FACULTY OF SCIENCE
UNIVERSITY OF COPENHAGEN

Storm surges of the West Coast

A statistical model of storm surges on the Danish West Coast

Student	Morten Hallas
Supervisor	Aslak Grindsted
Date of submission	20/12/2016

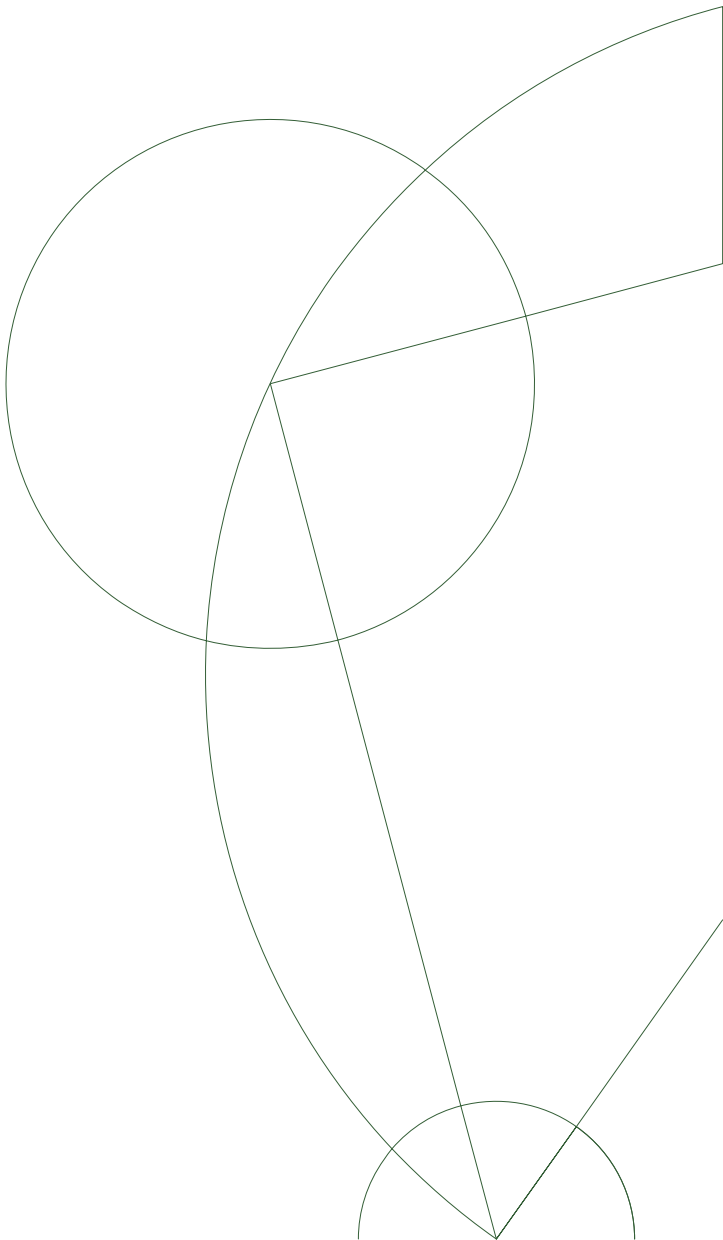

Abstract

Motivated by the mean sea level rise as described in IPCC [2013] and the declining trend in storm surges for Eastern England as described in Menéndez and Woodworth [2010], this paper aims to examine the distribution of storm surges on the West Coast of Denmark and test for non-stationary behavior. Only data from the stations facing the Wadden Sea has been acquired.
The peak over threshold method is used, the peaks modelled as a Generalized Pareto distribution and the rate of occurrences as a Poisson distribution. Through teleconnection analysis a connection is found with the 3rd EOF of the SST over the Northern Hemisphere, which Osborn [2010] identifies as a proxy for the NAO, although this is questioned at the end of the paper.
A range of different models, stationary and non-stationary, and predictors, such as NAO and the Global Temperature Anomaly, are fitted with Markov Chain Monte Carlo simulations and their performance intercompared with Bayes factors.
No evidence is found in favor of a non-stationary model for the Generalized Pareto distribution, but there is very strong evidence that the rate parameter of the Poisson distribution is dependent on the predictor identified in the teleconnection analysis.
The coefficient of this dependency is found to be negative, and by extrapolation with the CMIP5 ensemble, the rate of storm surges is found to be declining for all of the four emission scenarios, RCP2.6, 4.5, 6 and 8.5.

Abbreviations

AIC	Akaike Information Criterion, a measurement of the relative quality of statistical models
CI	Confidence interval. If nothing else is noted then it is set to 66%.
CMIP5	Coupled Model Intercomparisson Project 5. A framework for climate models allowing them to be used as an ensemble.
DIC	Deviance information Criterion. A generalization of AIC.
EOF	Empirical Orthogonal Function. Can both be referring to the method and the vectors. Introduced in Section 2.3.
GPD	Generalized Pareto Distribution. Introduced in Section 2.1.
KDI	Kystdirektoratet. The Danish coastal protection agency.
GDP	Gross Domestic Product.
GEV	Generalized Extreme Value. A distribution used to model block maxima.
MCMC	Markov Chain Monte Carlo. Introduced in Section 2.1.3.
ML	Maximum Likelihood. Introduced in Section 2.1.3.
MLE	Maximum Likelihood Estimate. The estimated parameter with ML.
NOAA	National Oceanic and Atmospheric Administration.
NAO	North Atlantic Oscillation. A climatic index defined in Section 3.3.
PC	Principal component. The temporal component of the EOF-analysis. Introduced in 2.3.
POT	Peak over threshold. An approach to modelling extreme events. Introduced in Section 2.1.
RCP	Representative Concentration Pathways. Four green house gas concentration scenarios used by IPCC [2013].
SLP	Sea Level Pressure.
SST	Sea Surface Temperature.

