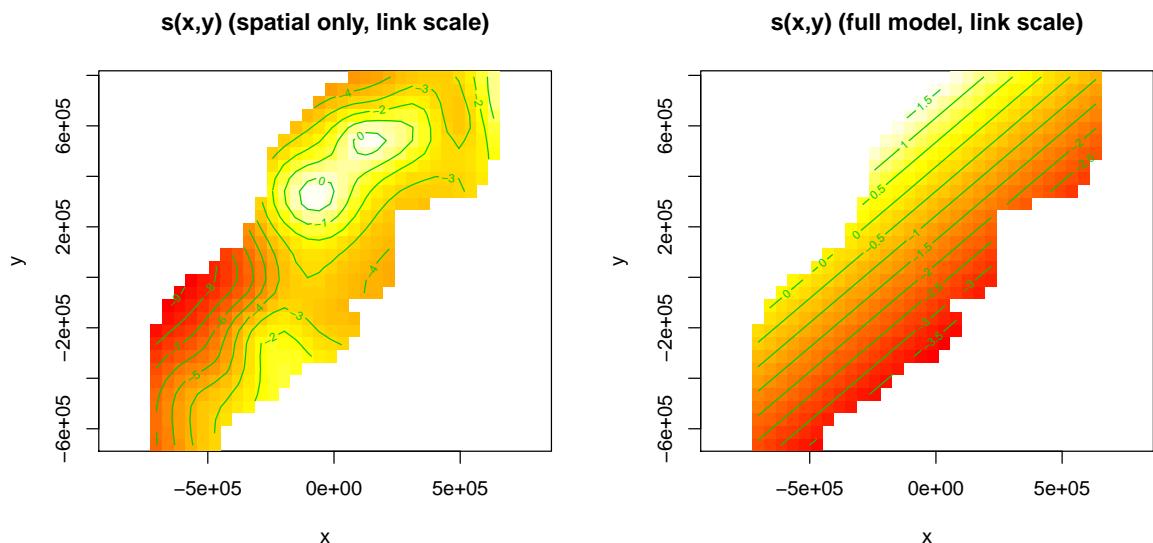
s(x,y) (link scale)



Compare this plot to Figure 1 generated in the previous practical when only x and y were included in the model.

We can load up that saved model and compare the bivariate smooth of location:

```
load("dsms-xy.RData")
par(mfrow=c(1,2))
vis.gam(dsm_nb_xy, view=c("x","y"), plot.type="contour", too.far=0.1,
         main="s(x,y) (spatial only, link scale)", asp=1)
vis.gam(dsm_nb_xy_ms, view=c("x","y"), plot.type="contour", too.far=0.1,
         main="s(x,y) (full model, link scale)", asp=1)
```

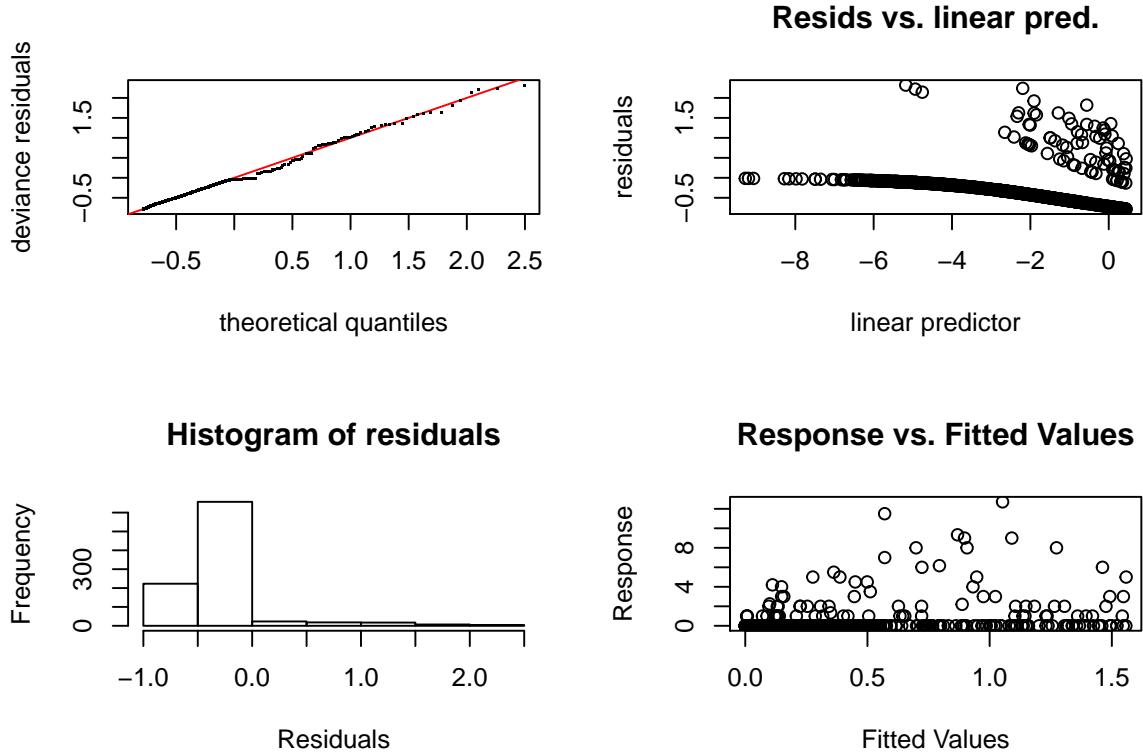


Once additional covariates are included in the model the spatial term is much simpler, as those terms are explaining variation that we'd previously only been able to describe via spatial variation.

Check

As before, we can use `gam.check()` and `rqgam.check()` to look at the residual check plots for this model. Do this in the below gaps and comment on the resulting plots and diagnostics.

```
gam.check(dsm_nb_xy_ms)
```



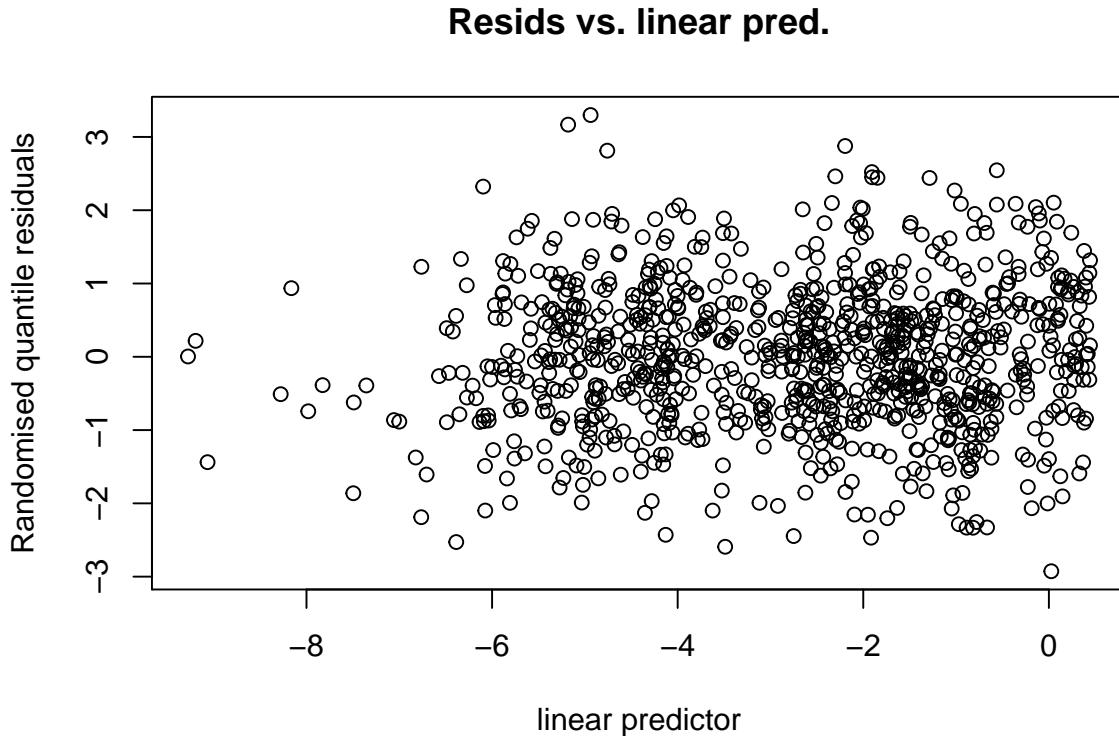
```

## 
## Method: REML   Optimizer: outer newton
## full convergence after 10 iterations.
## Gradient range [-4.045453e-05,1.043377e-05]
## (score 382.7591 & scale 1).
## Hessian positive definite, eigenvalue range [3.174778e-05,30.61621].
## Model rank = 75 / 75
##
## Basis dimension (k) checking results. Low p-value (k-index<1) may
## indicate that k is too low, especially if edf is close to k'.
##
##          k'      edf k-index p-value
## s(x,y)    2.90e+01 1.86e+00    0.55  <2e-16 ***
## s(Depth)   9.00e+00 3.42e+00    0.69   0.045 *
## s(DistToCAS) 9.00e+00 8.01e-05    0.67   0.005 **
## s(SST)     9.00e+00 2.08e-04    0.72   0.245
## s(EKE)     9.00e+00 8.56e-01    0.74   0.700
## s(NPP)     9.00e+00 1.02e-04    0.68   0.010 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

The left side of the `gam.check()` plot shows the Q-Q plot looks okay, the histogram less so.

