Annex C: Analysis of 7 wind farm data sets

Alain F. Zuur
Highland Statistics Ltd.
Newburgh, UK.
www.highstat.com

1.1 Introduction

Required knowledge for this text is data exploration, multiple linear regression, and generalized linear modelling (e.g., Zuur *et al*. 2010, 2012, 2014).

In Zuur *et al.* (2014) two detailed statistical analyses on wind farm data sets are presented. Both data sets consisted of data sampled before construction of the wind farm (T_0 data), and after building the wind farm (T_1 data). Advanced statistical techniques like zero inflated generalised additive mixed models (GAMM) with 2-dimensional spatial smoothers were applied.
In this report, data on Common Guillemot / Razorbills ("razormots" *Uria aalge / Alca torda*) from seven European offshore wind farms (OWEZ, PAWP, Horns Rev I, Horns Rev II, Apha Ventus, Blighbank, Robin Ridge) are analysed. The sizes of the wind farms and the number of turbines per wind farm differ. The density of turbines is defined as the number of turbines divided by the size (in km^2) of the wind park, and for the seven wind park parks used in the current analyses we have the following densities:

	Density	Wind farm
1	0.92	OWEZ
2	2.77	PAWP
3	2.89	HornsII
4	3.07	AlphaVentus
5	3.26	RobinRidge
6	4.07	HornsI
7	4.16	Blighbank

In this report we will investigate whether there is a turbine density effect on Common Guillemots and Razorbills. As compared to Zuur *et al*. (2014), a different statistical analysis is applied; we only use data sampled during the post-construction periods (whereas in Zuur *et al*. (2014) T_0 and T_1 data was used), since it is not always clear where the T_0 data stops and the T_1 period starts.

1.2 Analysis approach

In the next five sub-sections we discuss the details of the models that were applied on the post-construction data (T_1).

1.2.1 Distribution

The response variable is the density of Guillemots. Density is defined as observed numbers divided by survey area. Using a generalized linear model (GLM) allows us to model the number of birds with a Poisson, negative binomial or zero inflated distribution, while using (the log of) survey area as an offset variable (Zuur *et al*. 2007). This means that we assume a linear relationship between sampling effort and expected number of birds.
We will consider the following three statistical distributions for the number of birds:

- Poisson distribution.

- Zero inflated Poisson distribution.
- Negative binomial distribution.

Density of bird is equal to the number of birds divided by survey area. This allows us to model the number of birds using a GLM for count data with the log of the survey area as an offset variable. This approach assumes a linear relationship between sampling effort and expected number of birds.

1.2.2 Covariates

In all GLMs used in this report, the expected values of birds are modelled with a log link. For example, for the GLM with a Poisson distribution we use:

$$\begin{aligned} Birds_i &\sim Poisson(\mu_i) \\ E(Birds_i) &= \mu_i \\ \log(\mu_i) &= \text{Covariates}_i + \log(\text{Survey area}_i) \end{aligned}$$

Hence, there is an exponential relationship between the expected number of birds and the covariates, ensuring that the fitted values are positive. Similar expressions are used for the GLMs with a zero inflated Poisson (ZIP) or negative binomial (NB) distribution, see Zuur *et al*. (2012) for details.
We will consider the following three models in terms of the covariates:

- Covariates = Distance effect
- Covariates = Distance effect plus Year effect
- Covariates = Distance effect plus Year effect plus an interaction between distance and Year (the distance effect changes per year)

The term 'Distance' stands for distance of sampling location (bird count) to the wind farm.

We will use distance and year as covariates.

1.2.3 Correlation; approach 1

The GLMs that we introduced in the previous subsection do not take into account spatial and/or temporal correlation. One approach to include spatial correlation into these models is by adding a residual term ε_i to the predictor function (the predictor function is the term on the right hand side of the 'log(μ_i) ='), and allow it to be spatially correlated. Such a model is given by:

$$\begin{aligned} Birds_i &\sim Poisson(\mu_i) \\ E(Birds_i) &= \mu_i \\ \log(\mu_i) &= \text{Covariates}_i + \log(\text{Survey area}_i) + \varepsilon_i \end{aligned}$$

We considered the following three types of correlation structure:

- Spatial correlation between all sampled observations. This type of correlation was used in Zuur *et al*. (2014). It allows for spatial correlation between observation *i* in year *k* and observation *j* in year *l*, even if the two observations were taken in year *k* = 2000 and year *l* = 2010.
- Spatial correlation between all sampled observations from the same survey. We consider the spatial correlation from different surveys as *independent* realisations. Hence, we only allow for spatial correlation between observations from the same survey.
- Spatial correlation that changes over time (interaction between spatial and temporal correlation).