Contents

Chapter 1

Introduction

The climate is changing. IPCC [2013] predicts many and adverse changes in the climate of our earth due to the effects of increased concentrations of greenhouse gasses in our atmosphere. For Denmark, as a coastal nation, especially the change in mean sea level is of great interest as much of our land is not high above sea level and as such at risk of being flooded.
In the high emission scenario, RCP8.5, the IPCC predicts a mean sea level rise of $0.7\,\mathrm{m}$ for the West Coast of Denmark between now and 2100, with smaller changes for the rest of the country. $0.7\,\mathrm{m}$ in itself might not sound like much, but when added on top of storms and tidal effects the results can be devastating. It is assessed that Denmark is one of the countries in Europe that will see the greatest costs compared to their GDP because of the rising sea level [Christensen et al., 2014].
To put it into perspective the sea level in Esbjerg during the December Hurricane of 1999, the most devastating hurricane recorded in Denmark, was at $3.98\,\mathrm{m}$. This corresponds to a 100-year event according to Sørensen et al. [2013]. If the projected mean sea level rise of $0.7\,\mathrm{m}$ is added to the return level then the $3.98\,\mathrm{m}$ will be equal to a 5-year event in 2100. In other words such high sea levels would occur at a rate 20 times higher than today.
That is the naive extrapolation, based on the assumption that the statistics of storm surges remain stationary over the span of the next century. In Menéndez and Woodworth [2010] the trend in the distribution of yearly maxima is examined through a 40-year period, and it is shown that for the East Coast of England there is a statistically significant negative trend in the maxima once the tidal effects and mean sea level rise has been removed from the signal. No statistically significant trend is identified for any of the stations on the West Coast of Denmark or in the German Blight, but since they face the same sea as Eastern England, it does not seem unlikely that a trend exists of either the same sign, if it is caused by a calmer North Sea, or of opposite sign if it is due to stronger eastern winds.
In Figure 1.1 the number of storm surges over different heights/thresholds is plotted for each year. There is no clear trend for the low thresholds, but as they rise above $200\,\mathrm{cm}$ more events seem to happen in the first half of the period. In Figure 1.2 the total number of storm surges pre-1994 is compared with the same for post-1994. For low thresholds there seem to be slightly more events in the post-1994 era, though no statistically significant difference can be seen. As the threshold rises above 200 there appears to be a significantly higher number of

events in the pre-1994 era, indicating that there is a change in the dynamics of the storm surges.
This is the motivation for further examining the statistics of the storm surges on the West Coast and to try to set up a model that can explain the trend if any is found.

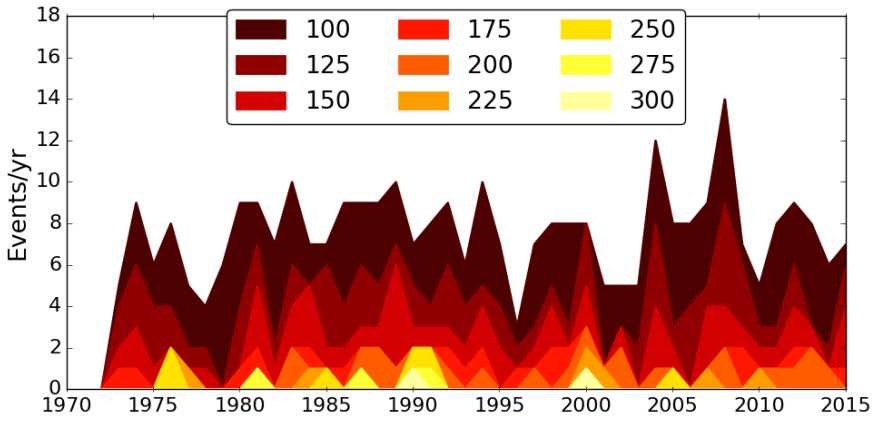

Figure 1.1: The number of yearly observations over a given threshold in centimeters. The data is processed to remove tidal effects and rising sea level.

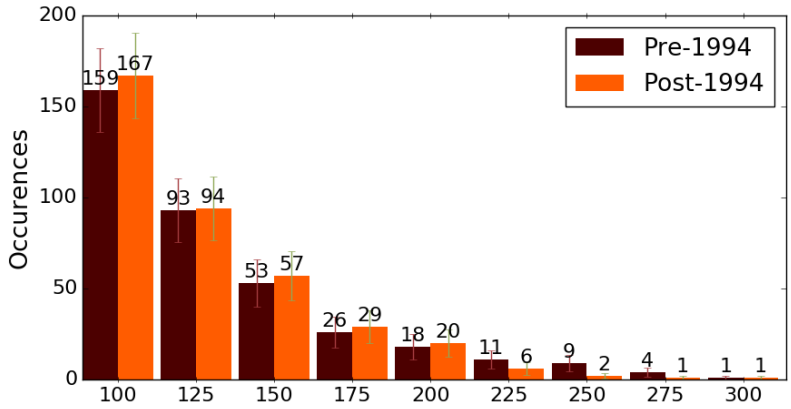

Figure 1.2: The total number of events above a given threshold in centimeters for the period 1973–1994 and the period 1994–2015. The error bars indicate 95%-CI.

Chapter 2

Theory

This chapter aims to lay the basic foundation on which to build the analysis of the storm surges as well as give some justifications for the choices made when multiple models or methods were available. The aim is not to give a thorough presentation of everything but to give an overview and references for further clarifications. Some sections in this chapter might as well has been placed in Section 3.2 and vice versa, so if something seems to be missing here, look there.

2.1 Model choice

When modelling the extreme events there were two approaches that was considered: The block maxima and the peak over threshold. Below they are both listed with their pros and their cons.

Block Maxima (BM)

The data is separated into blocks of a given size (e.g. one year) and the maximum values of these blocks are stored. If the data is drawn from the same distribution the maxima have one possible limiting distribution as the block-size grows. This distribution is known as the Generalized Extreme Value Distribution (GEV).

Pros: The main benefit of this approach is that there is no need to specify a threshold. This eliminates an otherwise problematic step (see sec 3.2.5) and thus decreases the bias that may have been introduced.