```
rqgam.check(dsm_nb_xy_ms)
```



`rqgam.check()` shows we don't need to be concerned with non-constant variance.

You might decide from the diagnostics that you need to increase `k` for some of the terms in the model. Do this and re-run the above code to ensure that the smooths are flexible enough. The `?choose.k` manual page can offer some guidance. Generally if the EDF is close to the value of `k` you supplied, it is worth doubling `k` and refitting to see what happens. You can always switch back to the smaller `k` if there is little difference.

It doesn't seem like `k` needs to be increased for any of the terms.

Select terms

As was covered in the lectures, we can select terms by (approximate) *p*-values and by looking for terms that have EDFs significantly less than 1 (those which have been shrunk).

Decide on a significance level that you'll use to discard terms in the model. Remove the terms that are non-significant at this level and re-run the above checks, summaries and plots to see what happens. It's helpful to make notes for yourself as you go.

It's easiest to either comment out the terms that are to be removed (using `#`) or by copying the code chunk above and pasting it below.

Having removed a smooth and reviewed your model, you may decide you wish to remove another. Repeat the process of removing a term and looking at plots and diagnostics again.

Looking again at the model summary:

```

summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.114)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(DistToCAS,
##      bs = "ts") + s(SST, bs = "ts") + s(EKE, bs = "ts") + s(NPP,
##      bs = "ts") + offset(offset)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7732     0.2295   -90.5 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x,y)      1.8636924    29 19.141 2.90e-05 ***
## s(Depth)     3.4176460     9 46.263 1.65e-11 ***
## s(DistToCAS) 0.0000801     9  0.000  0.9053
## s(SST)       0.0002076     9  0.000  0.5402
## s(EKE)       0.8563344     9  5.172  0.0134 *
## s(NPP)       0.0001018     9  0.000  0.7820
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0947  Deviance explained = 39.3%
## -REML = 382.76  Scale est. = 1          n = 949

```

Terms DistToCAS, SST and NPP seem to have been shrunk to be very small. We can start by removing DistToCAS from the model (as it has been shrunk the most):

```

dsm_nb_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(SST, bs="ts") +
                      s(EKE, bs="ts") +
                      s(NPP, bs="ts"),
                      df_hn, segs, obs,
                      family=nb())

```

Looking at the summary of our new model:

```

summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.114)
## Link function: log
##

```

```

## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(SST, bs = "ts") +
##       s(EKE, bs = "ts") + s(NPP, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7732     0.2295   -90.5 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x,y)    1.8636872    29 19.141 2.88e-05 ***
## s(Depth)  3.4176347     9 46.261 1.65e-11 ***
## s(SST)    0.0008110     9  0.000  0.5402
## s(EKE)    0.8563280     9  5.172  0.0134 *
## s(NPP)    0.0000292     9  0.000  0.7820
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0947  Deviance explained = 39.3%
## -REML = 382.76  Scale est. = 1           n = 949

```

Let's now remove NPP:

```

dsm_nb_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(SST, bs="ts") +
                      s(EKE, bs="ts"),
                      df_hn, segs, obs,
                      family=nb())
summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.114)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(SST, bs = "ts") +
##       s(EKE, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7732     0.2295   -90.5 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value

```

```

## s(x,y)    1.864e+00      29 19.145 2.90e-05 ***
## s(Depth)  3.418e+00      9 46.266 1.65e-11 ***
## s(SST)   3.518e-05      9  0.000  0.5403
## s(EKE)   8.563e-01      9  5.172  0.0134 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0947    Deviance explained = 39.3%
## -REML = 382.76  Scale est. = 1           n = 949

```

And now removing SST:

```

dsm_nb_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(EKE, bs="ts"),
                      df_hn, segs, obs,
                      family=nb())
summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.114)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(EKE, bs = "ts") +
##         offset(offset)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7732     0.2295   -90.5   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x,y)    1.8637     29 19.145 2.90e-05 ***
## s(Depth)  3.4176      9 46.266 1.65e-11 ***
## s(EKE)    0.8563      9  5.171  0.0134 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0947    Deviance explained = 39.3%
## -REML = 382.76  Scale est. = 1           n = 949

```

At this point we could leave EKE in the model, but I've decided that my cutoff for removing terms is that they are *** significant, so I'll remove that too:

```

dsm_nb_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts"),
                      df_hn, segs, obs,

```

```

family=nb())
summary(dsm_nb_xy_ms)

##
## Family: Negative Binomial(0.108)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.6731     0.2251  -91.83   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x,y)    1.856    29  21.25 9.06e-06 ***
## s(Depth)  3.621     9  45.74 3.27e-11 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0929  Deviance explained = 37.7%
## -REML = 384.78  Scale est. = 1          n = 949

```

Again note no major differences in the EDFs of the model. This model is also appealing as it contains only static covariates (we'll talk about this more in the lectures).

Compare response distributions

Use the `gam.check()` to compare quantile-quantile plots between negative binomial and Tweedie distributions for the response.

First we need to perform model selection again, when Tweedie is used as the response distribution:

```

dsm_tw_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(DistToCAS, bs="ts") +
                      s(SST, bs="ts") +
                      s(EKE, bs="ts") +
                      s(NPP, bs="ts"),
                      df_hn, segs, obs,
                      family=tw())
summary(dsm_tw_xy_ms)

##
## Family: Tweedie(p=1.277)
## Link function: log

```

```

## 
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(DistToCAS,
##      bs = "ts") + s(SST, bs = "ts") + s(EKE, bs = "ts") + s(NPP,
##      bs = "ts") + offset(off.set)
##
## 
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)    
## (Intercept) -20.675     0.234   -88.36 <2e-16 ***
## ---    
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
## 
## Approximate significance of smooth terms:
##             edf Ref.df    F p-value    
## s(x,y)      1.8875209   29 0.705 3.56e-06 ***
## s(Depth)    3.6794182    9 4.811 2.15e-10 ***
## s(DistToCAS) 0.0000934    9 0.000  0.6797    
## s(SST)      0.3826654    9 0.063  0.2160    
## s(EKE)       0.8196256    9 0.499  0.0178 *  
## s(NPP)      0.0003570    9 0.000  0.8359    
## ---    
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
## 
## R-sq.(adj) =  0.11    Deviance explained = 35.1%
## -REML = 385.04  Scale est. = 4.5486   n = 949

```

Again it doesn't seem like we need to increase `k` from the default values for any of the terms in the model, and again it looks like we should remove `DistToCAS` first:

```

dsm_tw_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(SST, bs="ts") +
                      s(EKE, bs="ts") +
                      s(NPP, bs="ts"),
                      df_hn, segs, obs,
                      family=tw())
summary(dsm_tw_xy_ms)