Statistical details and R implementation are discussed in Blangiardo *et al*. (2013). The package INLA (Rue *et al*. 2014) allows one to fit these models from within R (R Core Team 2013).
The full model specification is then given by:

$$
\begin{aligned}
& Birds_i \sim Poisson(\mu_i) \\
& E\left(Birds_i\right) = \mu_i \\
& \log\left(\mu_i\right) = Intercept + \beta_1 \times \text{Distance}_i + \log(\text{Survey area}_i) + \varepsilon_i \\
& \varepsilon_i \text{ spatially(-temporal) correlated noise}
\end{aligned}
$$

As discussed earlier, it is also an option to add year as a categorical covariate (and the interaction between distance and year).
When we ran these models on data from each wind farm we noticed that in most of the models the parameter β_1 (the slope for the covariate Distance) was not significant, indicating that there is no distance effect.
We then modelled distance as a categorical covariate, and also as a binary covariate (in- or not inside the wind park). In only a few models we obtained a significant distance effect. Note that this type of trial and error modelling has a certain fata phishing element.
A more detailed data exploration and initial modelling results showed a non-linear distance effect. We therefore used models of the form:

$$
\begin{aligned}
& Birds_i \sim Poisson(\mu_i) \\
& E\left(Birds_i\right) = \mu_i \\
& \log\left(\mu_i\right) = f(\text{Distance}_i) + \log(\text{Survey area}_i) + \varepsilon_i
\end{aligned}
$$

where the notation *f*(Distance) stands for smoothing function. Hence, this model allows for a non-linear distance effect (though keep in mind that the model is already non-linear due to the log-link function). INLA allows for Poisson, ZIP and NB GAMs with spatial, and spatial-temporal correlation, but we ended up with rather non-smooth smoothers. We also encountered various numerical optimisation errors. We therefore programmed a low rank thin plate regression spline, and also an O'Sullivan spline (these are more advanced smoothers as compared to the available smoothers in INLA) and used these in INLA. See Zuur *et al*. (2014) for examples and R code. INLA produced rather large confidence intervals for the smoothers and computing time was in the order of 24 hours per model on a modern computer. The main problem is the large data size; some wind farms contained 50,000 observations. For smaller data sets (< 5,000 observations) we did not encounter major problems.
The table below shows the number of observations per wind park.

```
AlphaVentus      49086
Blighbank         1238
HornsRevI         8590
HornsRevII        6247
OWEZ              6571
PAWP              5299
RobinRigg         9948
```

Although the software package INLA (which can be executed from within R) allows one to fit GAMs with spatial and/or temporal correlation structures, the tools for smoothing functions, in combination with the large data sets, means that we end up with excessive computing time and poor results. More time and research is needed in order to run INLA on such large data sets.

1.2.4 Correlation; approach 2

Instead of modelling the spatial correlation with a spatially correlated residual term, we can use a 2-dimensional smoothing function of the spatial coordinates (Xkm and Ykm). And we can also use Survey as a random intercept. This results in models of the form:

$$
\begin{aligned}
&Birds_i \sim Poisson(\mu_i) \\
&E(Birds_i) = \mu_i \\
&\log(\mu_i) = f(\text{Distance}_i) + \log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey
\end{aligned}
$$

The inclusion of the 2-dimensional smoother *f*(Xkm, Ykm) and the random intercept Survey is a quick and dirty way to capture the spatial correlation. Theoretically, its form should be similar to the residual spatial correlation estimated by INLA. The advantage of this approach is that we can use the `gamm4` function from the `gamm4` package (Wood 2006) to fit this model. The disadvantage is that this package does not have facilities to fit a ZIP distribution or a NB distribution with random effects.

We will use a 2-dimensional smoother *f*(Xkm, Ykm) and a random intercept to capture the spatial correlation.

1.2.5 Dealing with overdispersion

The model presented in Subsection 1.2.5 is in fact a generalized additive mixed effects model (GAMM) with a Poisson distribution, two smoothers and a random intercept. Once a Poisson GAMM has been fitted we need to check the dispersion parameter. Its value should ideally be 1, with values larger than 1 indicating overdispersion and values smaller than 1 underdispersion. If overdispersion (or underdispersion) occurs, then we need to figure out why this happens and solve the problem. Likely causes of overdispersion are zero inflation, correlation, or large variance, among many other possible causes. The wind farm data has a large number of zeros. However, one should not immediately apply zero inflated models. It is well possible that a Poisson GLM or Poisson GAMM can be used to analyse data with many zeros. Also, models that contain a zero inflation component and a spatial correlation term may encounter numerical estimation problems as both components may be fighting for the zeros. A negative binomial distribution allows for more variation than a Poisson distribution, but this mechanism may also capture the excess number of zeros. And a smoother may also be able to model the zeros. Hence, we have five components that could potentially model the large number of zeros; the smoothing function *f*(Distance), the 2-dimensional smoother *f*(Xkm, Ykm), a spatial correlation term, the zero inflation component in a ZIP model, and the negative binomial distribution. Suppose that we have lots of observations sampled close to each other, with lots of zero counts. This may either be considered as spatial correlation, zero inflation, or large variance. Or a covariate may explain the zero counts.

The problem of zero inflation can be dealt with in at least 5 different ways. It is unwise to fit a model that contains all approaches (e.g. a zero inflated negative binomial GAMM with spatial correlation and spatial smoothers). It is better to fit a model that only contains 1 or 2 approaches. In our GAMMs the 2-dimensional smoother *f*(Xkm, Ykm) can potentially model the zeros, and the same holds for the *f*(Distance) smoother.

If it turns out that the proposed Poisson GAMM is still overdispersed, then we will add an observation level random intercept (Elston *et al*. 2000). This is an extra latent variable that scopes up any extra variation not explained by the other covariates in the model.

1.3. Setup of the analyses

The total number of observations for the seven wind farms is around 100,000, which makes computing time rather long. We therefore analyse the data for each wind park separately.