Cons: Separating the data into blocks allows for two or more events, that could rightly be considered extreme, to be pooled together. As only the maxima of each block is kept, some extreme events may be discarded when using this method. This introduces the problem of choosing the right block size. A too large block size increases the chance of two extreme events being pooled together and a too low block size might break the asymptotic assumption of the GEV.

Peak Over Threshold (POT)

Each data point is compared to some threshold and, if the data point exceeds the threshold, it is stored as an extreme event. These events can be shown to have one possible limiting distribution as the threshold rises. This distribution is called the Generalized Pareto Distribution (GPD).

Pros: The main benefit of this approach is that it is less wasteful of the data. If two or more extreme events happen within the same block (as defined in the Block Maxima-method), they are both kept and thus more data is potentially available to fit the model.

Cons: The choice of the threshold is not a simple procedure as it is a balancing act between bias and variance. A low threshold will break the asymptotic assumption of the GPD introducing a bias, and a high value will give less exceedances and thus a higer variance of the estimates.

The benefit of no threshold value in the BM was compared with the benefit of no specified block size in POT. Although the first one is considered to be slightly more problematic, the efficiency of the use of data is a decisive factor in this paper, due to the inherent data scarcity of extreme values and therefore POT is chosen.

2.1.1 Generalized Pareto Distribution

The Generalized Pareto Distribution (GPD) is a family of distributions that is often used to model the tails of unknown distributions. An outline of the justification for this can be found in Coles [2001].
The probability density function of the GPD can be seen in Equation 2.1. It has a limiting form as the shape parameter $\xi \to 0$, where it simply becomes the exponential distribution. When $\xi \geq 0$ the only support for the GPD is $\mu \leq x$, this means that the distribution has no upper bound. If $\xi < 0$ the support changes to $\mu \leq x \leq \mu - \frac{\sigma}{\xi}$. This means that the distribution gets an upper bound.

Parameter		Range
Scale	σ	$(0, \infty)$
Shape	ξ	$(-\infty, \infty)$
Location	μ	$(-\infty, \infty)$

Table 2.1: The parameters of the GPD.

$$P(x) = \begin{cases} \frac{1}{\sigma}\left(1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right)^{-\left(1+\frac{1}{\xi}\right)} & \xi \neq 0 \\ \frac{1}{\sigma} e^{-\frac{x-\mu}{\sigma}} & \xi = 0 \end{cases} \tag{2.1}$$

Equation 2.1 is used to make inference about the extreme events and their distribution.

As can be seen from Figure 2.1 a lower value of ξ means a shorter tail on the distribution and likewise for the σ parameter. Noticeably would a $\xi < -1$ result in a distribution with a positive slope. In the case of water level this seems highly questionable, as that would make the most extreme water levels the most likely, and therefore we will expect that $\xi > -1$.

2.1.2 Non-stationary Modelling

A classical stationary approach to modelling would be to estimate the parameters of Equation 2.1 and 2.15 and assume that these parameters were fixed for all observations x. This would make sense if the observations were drawn from a stable distribution or the time span in which the observations were made were very short, but given the climatic changes as described in the introduction and the length of time series needed to assure enough extreme events, it seems reasonable to doubt these assumptions, and therefore non-stationary modelling is applied alongside the stationary.
There is no formal basis for modelling non-stationary processes [Coles, 2001], it

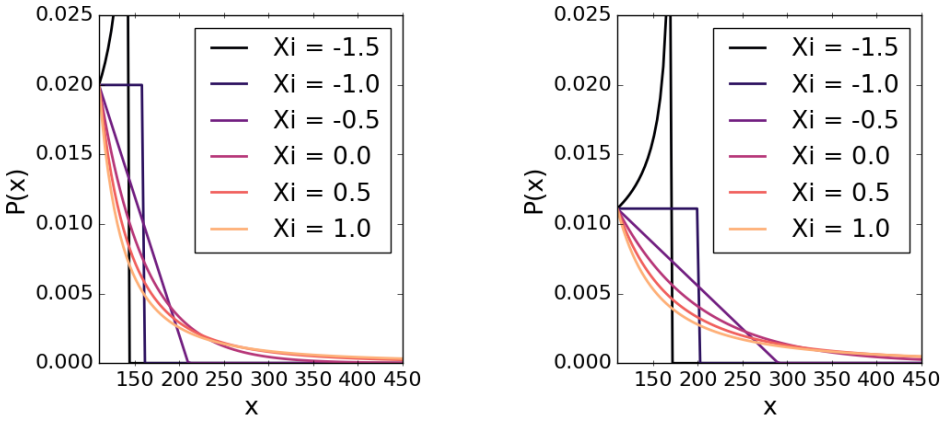

(a) The GPD for $\sigma = 50$ and for varying values of ξ.

(b) The GPD for $\sigma = 90$ and for varying values of ξ.

Figure 2.1: Two figures illustrating the influence of the parameters σ and ξ on the PDF of the GPD. $\mu = 110$ which is the threshold used for the rest of the paper.

is more a pragmatic approach where the parameters are allowed to vary between observations. This is implemented by letting $\sigma(t) = e^{a_\sigma H(t)+b_\sigma}$. Here $H(t)$ is the value of the predictor $H(t)$ at time t. The predictor could for instance be the global temperature. So the observation at time $x(t)$ would have a coupled predictor $H(t)$ which would be included in Equation 2.1:

$$P\left(x(t), H(t)\right) = \frac{1}{e^{a_\sigma H(t)+b_\sigma}}\left(1+\xi\left(\frac{x_i-\mu}{e^{a_\sigma H(t)+b_\sigma}}\right)\right)^{-\left(1+\frac{1}{\xi}\right)} \tag{2.2}$$

A similar substitution can be made for ξ and for the parameters introduced in Equation 2.15.
In theory any function can be chosen for the relationship between the parameters and the predictor, e.g. constant, linear, polynomial and exponential. For σ it was natural to choose an exponential relation so that $\sigma > 0$, as required by constraints of the GPD, but for ξ the linear relation $\xi(t) = a_\xi H(t) + b_\xi$ would be an option. Higher order dependencies are also possible but not considered in this paper as to keep the model as simple as possible.