## 
## Family: Tweedie(p=1.273)
## Link function: log
## 
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(SST, bs = "ts") +
##      s(EKE, bs = "ts") + s(NPP, bs = "ts") + offset(off.set)
## 
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)    
## (Intercept) -20.7242     0.2363   -87.71 <2e-16 ***

```

```

## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##          edf Ref.df   F p-value
## s(x,y)    1.3827805   29 0.138 0.053209 .
## s(Depth)  3.6629371    9 4.086 6.76e-09 ***
## s(SST)    0.0002933    9 0.000 0.546369
## s(EKE)    0.8388135    9 0.557 0.013516 *
## s(NPP)    3.0931492    9 1.530 0.000897 ***
## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) = 0.127 Deviance explained = 37.3%
## -REML = 385.11 Scale est. = 4.4291 n = 949

```

Now we see a bit of a change from the negative binomial model. Rather than having the EDFs stay roughly the same, we see instead that the *p*-value for **s(x,y)** has increased and NPP has become significant. Let's now remove SST and see what happens:

```

dsm_tw_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(EKE, bs="ts") +
                      s(NPP, bs="ts"),
                      df_hn, segs, obs,
                      family=tw())
summary(dsm_tw_xy_ms)

##
## Family: Tweedie(p=1.279)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(EKE, bs = "ts") +
##        s(NPP, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -20.673     0.234  -88.33  <2e-16 ***
## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##          edf Ref.df   F p-value
## s(x,y)    1.897e+00   29 0.707 1.76e-05 ***
## s(Depth)  3.695e+00    9 5.024 1.08e-10 ***
## s(EKE)    8.106e-01    9 0.470  0.0216 *
## s(NPP)    3.667e-05    9 0.000  0.6755
## ---

```

```

## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.105  Deviance explained = 34.9%
## -REML = 385.09  Scale est. = 4.5733    n = 949

Removing SST has moved us back to roughly the smooth terms EDF/p-values we had previously. Let's remove NPP...

dsm_tw_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts") +
                      s(EKE, bs="ts")),
                     df_hn, segs, obs,
                     family=tw())
summary(dsm_tw_xy_ms)

##
## Family: Tweedie(p=1.279)
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + s(EKE, bs = "ts") +
##       offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -20.673     0.234   -88.33  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df      F p-value
## s(x,y)    1.8969     29 0.707 1.76e-05 ***
## s(Depth)  3.6949      9 5.024 1.08e-10 ***
## s(EKE)    0.8106      9 0.470  0.0216 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.105  Deviance explained = 34.9%
## -REML = 385.09  Scale est. = 4.5733    n = 949

```

Again, we could leave EKE in the model but let's remove it:

```

dsm_tw_xy_ms <- dsm(count~s(x,y, bs="ts") +
                      s(Depth, bs="ts"),
                     df_hn, segs, obs,
                     family=tw())
summary(dsm_tw_xy_ms)

##
## Family: Tweedie(p=1.268)

```

```

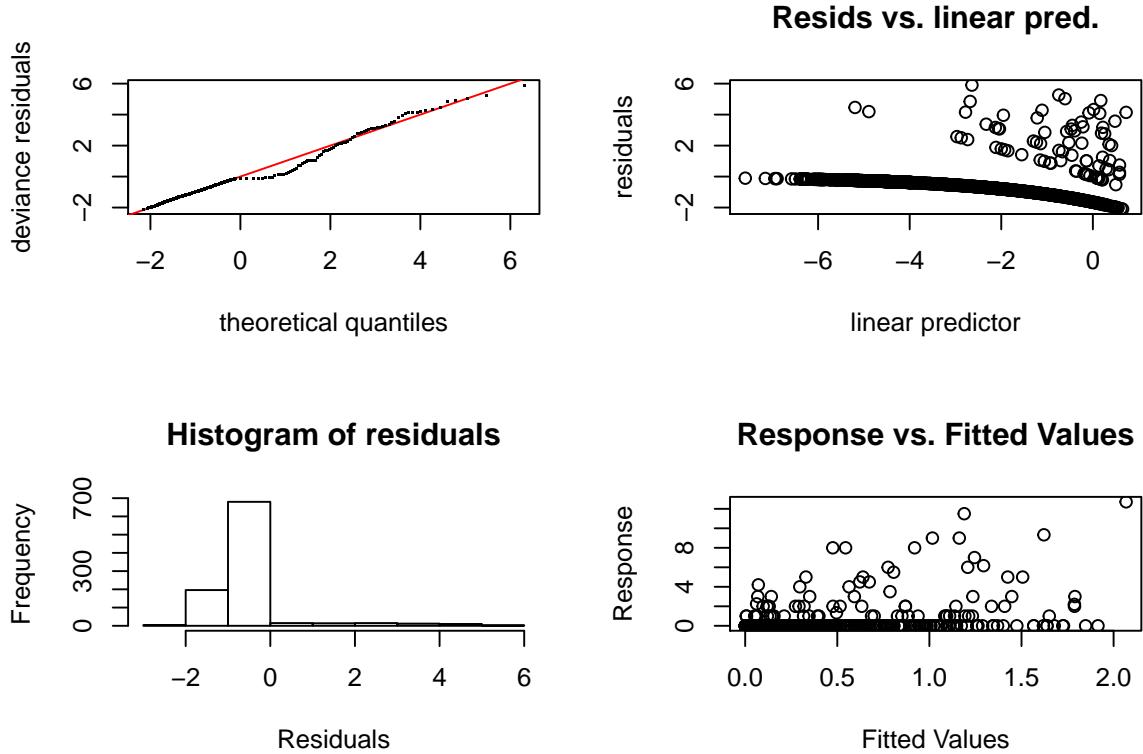
## Link function: log
##
## Formula:
## count ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -20.7240     0.2425  -85.46  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df      F p-value
## s(x,y)    6.443     29 1.322 4.75e-08 ***
## s(Depth)  3.611      9 4.261 1.49e-10 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.141  Deviance explained =   38%
## -REML = 389.86  Scale est. = 4.3516  n = 949

```

We are now left with the same model structure as `dsm_nb_xy_ms`, but by rather a different path. This demonstrates some potential instabilities in the model and the effect of using a different response distribution (and different mean-variance relationship).

Let's look at check plots for this model:

```
gam.check(dsm_tw_xy_ms)
```



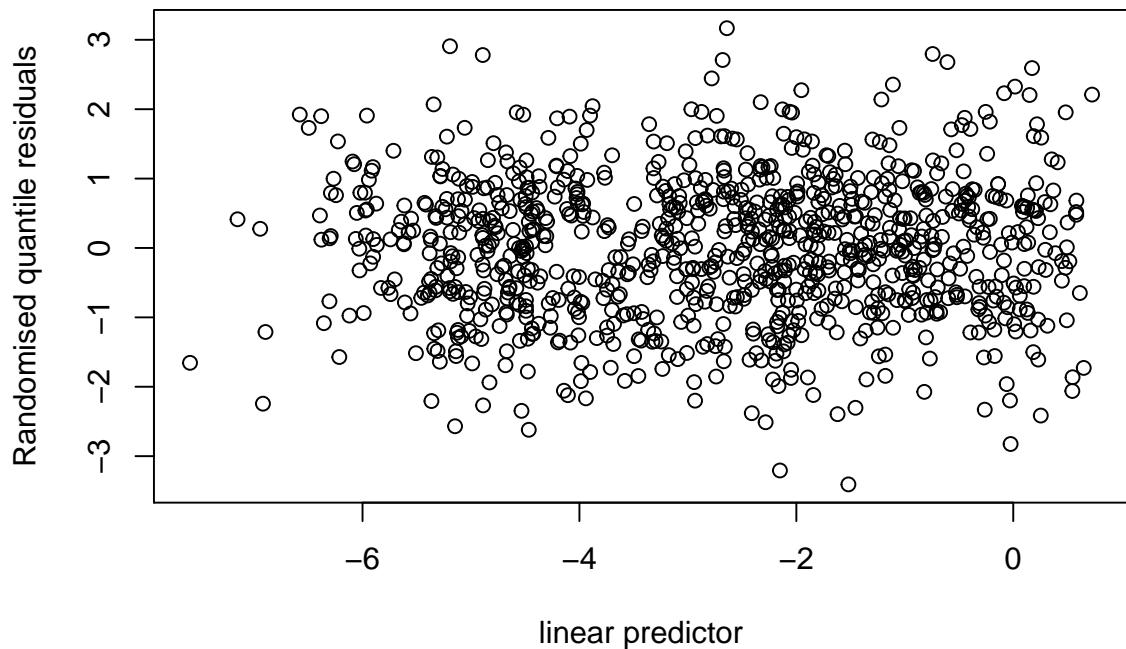
```

## 
## Method: REML   Optimizer: outer newton
## full convergence after 6 iterations.
## Gradient range [-1.604758e-05,4.86767e-06]
## (score 389.8638 & scale 4.351606).
## Hessian positive definite, eigenvalue range [0.1359107,303.1926].
## Model rank = 39 / 39
##
## Basis dimension (k) checking results. Low p-value (k-index<1) may
## indicate that k is too low, especially if edf is close to k'.
##
##          k'    edf k-index p-value
## s(x,y) 29.00  6.44    0.63  <2e-16 ***
## s(Depth) 9.00  3.61    0.81    0.42
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
rqgam.check(dsm_tw_xy_ms)

## Loading required namespace: tweedie

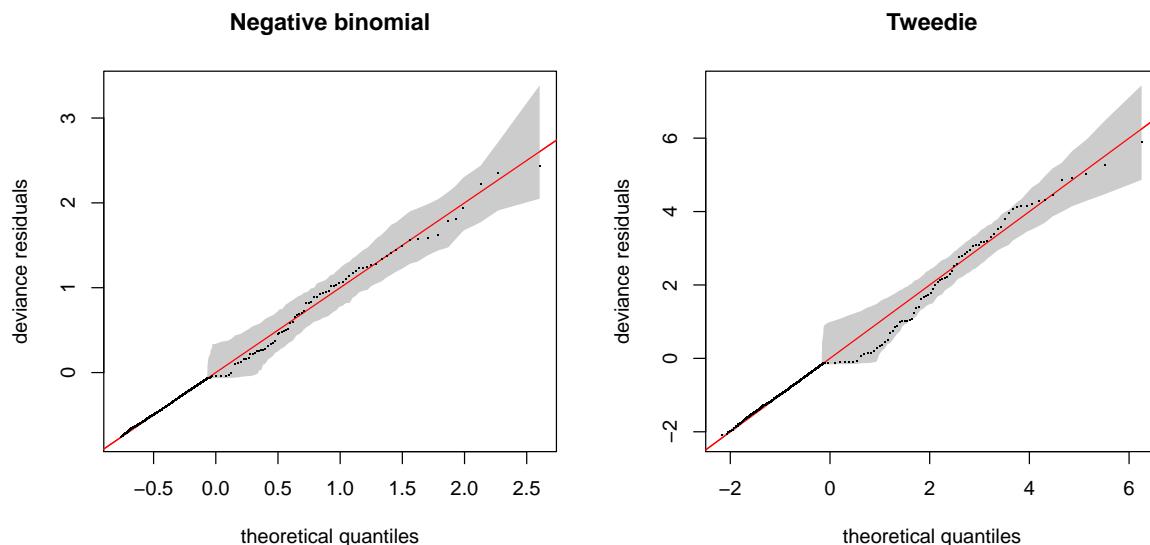
```

Resids vs. linear pred.



These plots look okay and don't show any major issues. Let's compare the Q-Q plots for the negative binomial and Tweedie models side-by-side:

```
par(mfrow=c(1, 2))
qq.gam(dsm_nb_xy_ms, main="Negative binomial", rep=200)
qq.gam(dsm_tw_xy_ms, main="Tweedie", rep=200)
```



There's not much between these two (all points inside the reference bands). I might be inclined to pick the negative binomial model as the points are slightly closer to the line in the Q-Q plot here.

Estimated abundance as a response

So far we have only looked at models with `count` as the response. Try using a detection function with covariates and use `abundance.est`, instead of `count`, as the response in the chunk below:

Let's test out both negative binomial and Tweedie distributions with estimated abundance as the response.

Starting with negative binomial. To save space I've commented the terms out as I've gone and annotated them with numbers to show the sequence in which terms dropped out of the model.

```
# Negative binomial
dsm_nb_xy_ms_ae <- dsm(abundance.est~s(x,y, bs="ts") +
                           s(Depth, bs="ts"), # +
                           # s(DistToCAS, bs="ts") + # 3
                           #s(SST, bs="ts") + # 1
                           #s(EKE, bs="ts"), #+ # 4
                           #s(NPP, bs="ts"), # 2
                           df_hr_ss_size, segs, obs,
                           family=nb())
summary(dsm_nb_xy_ms_ae)

##
## Family: Negative Binomial(0.046)
## Link function: log
##
## Formula:
## abundance.est ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.268     0.206   -98.4    <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x,y)     1.317     29  14.73 6.63e-05 ***
## s(Depth)   3.605      9  56.79 1.64e-14 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.0579  Deviance explained = 31.7%
## -REML = 506.07  Scale est. = 1          n = 949
```

Again the model structure is the same at the end of the selection, but it takes a different route.

```

# Tweedie
dsm_tw_xy_ms_ae <- dsm(abundance.est ~ s(x, y, bs="ts") +
                           s(Depth, bs="ts"), # +
                           #s(DistToCAS, bs="ts") + # 2
                           #s(SST, bs="ts") + # 3
                           #s(EKE, bs="ts"), # +
                           #s(NPP, bs="ts"), # 1
                           df_hr_ss_size, segs, obs,
                           family=tw())
summary(dsm_tw_xy_ms_ae)

##
## Family: Tweedie(p=1.216)
## Link function: log
##
## Formula:
## abundance.est ~ s(x, y, bs = "ts") + s(Depth, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -20.3022    0.2435  -83.37   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df   F p-value
## s(x,y)    7.567    29 1.445 2.40e-08 ***
## s(Depth)  3.639     9 4.150 2.09e-10 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.163  Deviance explained = 38.6%
## -REML = 460.55  Scale est. = 8.2957   n = 949

```

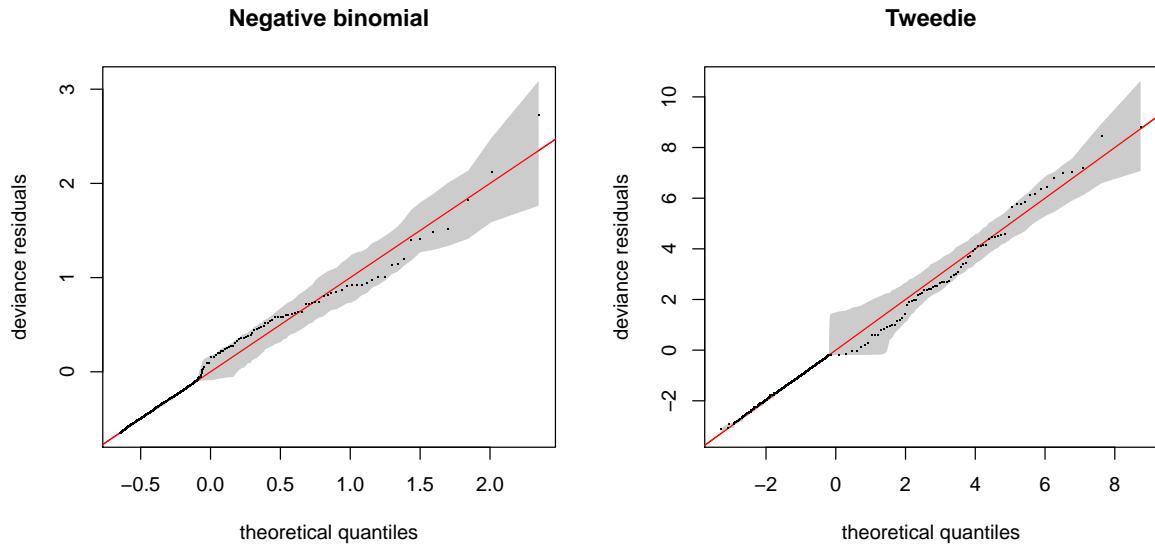
We see the same terms get selected here too, reassuring!

Now quickly plotting the Q-Q plots to compare Tweedie and negative binomial models:

```

par(mfrow=c(1,2))
qq.gam(dsm_nb_xy_ms_ae, main="Negative binomial", rep=200)
qq.gam(dsm_tw_xy_ms_ae, main="Tweedie", rep=200)

```



This is a close call, neither distribution appears to perform considerably better than the other in Q-Q plot terms. We might prefer the Tweedie model as there is a smaller distance between point and line at larger values.

Concurvity

Checking concurvity of terms in the model can be accomplished using the `concurvity()` function.