1.4 Data exploration and model validation

Prior to the analysis of the data, a data exploration following the protocol described in Zuur *et al*. (2010) is applied. Once models have been fitted, model validation is applied to inspect the residuals for any non-linear patterns.

2 Results

We present the results for each wind park. The order of the results per wind park is based on the order of the turbine density per park. To compare like-with-like we only use the data sampled up to 12 kilometer from the wind parks.

2.1 Results for OWEZ

The data set for this park contains around 6,500 observations. Sampling took place from 2002 to 2012, though not every year was sampled. Figure 1 shows a so-called Cleveland dotplot of the birds sampled at OWEZ. In this graph the number of birds are plotted along the *x*-axis and the row number (as imported from the data file) is along the *y*-axis. A Cleveland dotplot allows one to check the data for outliers. In this case we can see the large number of zeros (all the dots on the left), and there are only a few observations of relative large numbers. In our experience, when the majority of the observations are between 0 and 25-ish, a Poisson distribution tends to work well. If the majority of the observations are considerably larger than 25-ish, we tend to end up with a negative binomial distribution.

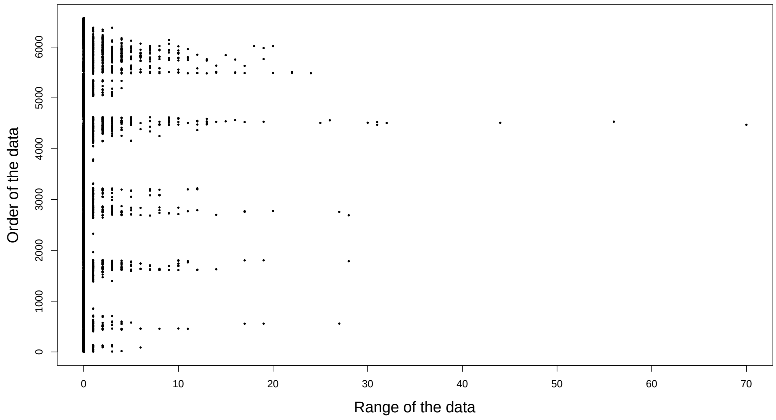

Figure 1. Cleveland dotplot of the number of birds sampled at OWEZ.

We also made scatterplots of distance (in kilometers) versus bird density; see Figure 2. Note the differences in patterns between the years.

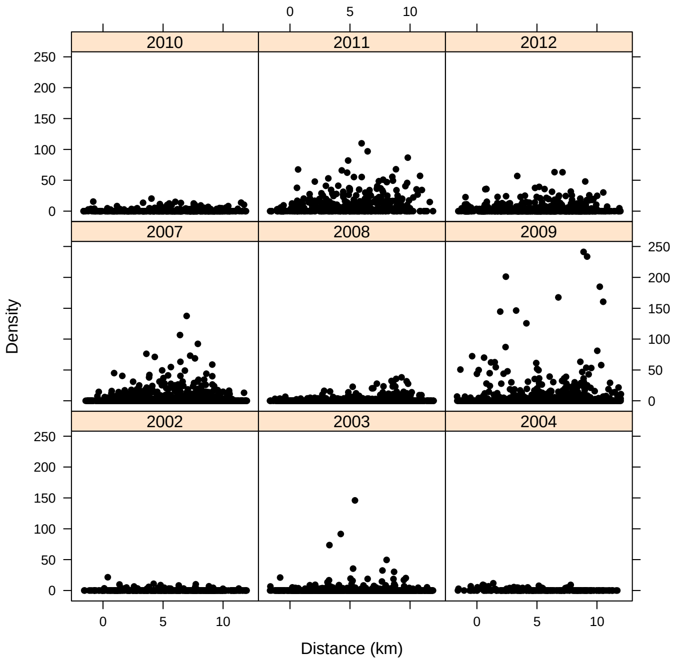

Figure 2. Scatterplot of bird density versus distance for the OWEZ data.

We first applied a GAMM of the form:

$$Birds_i \sim Poisson(\mu_i)$$
$$E(Birds_i) = \mu_i$$
$$\log(\mu_i) = f(\text{Distance}_i) + \log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey$$

The smoothing function *f*(Distance) is presented in Figure 3. Note the wide 95% point-wise confidence bands for the smoother; this indicates that the smoother is not significantly different from 0. This is confirmed by the numerical output of the model (not presented here). The model can be rewritten as:

$$Birds_i \sim Poisson(\mu_i)$$
$$E(Birds_i) = \mu_i$$
$$\mu_i = e^{f(\text{Distance}_i) + \log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey}$$
$$= e^{f(\text{Distance}_i)} \times e^{\log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey}$$

Hence, a non-significant smoother *f*(Distance) in the GAMM means that

$$\exp(f(\text{Distance})) \approx \exp(0) \approx 1.$$

Therefore, a non-significant smoother *f*(Distance), as in Figure 3, means that we can state that expected numbers of birds do not increase, or decrease, when we move away from the wind farm.