2.1.3 Methods of parameter estimation

Once the model has been set up as in the previous sections, a method is needed to make inferences about the true parameters based on the observations. Given a set of n exceedances of some threshold, μ, $\mathbf{x} = (x_1, \ldots, x_n)$ one can set up a likelihood function based on Equation 2.1:

$$P\left(\xi, \sigma|\mathbf{x}\right) = \prod_{i=1}^{n}\frac{1}{\sigma}\left(1+\xi\left(\frac{x_i-\mu}{\sigma}\right)\right)^{-\left(1+\frac{1}{\xi}\right)} \tag{2.3}$$

Often it is more convenient to look at the log-likelihood-function (LLH). This converts the product into a sum and also makes it more convenient to manipulate

very small probabilities numerially.

$$\mathcal{L}(\xi, \sigma|\mathbf{x}) = \sum_{i=1}^{n} -\log(\sigma) - \left(1 + \frac{1}{\xi}\right) \log\left(1 + \xi \frac{x_i - \mu}{\sigma}\right) \tag{2.4}$$

A similar transformation is made for the case $\xi = 0$ as in Equation 2.1.
Additional parameters and predictors can be added as described in Section 2.1.2.
Two methods will be used to make inference about the parameters: Maximum Likelihood (ML) and Markov Chain Monte Carlo (MCMC).

Maximum Likelihood

The ML-method is simply an optimisation of Equation 2.4 with regards to the parameters σ and ξ or their sub parameters (a_σ, b_σ) and (a_ξ,b_ξ). There is a plethora of methods to do this optimization and the one chosen is the MIGRAD fitting method, which uses a variable-metric method with inexact linear search[Hanlon, 2008]. A more full explanation of the method can be seen in James [2004]. Importantly it is heavily dependant on the first derivative and will likely be very inaccurate when this does not behave properly and like all other numerical minimizers it is never certain whether it locates a global or a local minima.
The first derivatives of the LLH can be calculated analytically and have no apparent problematic areas in the appropriate range of the parameters. To examine the LLH-landscape and check whether it seems feasible that there is only one local minimum a raster scan is made as can be seen in Figure 2.2.
When looking at the scans there are big parts of the parameter space that is invalid. This is due to the constraint $x \leq \mu - \frac{\sigma}{\xi} \forall \xi < 0$. Apart from this constraint the landscape looks smooth and unimodal, and it seems reasonable to assume that only one maximum exists.
The parameters that maximize the LLH-function are called the Maximum Likelihood Estimates (MLE) of the parameters. The notation is to put a hat on the parameter, so the MLE of ξ is $\hat{\xi}$. The approach can be quite fast, but does in its simple form only give the MLE and no inference about the uncertainty. By use of Wilks' theorem a confidence interval of arbitrary size can be created and thus inference about the uncertainty of the MLE can be made [Wilks, 1938].

Markov Chain Monte Carlo

The MCMC method is a more exploratory approach. The main principle is to sample the posterior distribution by creating a Markov Chain that has the desired distribution as its equilibrium distribution. The step-function of the Markov Chain is chosen to be the Metropolis-Hastings due to technical limitations. More sophisticated methods exist but the resulting distribution should be the same as the length of the chain becomes large enough.
Uninformative uniform priors have been chosen for the parameters as no knowledge about their distribution has been derived. The range of the parameters have been chosen based on prior MCMC runs so that the sampled posterior is at a comfortable distance from the upper and lower bounds. The convergence of the

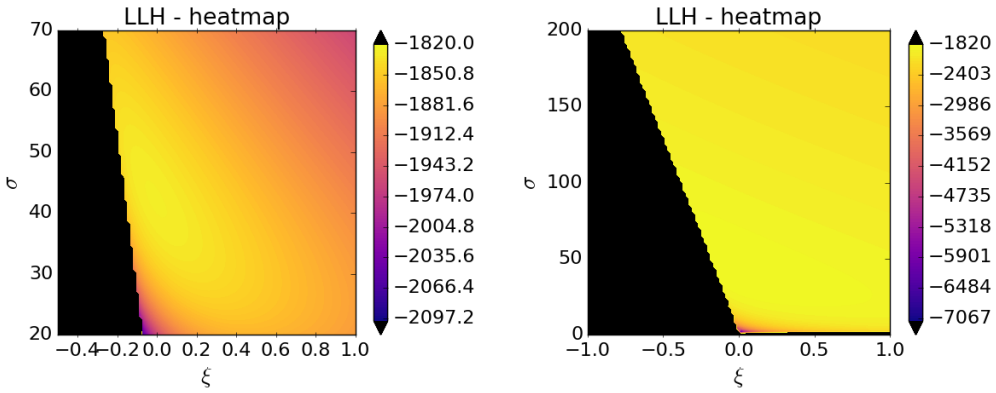

Figure 2.2: Two plots of the LLH-landscape of the Ballum data. The plot on the right shows a macroscopic picture of what is deemed to be the realistic parameter landscape. The black area is invalid parameter-combinations. The plot on the left is a zoom into the more likely area of parameters.

chains is assured by inspecting the chains and a burn of 5000 steps was chosen. The total length of the chains were 20000 and 4 chains were made. More stringent convergence criteria could have been chosen but the chains were deemed to be sufficiently well behaved to choose a fixed burn.
This method is a lot slower than ML at getting a good estimate of the parameters, but in return it gives an approximation of the posterior distribution, which makes it easy to get an estimate of the distribution and covariance of the parameters. It also makes it simpler to propagate the uncertainty of the parameters numerically to other calculations as you can make the calculation for the entire ensemble.

So the ML is a lot faster but gives close to no information about the posterior distribution compared to MCMC. Therefore the ML-method is well suited for superficially exploring large amount of data, where only a basic understanding is needed (as in the teleconnection analysis) while the MCMC is preferred when in-depth analysis is the goal.