```
concurvity(dsm_nb_xy_ms)
```

```
##           para      s(x,y)    s(Depth)
## worst  1.331907e-23 0.9780982 0.9780982
## observed 1.331907e-23 0.3768038 0.7112880
## estimate 1.331907e-23 0.2247389 0.9067347
```

By default the function returns a matrix of a measure of concurvity between one of the terms and the rest of the model.

Compare the output of the models before and after removing terms.

We can see from the above that by the “worst” measure of concurvity, `s(x,y)` and `s(Depth)` are highly concurve with the rest of the model, respectively. We expect to see this in our models in general, since depth can be expressed as a function of location pretty easily (thinking about how bathymetry changes as we move away from the shore).

Looking at how bad things could be, let’s re-fit our model with all the terms in and look at those results:

```
dsm_nb_xy_all <- dsm(count~s(x,y, bs="ts") +
                         s(Depth, bs="ts") +
                         s(DistToCAS, bs="ts") +
                         s(SST, bs="ts") +
                         s(EKE, bs="ts") +
```

```

    s(NPP, bs="ts"),
    df_hn, segs, obs,
    family=nb())
concurvity(dsm_nb_xy_all)

##           para      s(x,y)   s(Depth)  s(DistToCAS)   s(SST)      s(EKE)
## worst     2.539199e-23 0.9963493 0.9836597 0.9959057 0.9772853 0.7702479
## observed  2.539199e-23 0.8597372 0.8277050 0.9879372 0.9523512 0.6746585
## estimate   2.539199e-23 0.7580838 0.9272203 0.9642030 0.8978412 0.4906765
##           s(NPP)
## worst     0.9727752
## observed 0.9525363
## estimate 0.8694619

```

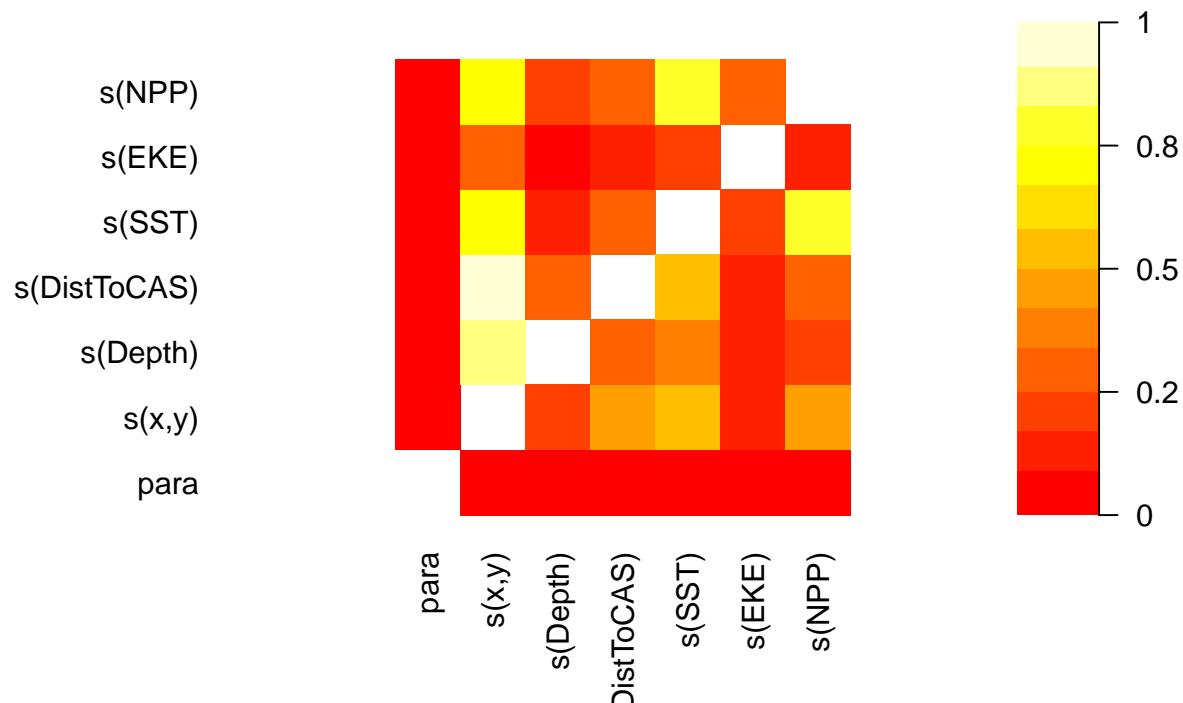
Here we see (generally speaking) most of the terms are highly concurve with the rest of the model using most of the measures of concurvity. We'll look at term-term concurvity next...

Reading these matrices can be a little laborious and not much fun. Here is a little bit of code to visualise the concurvity *between terms* in a model by colour coding the matrix.

Again compare these plots for models with different terms.

Visualising the term-term concurvity in the model with all terms:

```
vis.concurvity(dsm_nb_xy_all)
```



Here we can see that all terms are strongly concurve with the spatial smooth – to be

expected. We also see that NPP and SST are highly concurve, which seems biologically plausible.

Sensitivity

Compare bivariate and additive spatial effects

If we replace the bivariate smooth of location ($s(x, y)$) with an additive terms ($s(x) + s(y)$), we may see a difference in the final model (different covariates selected). For the negative binomial distribution:

```
dsm_nb_x_y_ms <- dsm(count~s(x, bs="ts") +
                         s(y, bs="ts") +
                         s(Depth, bs="ts") +
                         s(DistToCAS, bs="ts") +
                         s(SST, bs="ts") +
                         s(EKE, bs="ts") +
                         s(NPP, bs="ts"),
                         df_hn, segs, obs,
                         family=nb())
summary(dsm_nb_x_y_ms)

##
## Family: Negative Binomial(0.116)
## Link function: log
##
## Formula:
## count ~ s(x, bs = "ts") + s(y, bs = "ts") + s(Depth, bs = "ts") +
##       s(DistToCAS, bs = "ts") + s(SST, bs = "ts") + s(EKE, bs = "ts") +
##       s(NPP, bs = "ts") + offset(offset)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.7743     0.2274  -91.37  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(x)        2.875e-01    9  0.337  0.2697
## s(y)        2.338e-04    9  0.000  0.6303
## s(Depth)    3.391e+00    9 37.218 9.87e-10 ***
## s(DistToCAS) 1.307e-05    9  0.000  0.5246
## s(SST)      3.809e-05    9  0.000  0.8175
## s(EKE)      8.670e-01    9  5.582  0.0103 *
## s(NPP)      2.844e+00    9 23.233 9.15e-07 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```

## 
## R-sq.(adj) = 0.0993 Deviance explained = 40.8%
## -REML = 383.03 Scale est. = 1 n = 949

```

Try performing model selection as before from this base model and compare the resulting models.

Again saving space by doing this in one code chunk, I've numbered the order in which the terms were removed:

```

dsm_nb_x_y_ms <- dsm(count ~ s(x, bs="ts") + # 5
                        s(y, bs="ts") +
                        s(Depth, bs="ts"), # +
                        #s(DistToCAS, bs="ts") + # 1
                        #s(SST, bs="ts") + # 2
                        #s(EKE, bs="ts"), # + # 4
                        #s(NPP, bs="ts"), # 3
                        df_hn, segs, obs,
                        family=nb())
summary(dsm_nb_x_y_ms)