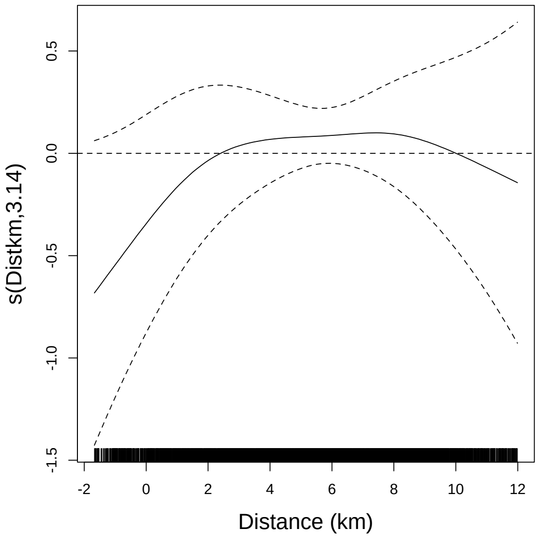

Figure 3. Results for OWEZ data. One distance smoother was used.

The GAMM assumes that the distance effect is the same in each year. This may be a plausible assumption if the data has been sampled in only 1 or 2 sequential years, but sampling at OWEZ took place between 2002 and 2012. The model can easily be extended to allow for a different distance effect per year. Such a GAMM is specified below.

$$
\begin{aligned}
& Birds_i \sim Poisson(\mu_i) \\
& E\left(Birds_i\right) = \mu_i \\
& \mu_i = e^{f_k(\text{Distance}_i) + \log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey} \\
& \quad = e^{f_k(\text{Distance}_i)} \times e^{\log(\text{Survey area}_i) + f(Xkm_i, Ykm_i) + Survey}
\end{aligned}
$$

Note the subscript *k* for the f_k(Distance) smoother. We now have one smoother for each year *k*. The estimated smoothers are presented in Figure 4. The AIC indicated that the model with 9 smoothers is better than the model with 1 smoother.

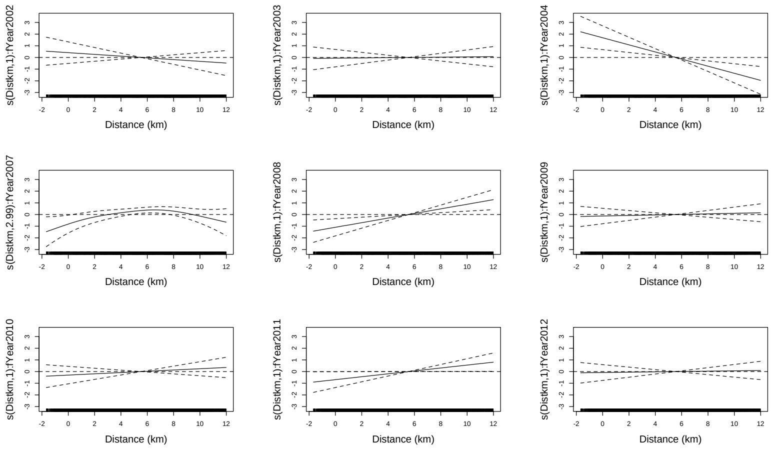

Figure 4. Smoothers for each year for the OWEZ data.

The numerical output for the GAMM with 9 smoothers is given below.

Parametric coefficients:

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-2.7713	0.6448	-4.298	1.73e-05

Approximate significance of smooth terms:

	edf	Ref.df	Chi.sq	p-value
s(Distkm):fYear2002	1.00	1.00	0.799	0.371319
s(Distkm):fYear2003	1.00	1.00	0.025	0.874036
s(Distkm):fYear2004	**1.00**	**1.00**	**10.964**	**0.000929**
s(Distkm):fYear2007	**2.99**	**2.99**	**13.310**	**0.003994**
s(Distkm):fYear2008	**1.00**	**1.00**	**8.847**	**0.002936**
s(Distkm):fYear2009	1.00	1.00	0.151	0.697538
s(Distkm):fYear2010	1.00	1.00	0.650	0.420167
s(Distkm):fYear2011	1.00	1.00	4.193	0.040586
s(Distkm):fYear2012	1.00	1.00	0.057	0.811017
s(Xkm,Ykm)	15.37	15.37	93.279	4.12e-13

Note that the distance smoother is only significant for the years 2004, 2007 and 2008. The number under edf is the degrees of freedom of a smoother. A value of 1 means a straight line and the larger the value, the more non-linear is a smoother. The optimal edf is estimated using a process called cross-validation. By the way, the flexibility of smoothers to estimate the optimal degrees of freedom is yet another way how a GAMM can fit excessive number of zeros.

In all years, except for 2007, the distance effect is linear. Let us zoom in on the smoother for 2004, 2007 and 2008; see Figure 5. When the smoother *f*(Distance) is negative, the exp(*f*(Distance)) term is smaller than one, which implies a decrease in expected number of birds. If the smoother is larger than 0, there is an increase.

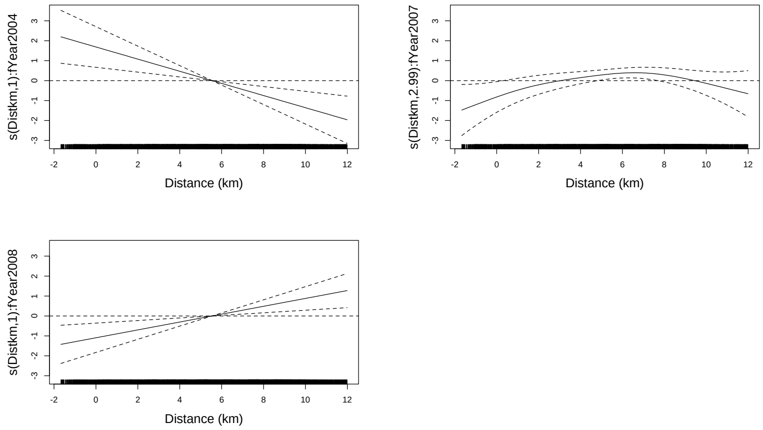

Figure 5. The smoothers for 2004, 2007 and 2008 for OWEZ. Smoothers are centered around 0.