2.1.4 Model comparison

As described in Section 2.1.2 a set of different models can be created by using different functions to express the relation between the predictor and the parameters. When there is several models to choose from a method to decide which model performs best is needed. There is a lot of different methods for this. Different variations of the likelihood-ratio test which compares $\hat{\mathcal{L}}$ for the different models. Information Criterion, such as DIC, AIC, WAIC which aims to penalize model complexity in different ways. Various Cross-Validation-techniques that quantifies the predictive ability of the model. Bayes factor which is based on sampling the posterior distribution. This is just to illustrate that there is no single way to do it.
The likelihood-ratio tests is for nested models. The leave-one-out cross-validation and WAIC technique were tried but they were overly to be sensitive to the out-

liers of the data[1]. The Bayes factor were chosen as a method as it is both simple and flexible.
The Bayes factor is simply the ratio between the probability of the observed data given one model H_1 and the probability of the observed data given another model H_2:

$$B_{12} = \frac{P(\mathbf{D}|H_1)}{P(\mathbf{D}|H_2)} \tag{2.5}$$

$$\rightarrow \ln B_{12} = \ln P\left(\mathbf{D}|H_1\right) - \ln P\left(\mathbf{D}|H_2\right) \tag{2.6}$$

The probability $P(\mathbf{D}|H)$ can be approximated as:

$$\hat{P}(\mathbf{D}|H) = \left(\frac{1}{m}\sum_{i=1}^{m}\left(P(\mathbf{D}|\theta_i\right)^{-1}\right)^{-1} \tag{2.7}$$

$$\rightarrow \ln \hat{P}\left(\mathbf{D}|H\right) = \ln m + \ln P\left(\mathbf{D}|\theta_0\right) - \ln\left(1 + \sum_{i=1}^{m} e^{\ln P(\mathbf{D}|\theta_0) - \ln P(\mathbf{D}|\theta_i)}\right) \tag{2.8}$$

Here θ_i is drawn from the posterior distribution $P(\theta_i|\mathbf{D})$ which is simulated with MCMC for each model as described in Section 2.1.3 and $P\left(\mathbf{D}|\theta_0\right) = \max P\left(\mathbf{D}|\theta_i\right)$. A more elaborate derivation of these relations can be seen in Kass and Raftery [1995].
The factor B_{12} can then be used to determine how much faith is put into H_1 compared to H_2. The factor can then be interpreted according to the table 2.2. Note that these are just guidelines and one can weight in prior faith in one model over the other when interpreting the results.

2ln (B_{12})	B_{12}	Evidence against H_2
0 to 2	1 to 3	Not worth more than a mention
2 to 6	3 to 20	Positive
6 to 10	20 to 150	Strong
> 10	> 150	Very strong

Table 2.2: Bayes factors and their interpretations as described in Kass and Raftery [1995]

When comparing n different models to one basis-model H_1 the fraction of belief in one model out of all models can be expressed as:

$$P(H_k|\mathbf{D}) = \frac{B_{k1}}{\sum_{i=1}^{n} B_{i1}} \tag{2.9}$$

$$\rightarrow \ln P\left(H_k|\mathbf{D}\right) = \ln B_{k1} - \ln B_{i1} - \ln\left(1 + \sum_{j\neq i}^{n} e^{\ln B_{j1} - \ln B_{i1}}\right) \tag{2.10}$$

Where $B_{i1} = \max B_{j1}$.
This approach can be used to quantify model uncertainty and would allow to make hybrid-estimations of parameters shared by the models, but in this paper it is only used as a convenient way to show how much faith there should be put in one model compared to the others.

[1] such as the December storm of 1999

2.1.5 Return levels

When interpreting the model it is often more illustrative to look at return levels than the estimated parameter-values. Return levels is a way of inferring how likely an event of a certain magnitude is. The return level can be calculated from the tail distribution of the GPD:

$$P(X > x|x > \mu) = \left(1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right)^{\frac{-1}{\xi}} \tag{2.11}$$

The conditional probability $P(X > x|x > \mu)$ can be rewritten as $\frac{P(X>x)}{P(x>\mu)}$. Since $P(x > \mu)$ is simply the probability of an exceedance occurring, λ, in one time step, τ, the equation can be rewritten:

$$P(X > x) = \lambda\left(1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right)^{\frac{-1}{\xi}} \tag{2.12}$$

So if a water level x_m is expected to be exceeded every m'th time step it must follow that the probability for it to be exceeded at every time step is $P(X > x_m) = \frac{1}{m}$.

$$\frac{1}{m} = \lambda\left(1 + \xi\left(\frac{x_m-\mu}{\sigma}\right)\right)^{\frac{-1}{\xi}} \tag{2.13}$$

$$\Rightarrow x_m = u + \frac{\sigma}{\xi}\left((m\lambda)^{\xi} - 1\right) \tag{2.14}$$

This value is called the m-period return level
The model for the exceedance rate λ is chosen to be the Poisson distribution (PD). The probability mass function of the PD can be seen in Equation 2.15. τ is chosen to be the number of observations in a month and k is then the number of exceedances in the given month. Thus we can make inference about the rate parameter λ and use it when calculating the return level.

Parameter		Range
Occurrences	k	$(0, \infty)$
Rate	λ	$(0, \infty)$
Interval	τ	(k, ∞)

Table 2.3: The parameters of the PD.

$$P(k, \lambda, \tau) = \frac{(\lambda\tau)^{k} e^{-k}}{k!} \tag{2.15}$$

Notice that λ is required to be positive so a natural choice when making a non-stationary model is to model λ as an exponential function of the predictor:

$$\lambda(t) = e^{a_\lambda H(t) + b_\lambda} \tag{2.16}$$

2.2 Teleconnection

Teleconnections are the discipline of linking together climate anomalies over large distance, sometimes thousands of kilometers. This is done by looking at the correlation between 2 different climatic variables, for instance the water level on the West Coast of Denmark and the sea surface temperature on the east coast of North