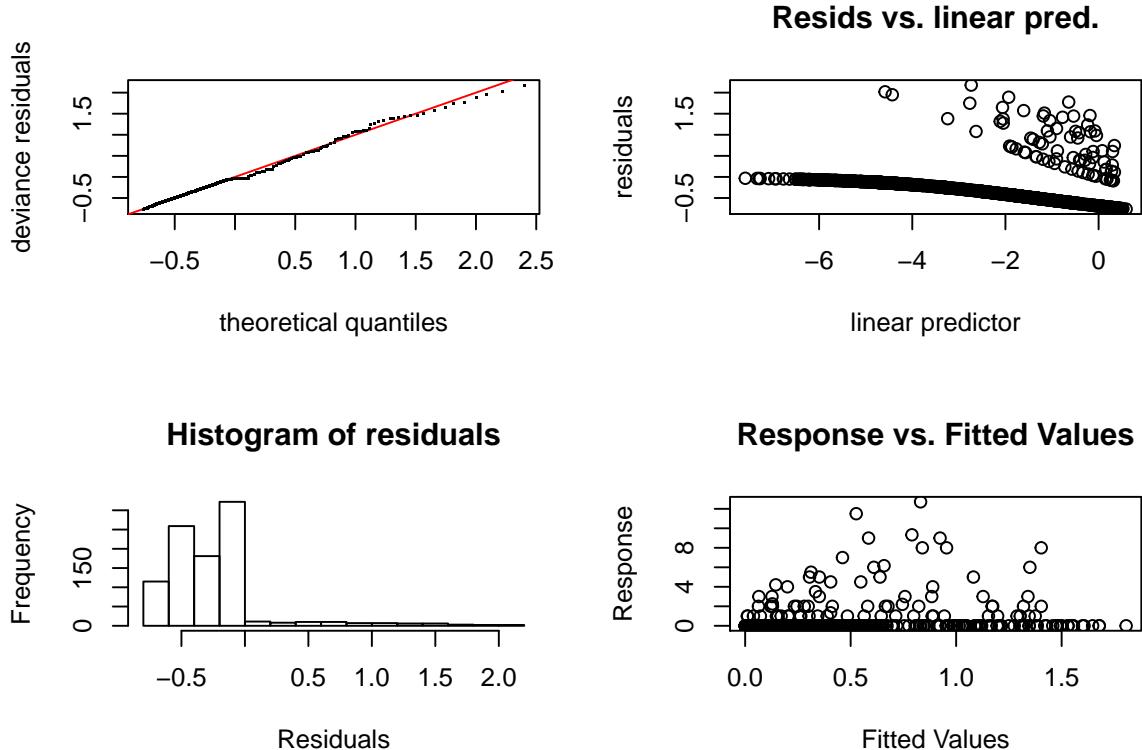
## 
## Family: Negative Binomial(0.098)
## Link function: log
##
## Formula:
## count ~ s(y, bs = "ts") + s(Depth, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -20.6381    0.2317   -89.08   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##             edf Ref.df Chi.sq p-value
## s(y)      1.036     9 17.63 1.54e-05 ***
## s(Depth)  3.850     9 52.74 2.05e-12 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) = 0.0486 Deviance explained = 35.9%
## -REML = 386.11 Scale est. = 1 n = 949

```

Here we actually remove the `x` term (as it's p -value was 0.00269, which was larger than the significance I'd decided on).

Looking at check plots for this model

```
gam.check(dsm_nb_x_y_ms)
```

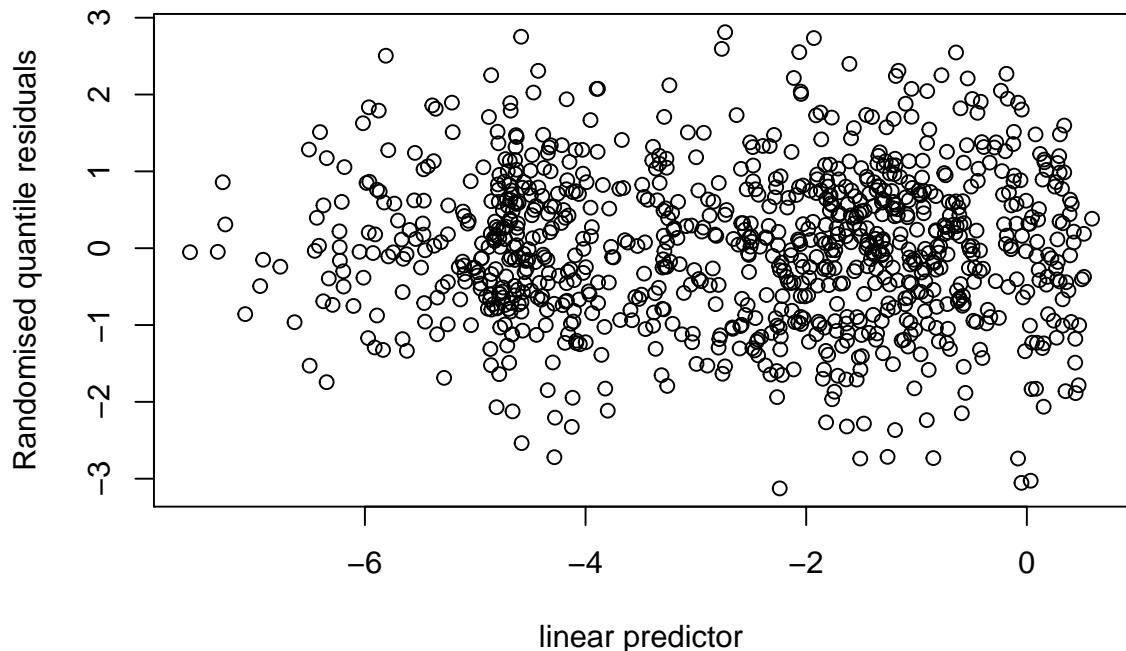


```

## 
## Method: REML   Optimizer: outer newton
## full convergence after 7 iterations.
## Gradient range [-3.687675e-06,5.114702e-06]
## (score 386.11 & scale 1).
## Hessian positive definite, eigenvalue range [0.2337323,33.5534] .
## Model rank = 19 / 19
##
## Basis dimension (k) checking results. Low p-value (k-index<1) may
## indicate that k is too low, especially if edf is close to k'.
##
##          k'    edf k-index p-value
## s(y)      9.00  1.04     0.64   0.005 **
## s(Depth)  9.00  3.85     0.65   0.020 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
rqgam.check(dsm_nb_x_y_ms)

```

Resids vs. linear pred.



All looks fine here.

Let's do the same with the Tweedie distribution for the sake of completeness.

```
dsm_tw_x_y_ms <- dsm(count ~ s(x, bs="ts") +
                         s(y, bs="ts") +
                         s(Depth, bs="ts") +
                         s(DistToCAS, bs="ts") + # 2
                         s(SST, bs="ts") + # 1
                         s(EKE, bs="ts") +
                         s(NPP, bs="ts"),
                         df_hn, segs, obs,
                         family=tw())
summary(dsm_tw_x_y_ms)

##
## Family: Tweedie(p=1.282)
## Link function: log
##
## Formula:
## count ~ s(Depth, bs = "ts") + s(NPP, bs = "ts") + offset(off.set)
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -20.6727    0.2392   -86.41  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```