To summarise the results of the GAMM with 9 smoothers, there is a distance effect but only in 3 years. In 2004 there is a negative effect; the further away from the wind farm the fewer (!) birds. In 2007 there is a non-linear effect, but confidence bands around the distance smoother are such that the distance effect is up to about 1 km. In 2008 there is a positive effect.

2.2 Results for PAWP

The second smallest wind park is PAWP. We applied a GAMM with one distance smoother, and a GAMM with a distance smoother per year and compared the two models using the AIC. The AIC indicated that the model with one distance smoother for all years is better. The estimated smoother is presented in Figure 6.

Note the linear shape of the smoother to about 3 km. In this distance range the smoother is negative. That means that the further we more away from the PAWP wind farm, the more birds we sample, but from 3 km onwards this effect plateaus.

At the distance of -2 km (this is at the centre of the wind park) the value of the smoother is around -1. The value of exp(-1) is around 0.36. This means that from 3 km to -2 km there is a decrease of 74% in bird numbers.

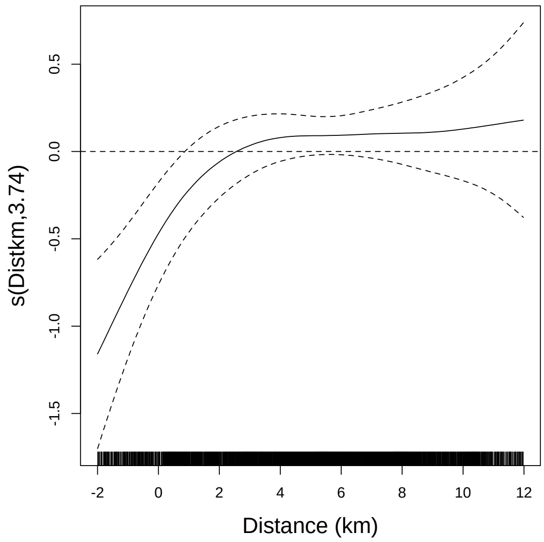

Figure 6. Smoother for distance for the PAWP data.

For the PAWP data there is a distance effect up to about 3 km distance from the wind park. At the centre of the wind park there is a 74% decrease in abundance as compared to the 3^{+} km values.

2.3 Results for HornsRevII

The third wind farm in terms of turbine density gave a non-significant distance effect. The smoother is presented in Figure 7. Although the smoother is not significant, it is interesting to note that its shape again indicates a plateau pattern. It may be an option to add seasonal information in order to reduce the width of the confidence bands. If the confidence bands would be smaller, then the interpretation would be identical to the PAWP data (though the distance where it reaches the plateau is slightly further away).

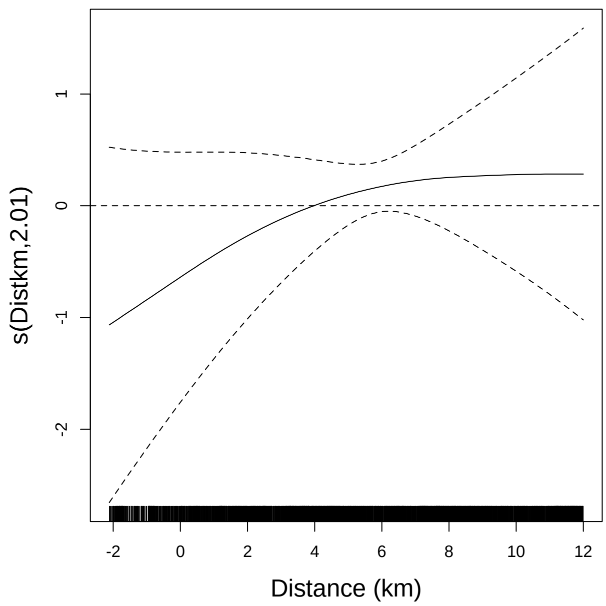

Figure 7. Smoother for the HornsRevII data.

The shape of the HornsRevII distance smoother looks similar to that of the PAWP data, but unfortunately the confidence bands are wider. Further research may result in a better model.

2.4 Results for AlphaVentus

This data set contains more than 50,000 observations and we encountered various numerical estimation problems with the GAMM. We fitted a model with one smoother for distance, and also a model with 3 smoothers (sampling took place in 3 years, though in the third year only November was sampled). The AIC indicated that a model with 3 distance smoothers is better. The three estimated smoothers are presented in Figure 8. The results for 2010 and 2011 show that there is a negative wind farm effect up to about 4 km. As compared to the previous wind farms, the distance effect is stronger. Inside the wind park the value of the smoother is around -3 and -2 for 2010 and 2011, respectively. That is a 90% reduction! Also note that the smoother does not plateau. Instead, observed numbers increase for larger distance.