America. By looking at long time series one can get a good estimate of how correlated the variables are [Schneider et al., 2011].
Normally one would determine the correlation coefficient for the two time series, but in this paper the GPD- and POI-distribution is fitted instead using the water level as the primary time series and the other variable as the predictor, as described in Section 2.1.2. If for instance the water level in Ballum and the sea surface temperature is chosen, one could look at a grid of temperatures all over the world and generate a map of fitted values as in Figure 4.1a. Here the a_λ is plotted and it shows which areas the fit returns a non-zero dependency, which can be interpreted as a good correlation.
This approach has the downside that it assumes the chosen model is correct or at least a good approximation, which is not necessarily true. But as the further analysis assumes that the model is a good approximation it was deemed more informative to fit the distributions than to calculate the correlation coefficients.
Teleconnection analysis are used as a tool to locate patterns that are believable, but it is important to be critical, for instance could simple seasonal residuals in the variables lead to unwarranted correlations. If a feasible pattern is located it is also worth investigating whether it might be a proxy for another pattern, for instance in Section 5.1 a teleconnection pattern is identified as a proxy for the NAO. This does not necessarily invalidate the pattern, but it should be considered before declaring that a new pattern has been found.

2.3 Empirical Orthogonal Functions

Once a teleconnection analysis has been made it is desirable to isolate the signal that caused the correlation. One approach to do this is through Empirical Orthogonal Function-analysis (EOF-analysis). The main idea is to take a temporal-spatial field and dissolve its covariance matrix into eigenvectors. Each eigenvector (an EOF) has a corresponding time series, called a principal component (PC), and together they form a temporal-spatial field. Each of these fields explains a fraction of the variance of the original field, and any major patterns in the original field is therefore expected to be represented in one or more of the EOFs. This makes it possible to isolate the patterns without background noise.
A thorough derivation and explanation of its application can be found in Navarra and Simoncini [2010], but a short summary of how the method is applied in this project will be given here.
Given a set of n time series $\mathbf{x}$ forming a grid around the world[2], one can create the matrix $\mathbf{X} = [\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_n}]$.
From here the covariance-matrix, $\mathbf{R} = \mathbf{X}^t\mathbf{X}$, is created and one solves the eigenvalue problem:

$$\mathbf{RC} = \mathbf{C\Lambda} \tag{2.17}$$

Here the diagonal of Λ is the eigenvalues $[\lambda_1, \lambda_2, \ldots, \lambda_n]$ and the rows of $\mathbf{C}$ is the corresponding eigenvectors. By normalizing the eigenvectors the magnitude of

[2]Each time series is for instance the observations from one observational post, or as in HadCrut a mix between interpolations and observations

the eigenvalues will represent how dominant the eigenvector is in the composition of the covariance-matrix, and thus the variance in the field $\mathbf{X}$. By ranking the eigenvectors by the magnitude of their corresponding eigenvalues one gets a sequence $[\mathbf{c_1}, \mathbf{c_2}, \ldots, \mathbf{c_n}]$. The vector $\mathbf{c_1}$ is referred to as the first EOF and so on. An example of the first 3 EOFs of the SST can be seen in Figure 4.6
By projecting the EOFs back unto the original field a sequence of time series (PCs) is acquired $[\mathbf{a_1}, \mathbf{a_2}, \ldots, \mathbf{a_n}]$, an example of such can be seen in Figure 5.2.
If the pattern from the teleconnection analysis matches one of the EOFs well, then the PC of that EOF can be used as the predictor. Alternatively one can reconstruct the original field:

$$F_i = \sum_{j=1}^{n} \mathbf{c_j} a_{j_i} \tag{2.18}$$

Here F_i is the reconstructed field at time i, a_{j_i} is the i'th value of a_j and n is the number of EOFs used in the reconstruction. The choice of n can be used to filter out the less important components and thus simplify the signal, resulting in a simplified predictor.
Another application of the EOFs is to extrapolate the predictors derived from them. If a series of EOFs is derived from for instance the HadCRUT (see Section 3.3) they run up until today. But by projecting model runs for future SST (see Section 5.4) unto the EOFs, a series of pseudo-PCs can be generated that can be used to reconstruct a field that has been extrapolated in time, and thus make inference about how the climate is expected to evolve if the model is correct.

Chapter 3

Data

This chapter introduces the data and the explains the pre-processing of the data before it is used for statistical inference. They are separated into two categories: Water level measurements and predictors.

3.1 Water level

The water level data is the basis of the analysis. The data is gathered from tide gauges on the southern part of the West Coast of Denmark. In Figure 3.1 the location of the four tide gauges can be seen. They are all placed in the same area and as such will only give an idea about the behaviour of the water level here. The ambition was to include tide gauges along all of the West Coast, but no success was had in acquiring the data.

Figure 3.1: The location of the four tide gauges from which the water level measurements in this paper originate

In Table 3.1 the meta-data for the water level measurements can be seen. It is worth noting the relatively high amount of missing data from Ballum. This should be weighted in when analysing results based on this station. Curiously enough the number of events is the second highest in spite of the missing data. This could either indicate that Ballum has more storm surges than the other stations or that the missing data is not masking many storm surges. Since all the stations are placed in the same region, and therefore expected to have similar water level, the latter seems more likely.

The station at Højer was moved in 1980 from Højer Sluse to the front of the levee [Sørensen et al., 2013], so it seems likely that there is some different dynamics in the first 8 years of the time series. This should also be taken into account when analysing the results.

City	Coordinates	Period	Range/cm	Missing data	Events	Source
Esbjerg	N55°46′ E8°45′	05/12/1972 to 23/02/2016	-294 to 409	9.6%	271 (1.73%)	KDI
Ballum	N55°13′ E8°69′	05/12/1972 to 23/09/2016	-49 to 430	25.9%	298 (1.91%)	KDI
Havneby	N55°9′ E8°69′	05/12/1972 to 18/05/2016	-296 to 437	13.3%	253 (1.62%)	KDI
Højer	N54°96′ E8°66′	05/12/1972 to 18/05/2016	-191 to 472	12.5%	354 (2.26%)	KDI

Table 3.1: Meta-data for the water level measurements. The range is for the preprocessed data while the missing data and events is for the post-processed data. Missing data indicates how big a fraction of the days that do not have any observations in them. Events is number of surges above $110\,\mathrm{cm}$*, and the percentage is how big a fraction of the days contain an event.*

3.2 Processing of water level measurements

Before further analyzing the data it is needed to define exactly which trends of the water level that are desired in the analysis. For instance is it desirable to include tidal effects or should they be filtered out pre-analysis? In the case of tidal effects, they are well understood and complicates the analysis unnecessarily. In general all predictable trends should be removed, so that we are only left with a signal that is mostly unpredictable by common methods. In the following sections different pre-analytic considerations on how to prepare the data for later modelling are presented.