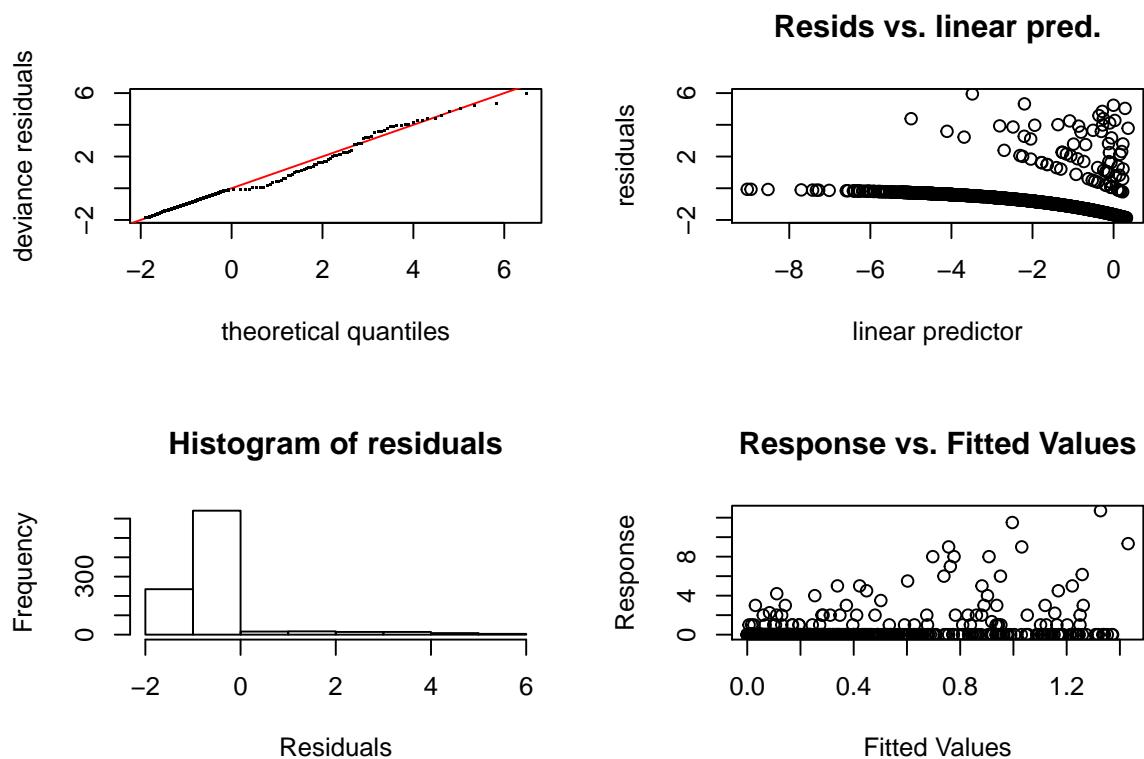
## 
## Approximate significance of smooth terms:
##          edf Ref.df    F  p-value
## s(Depth) 3.880     9 5.312 5.54e-11 ***
## s(NPP)   3.138     9 3.024 1.15e-06 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
## 
## R-sq.(adj) =  0.107  Deviance explained = 34.7%
## -REML = 387.86  Scale est. = 4.6361   n = 949

```

Now we pick rather different covariates in the model...

Looking at check plots for this model

```
gam.check(dsm_tw_x_ms)
```



```

## 
## Method: REML  Optimizer: outer newton
## full convergence after 7 iterations.
## Gradient range [-1.659687e-06,8.286128e-07]
## (score 387.8614 & scale 4.636114).
## Hessian positive definite, eigenvalue range [0.8391417,299.7395].
## Model rank = 19 / 19
## 
## Basis dimension (k) checking results. Low p-value (k-index<1) may
## indicate that k is too low, especially if edf is close to k'.
## 

```

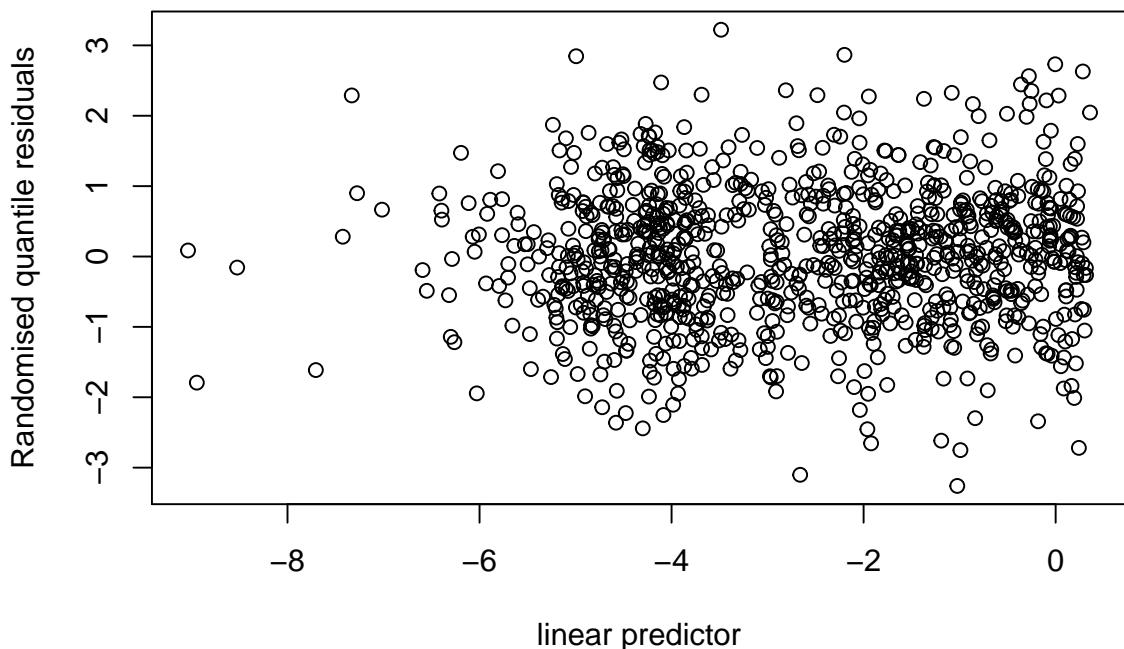
```

##          k' edf k-index p-value
## s(Depth) 9.00 3.88     0.80    0.34
## s(NPP)   9.00 3.14     0.77    0.07 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

rqgam.check(dsm_tw_xy_ms)

```

Resids vs. linear pred.



The plots look okay for this model, perhaps there is a slight issue with non-constant variance for the smaller values of the linear predictor.

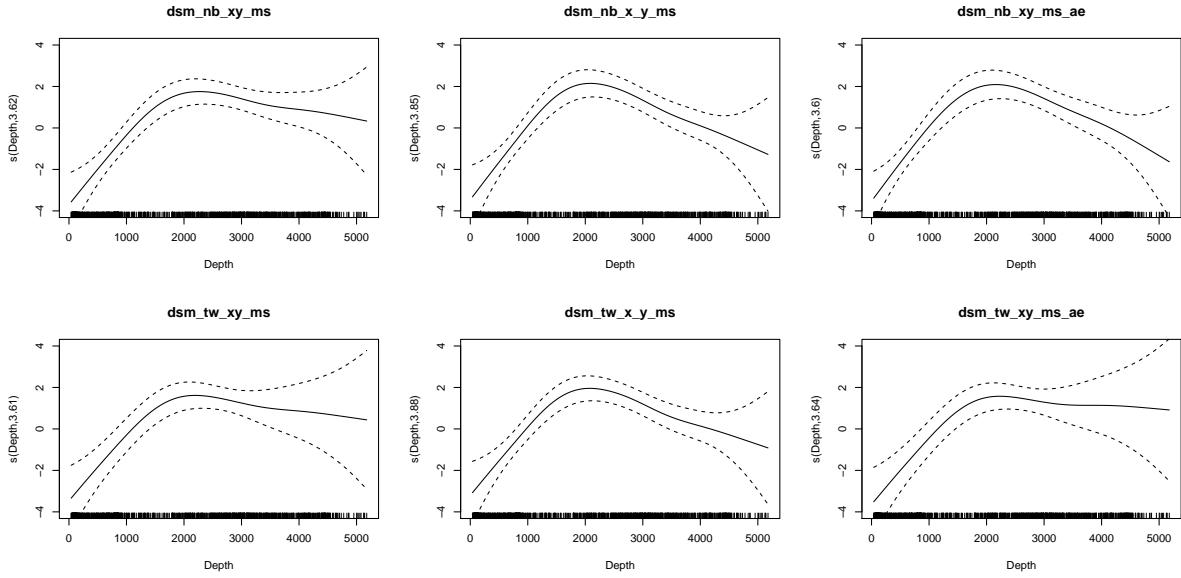
Compare the resulting smooths from like terms in the model, for example, if depth were selected in both models, compare EDFs and plots, e.g.:

Plotting all of the Depth terms, which are common to all models:

```

par(mfrow=c(2,3))
plot(dsm_nb_xy_ms, main="dsm_nb_xy_ms", select=2, scale=0, ylim=c(-4,4))
plot(dsm_nb_x_y_ms, main="dsm_nb_x_y_ms", select=2, scale=0, ylim=c(-4,4))
plot(dsm_nb_xy_ms_ae, main="dsm_nb_xy_ms_ae", select=2, scale=0, ylim=c(-4,4))
plot(dsm_tw_xy_ms, main="dsm_tw_xy_ms", select=2, scale=0, ylim=c(-4,4))
plot(dsm_tw_x_y_ms, main="dsm_tw_x_y_ms", select=1, scale=0, ylim=c(-4,4))
plot(dsm_tw_xy_ms_ae, main="dsm_tw_xy_ms_ae", select=2, scale=0, ylim=c(-4,4))

```



As we can see, these are pretty similar between models. In general showing a peak at around 2000m.

Note that the argument `select=` picks just one term to plot. These are in the order in which the terms occur in the `summary()` output (so you may well need to adjust the above code).

Comparing models

As with the detection functions in the earlier exercises, here is a quick function to generate model results tables with appropriate summary statistics:

```
summarize_dsm <- function(model){

  summ <- summary(model)

  data.frame(response = model$family$family,
              terms      = paste(rownames(summ$s.table), collapse=", "),
              AIC        = AIC(model),
              "Deviance_explained" = paste0(round(summ$dev.expl*100,2), "%")
            )

}
```

We can, again, make a list of the models and give that to the above function

Note here we're just going to compare the models that have count as the response:

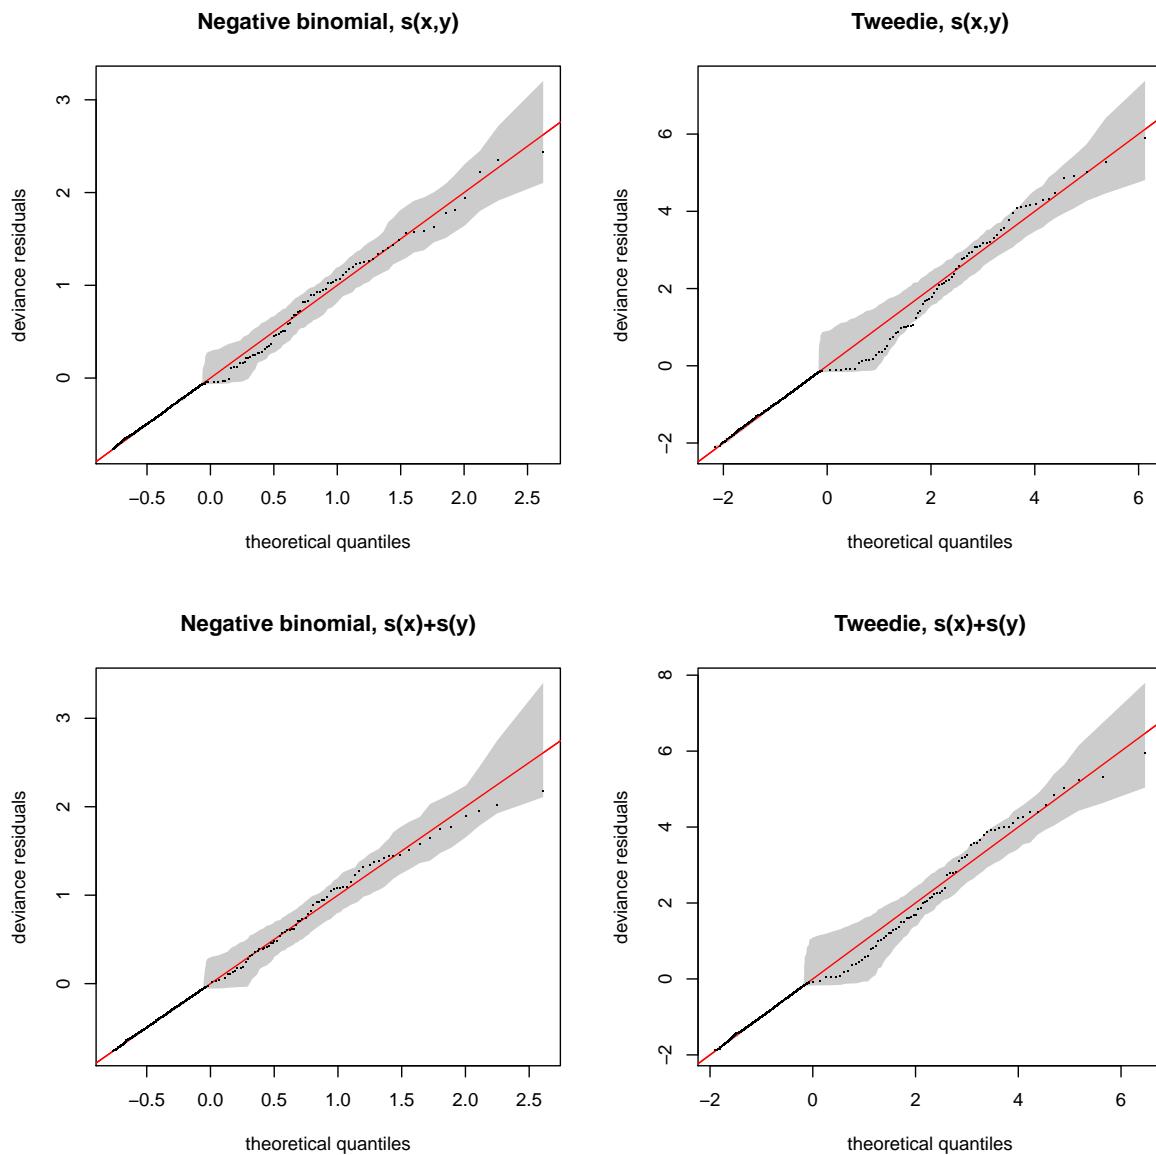
```
# add your models to this list!
model_list <- list(dsm_nb_xy, dsm_nb_x_y_ms, dsm_tw_xy_ms, dsm_tw_x_y_ms)
summary_table <- ldply(model_list, summarize_dsm)
row.names(summary_table) <- c("dsm_nb_xy", "dsm_nb_x_y_ms", "dsm_tw_xy_ms", "dsm_tw_x_y_ms")
```

```
summary_table <- summary_table[order(summary_table$AIC, decreasing=TRUE),]
kable(summary_table)
```

	response	terms	AIC	Deviance_explained
dsm_tw_xy_ms	Tweedie(p=1.268)	s(x,y), s(Depth)	1229.8875	37.96%
dsm_tw_x_y_ms	Tweedie(p=1.282)	s(Depth), s(NPP)	1227.3139	34.68%
dsm_nb_xy	Negative Binomial(0.105)	s(x,y)	775.2593	40.63%
dsm_nb_x_y_ms	Negative Binomial(0.098)	s(y), s(Depth)	762.5979	35.88%

We can also plot the Q-Q plots for these models to compare their distributional assumptions:

```
par(mfrow=c(2,2))
qq.gam(dsm_nb_xy_ms, main="Negative binomial, s(x,y)", rep=200)
qq.gam(dsm_tw_xy_ms, main="Tweedie, s(x,y)", rep=200)
qq.gam(dsm_nb_x_y_ms, main="Negative binomial, s(x)+s(y)", rep=200)
qq.gam(dsm_tw_x_y_ms, main="Tweedie, s(x)+s(y)", rep=200)
```



We'll save all four models below, but our selection were based on the deviance explained and the Q-Q plot would be for `dsm_nb_xy_ms`.

Saving models

Now save the models that you'd like to use to predict with later: I recommend saving as many models as you can so you can compare the results in the next practical.

```
# add your models here
save(dsm_nb_xy_ms, dsm_nb_x_y_ms, dsm_tw_xy_ms, dsm_tw_x_y_ms,
      file="dsms.RData")
```