There is a small amount of overdispersion in these models that is not accounted for yet. Hence, further model improvement is needed.

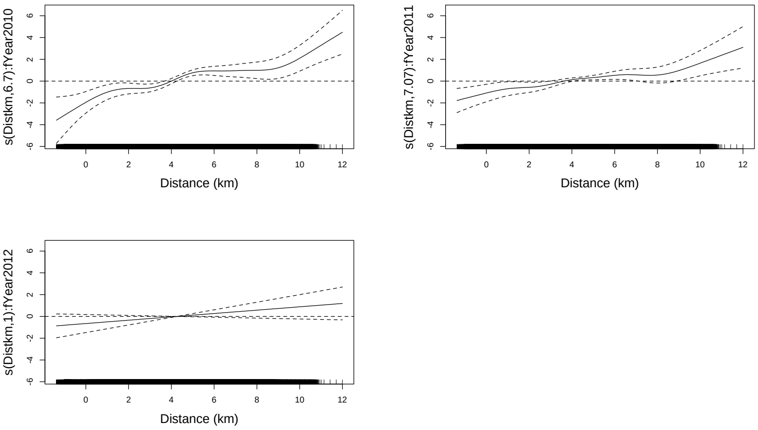

Figure 8. Distance smoothers for the AlphaVentus data. It may me an option to rerun the model without the 2012 data as sampling in this year took only place in November.

For the AlphaVentus data there is a distance effect up to about 4 km distance from the wind park. At the centre of the wind park there is a 90% decrease in abundance as compared to the 4^+ km values. We recommend rerunning the models without the 2012 data.

2.5 Results for RobinRidge

For the RobinRidge data we have around 10,000 observations from 4 years. In 2013, we only have January and February data. In 2011 and 2012 all months were sampled, and in 2010 sampling took place from March onwards.

The estimated distance smoothers are presented in Figure 9. The estimated degrees of freedom for all smoothers is 1 in 2010, 2011 and 2013, indicating that we have straight lines (on the predictor scale) in these years. In 2012 the distance effect is slightly non-linear. Note that the distance effect is significant in all years.

The distance effect is negative up to around 5 km, and inside the wind park the value of the smoother is around -5, which means considerable lower numbers as compared to observations made at 5 km distance (exp(-5)) to be precise).

The optimization routines gave some warning messages for the optimal model, hence we recommend rerunning the model without the 2013 data.

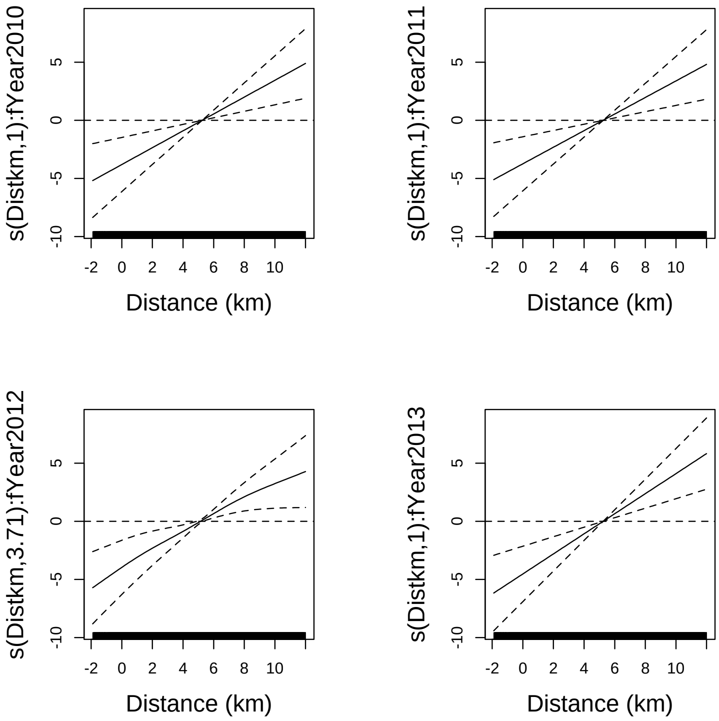

Figure 9. Distance smoothers for the RobinRidge data.

2.6 Results for HornsRevI

The distance smoother for the HornsRevI data is presented in Figure 9. Again, there is a decrease in number of birds close to the wind farm. At about 5 km the expected numbers of birds are approximately a factor exp(0.4) ≈ 1.50, which is 50%, higher.

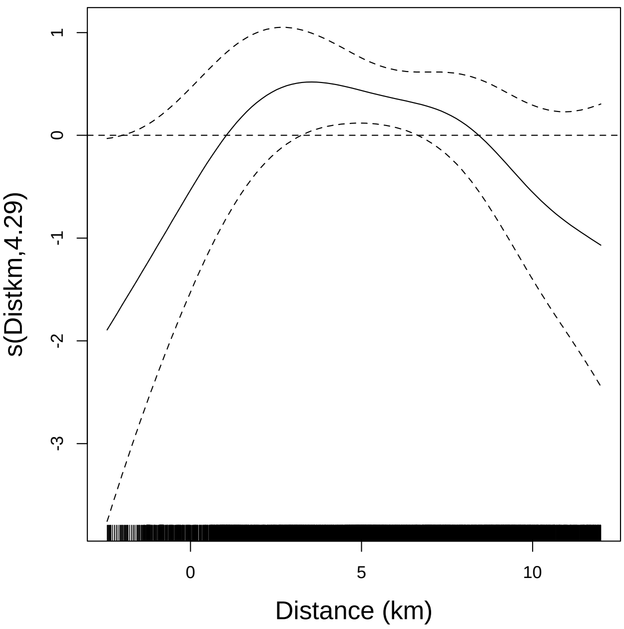

Figure 10. Distance smoother for the HornsRevI data.