3.2.1 Data quality

The data that was acquired from KDI had anomalies in them, presumably due to errors in the sensors. There were both obvious and less obvious anomalies in the dataset. Four categories of these can be seen in figure 3.2.
The extreme signal as seen in figure 3.2a where a 600 m high surge appeared is obviously an artifact of instrumental malfunction. It is easy to detect both visually and algorithmically by looking for the flat signal preceding or following the peak, which is easily found by looking at the variance of the signal, which drops to 0.
The second kind is dubbed ”signal cutoff” as seen in figure 3.2b. This could be due to a floating device getting stuck and thus setting an upper limit. It could potentially mask extreme events, but the artifact is not in itself problematic for the analysis, as there is no fallouts that could be mistaken as an extreme event.
The artifacts as seen in figure 3.2c is the more common in the dataset and can vary in size from a few hours to a few days as here. They are characterized by a stepwise change in the data and clearly doesn't follow the expected tidal behaviour. In the case of the Ballum dataset there is around 150 of these artifacts,

but in this paper it has been decided to (in most cases) ignore them, as they are in the low register and therefore will not mask any extreme highs included in the analysis. They will have an influence on the analysis of tidal effects though, but it is assumed that this is negligible as they only account for a miniscule fraction of the time series.
The last type of artifact as seen in figure 3.2d is a bit more problematic. The signal flattens at a high water level and then start again after some time. This makes them easy to detect as the variance is 0 when the signal is flat lining. These artifacts seem to coincide with a rising water level as seen in the figure, and rightly so there was a storm on the 28th Nov. 2011. So there is a high risk that the missing data masks an extreme event. Since the anomaly is suspected to mask an extreme event it is not feasible to make any sort of interpolation or estimation of the water level based on the rest of the time series. Data from other nearby stations may be used to make an estimate of the highest water level during these anomalies but that has not been done in this analysis.

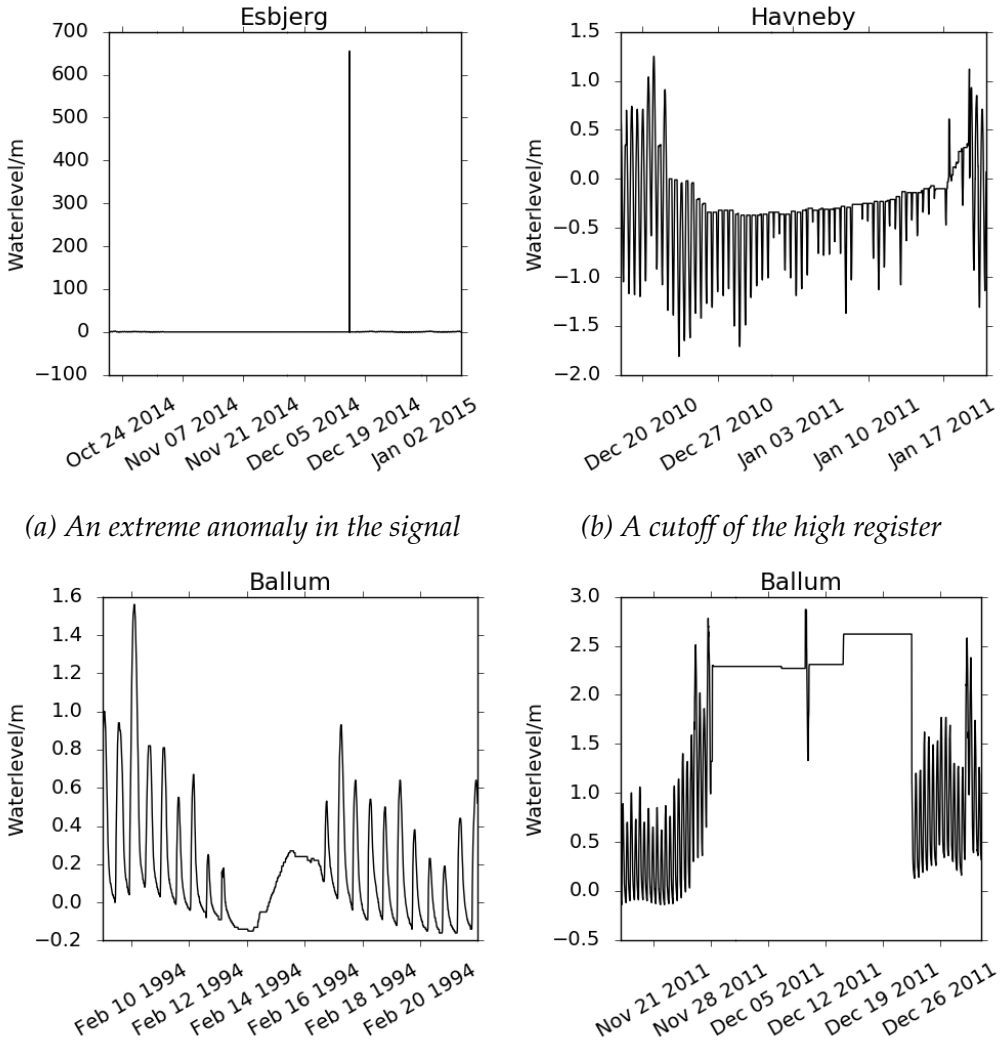

(a) An extreme anomaly in the signal

(b) A cutoff of the high register

(c) Step-like signal in the low register

(d) Signal dropout in the high register

Figure 3.2: Four examples of anomalies in the data.