2.6 Results for Blighbank

For the Blighbank we have 1238 observations, made in 2010, 2011 and 2012. For the 2010 data we only have autumn data. The AIC indicated that the model with 3 smoothers is the best. The estimated smoothers are presented in Figure 11. Only the smoother for 2010 is significant. Up to about 4 km the numbers are lower. At the centre of the wind park there is a reduction of 70%.

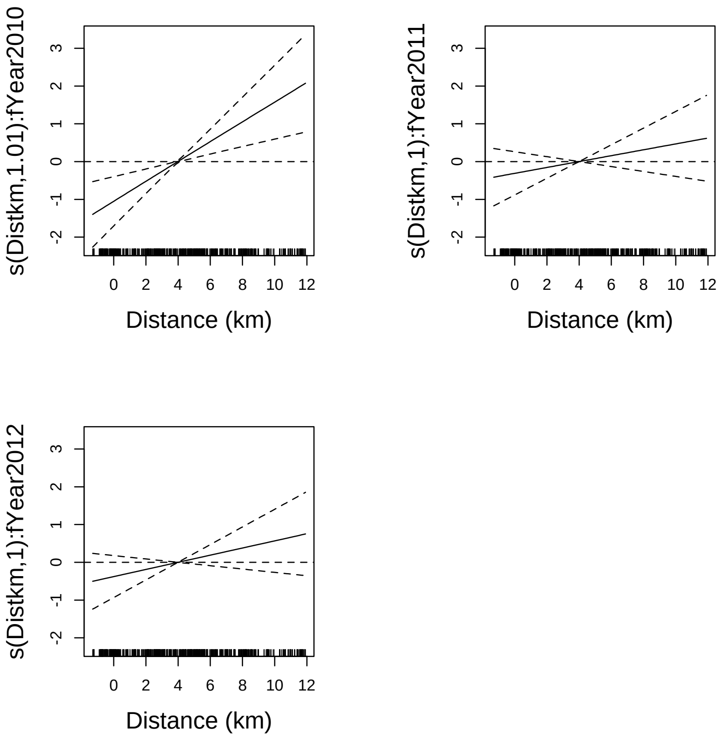

Figure 11. Smoothers for the Blighbank data.

3. Discussion and conclusions

In this report we showed the results of various different statistical analyses applied on the wind farm data. Due to the large number of observations and the limited amount of time available for this project, we settled for a Poisson GAMM using a 2-dimensional spatial smoother and random intercepts to capture the spatial correlation. Without doubt, the INLA approach with a residual spatial-temporal correlation is the most statistically advanced, and preferred approach. It allows one to fit Poisson, zero inflated and negative binomial GAMs, and compare these three distributions. As long as only parametric terms are involved, computing time is not an issue. However, applying GAMMs on data sets with more than 10,000 observations is a frustrating process in INLA. Quite often we ran the code for 8 hours only to end up with an error message related to numerical optimization problems. These error messages can be avoided by tweaking the value of a specific INLA parameter (e.g. the value of a prior, or change the distribution of a prior). But this requires a time-consuming trial and error process. However, the INLA models are certainly the way to go.

The results presented in this report are based on Poisson GAMMs. We dealt with the spatial correlation in a pragmatic way (spatial smoother), and the same holds for the overdispersion that was present in some of the models (i.e. we used an observation level random intercept). Both solutions are accepted tools in the literature, though we would label them as cumbersome from a purist statistical point of view.

Some of the models can be improved by dropping data for certain years (e.g. years with only 1 or 2 months of data), and results may change accordingly.

So what do the analyses tell us? For most wind parks the models indicate that (i) expected number of birds are lower inside the parks, and (ii) the distance effect is up to about 4 km (for some parks there was a plateau effect). By taking the exponential of the distance smoother we can quantify the reduction. For some wind farms the reduction was up to 70%. There is some indication that wind farms with higher turbine density have a large reduction. Additionally, for some of the larger wind parks the distance to which there was a distance effect was further away as compared to the smaller parks. However, these patterns were not consistent, and we doubt whether they can be used for extrapolation. For example, Figure 12 shows a scatterplot of turbine density versus the distance at which the smoother equals 0 (we called this the plateau point). When the smoother of distance equals 0, exp(0) = 1 and there is no effect of the distance smoother on expected numbers. The figure does not show a clear pattern.

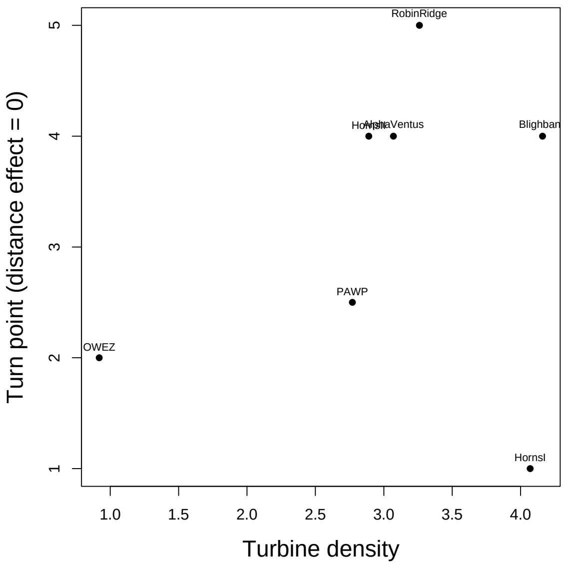

Figure 12. Turbine density versus the distance value at which the distance smoother = 0 (and therefore exp(smoother) = 1). We called this the plateau point, but not all smoothers had a plateau shape. Results based on non-significant smoothers are included.

It may be an option to refit the models using MCMC in JAGS, and estimate the distance values at which the distance effect plateaus. Such an approach would allow us to obtain a 95% credible interval for this specific distance value. Alternatively, we can fit a GLM with a breakpoint, and try to estimate the optimal breakpoint value for distance (like an elbow effect).
It may be an option to investigate the relationship between turbine density and the plateau distance. And it may also be an option to look at the relationship between turbine density and the reduction in expected number of birds the closer one gets to a wind farm.
Another issue that we need to investigate is the effect size. How much variation is explained by the distance smoother?

A pragmatic statistical analysis revealed to main findings:

- For some wind farms there is a reduction in expected number of birds the closer one gets to a wind farm.
- For some wind farms the distance effect plateaus at about 3 – 5 km. It is unclear whether this effect is related to turbine density.

Care is needed with the results of the models presented in this report and further research is needed.

We finally present a table showing the turbine density per park, the distance at which the *f*(Distance) value equals 0 (called Turnpoint), and the percentage of change between the centre and the point where *f*(Distance) = 0. As mentioned before, these results should be used with care. The density column gives the turbine density per park. Turnpoint is the distance where *f*(Distance) = 0. We read this from the graphs, and in case we have multiple distance smoothers per wind park, we took an average. Note that results for non-significant smoothers are given as well. Reduction is the value of 1 – exp(*f*(Distance)), calculated at the center of a wind farm, and represents the reduction in expected numbers of birds. Values are read of the graph, results for non-significant smoothers are included, and no measure of uncertainty is included. Formulated differently: Use with great care!

	Density	Names	Turnpoint	Reduction
1	0.92	OWEZ	2.0	0.45
2	2.77	PAWP	2.5	0.67
4	2.89	HornsII	4.0	0.63
5	3.07	AlphaVentus	4.0	0.95
7	3.26	RobinRidge	5.0	0.95
3	4.07	HornsI	1.0	0.86
6	4.16	Blighbank	4.0	0.53

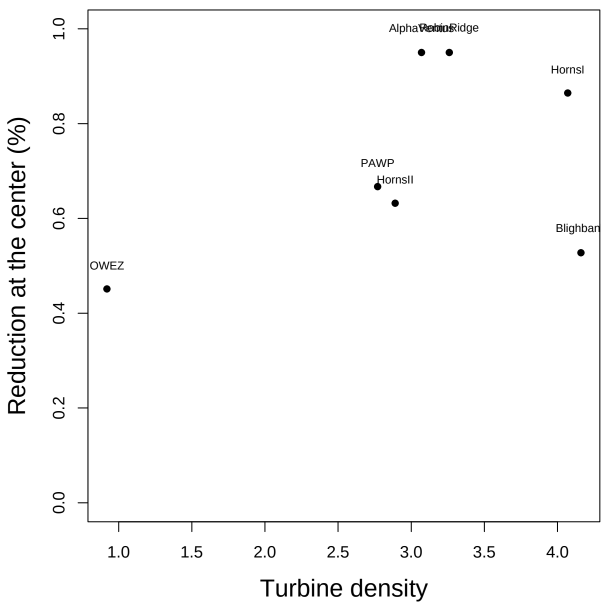

Figure 13. Reduction at the centre versus turbine density. Results based on non-significant smoothers are included.

References

Blangiardo M, Cameletti M, Baio G, Rue H (2013). Spatial and spatio-temporal models with R-INLA. Spatial and Spatio-temporal Epidemiology 7: 39–55

Elston DA, Moss R, Boulinier T, Arrowsmith C, Lambin X (2001) Analysis of aggregation, a worked example: numbers of ticks on red grouse chicks. Parasitology 122: 563-569

R Core Team (2013). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/

Rue H, Martino S, Lindgren F, Simpson D, Riebler A (2013) INLA: Functions which allow to perform full Bayesian analysis of latent Gaussian models using Integrated Nested Laplace Approximation. R package version 0.0-1379661604.

Wood SN (2006) Generalized Additive Models: An Introduction with R. Chapman & Hall/CRC: London

Zuur AF, Ieno EN, Smith GM (2007) Analysing Ecological Data. Spring-er. New York

Zuur AF, Ieno EN, Elphick CS (2010) A protocol for data exploration to avoid common statistical problems. Methods Ecology and Evolution 1: 3-14

Zuur AF, Saveliev AA, Ieno EN (2014) A Beginner's Guide to Generalized Additive Mixed Models with R. Highland Statistics Ltd.