3.2.2 Tidal effects

The tidal effects are estimated by doing a least squares regression of Equation 3.1 with regards to the observed data, here a_i is the angular speeds predefined by established tidal constituents (i), and thus the fitted components are the amplitude (D_i) and the phase (θ_i)[Foreman, 1977].

$$F(t) = F_0 + \sum_{i=1}^{n} D_i \cdot \cos\left(a_i \cdot t + \theta_i\right) \tag{3.1}$$

The tidal constituents used are the ones used by the NOAA[Hicks, 2006], as opposed to choosing from the full spectrum of constituents. This is done for convenience and speed. The resulting tidal filter seems reasonably good.
An example of the data and the corresponding tidal fit can be seen in figure 3.3a and the resulting residual after subtracting the tidal model from the data can be seen in figure 3.3b.

3.2.3 Changing sea level

This paper does not aim to do any inference about the change in mean sea level, as that is examined in other papers in greater detail. To negate the changing sea level a rolling average of some interval is subtracted from the data. This also serves to remove any seasonal variation in the water level still left over from the tidal filter.
One unwanted effect from this is that it compress the dynamic range slightly. A great surge in water level is not paired with an equal low, and as such a month with one or more surges will likely have a higher mean than one without any, thus after removing the mean the peak values will be closer to the ground level. A balance between catching seasonal and other temporal change in the mean sea level and keeping the peaks as distinct features as possible is sought.
As can be seen in figure 3.4 three fitted parameters is plotted against the size of the rolling mean. There is a clear effect in both ξ and a_λ up until an interval of 30 days, but in a_σ there's an influence up until around 50 days. It seems reasonable to assume that the sea level change and seasonal change is constant through a 50 day period, so that is the interval chosen for this paper.

3.2.4 Declustering

Storminess is the primary driver for storm surges in Denmark and they can have durations of several days [Sørensen et al., 2013], thus raising the sea level for a prolonged amount of time. This means that peaks have a tendency to cluster. As one of the criteria for the GPD is independent data points, some declustering has to be done. This is done by setting up a rule for what is considered a cluster and then discard all threshold-exceedances except the cluster-maxima[Coles, 2001]. A simple way to do it is to decide on a minimum interval that has to separate two events. In figure 3.5 three fitted parameters are plotted against the minimum interval seperating two events.
The uncertainty on the MLEs is too big to pinpoint a value for which the cluster

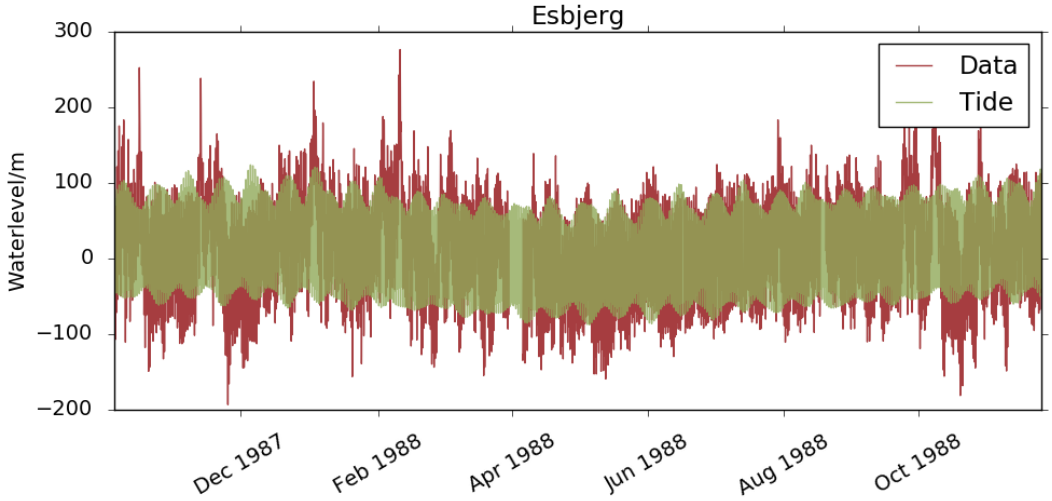

(a) The original data (red) and the fitted tidal model (yellow)

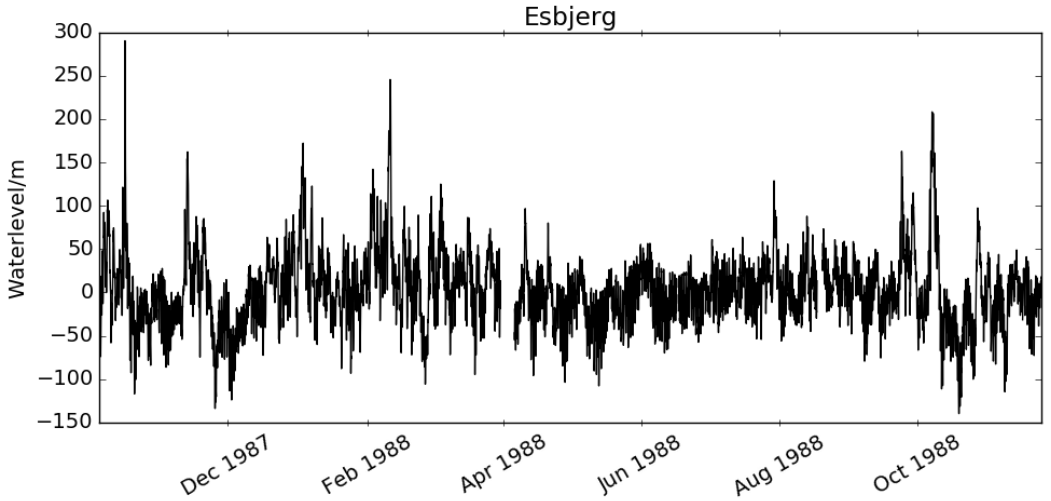

(b) The residual after subtracting the tidal model

Figure 3.3: The signal at Esbjerg being treated with the tidal filter proposed in this section. Importantly the peaks of the original signal is still clearly distinguished in the residual.

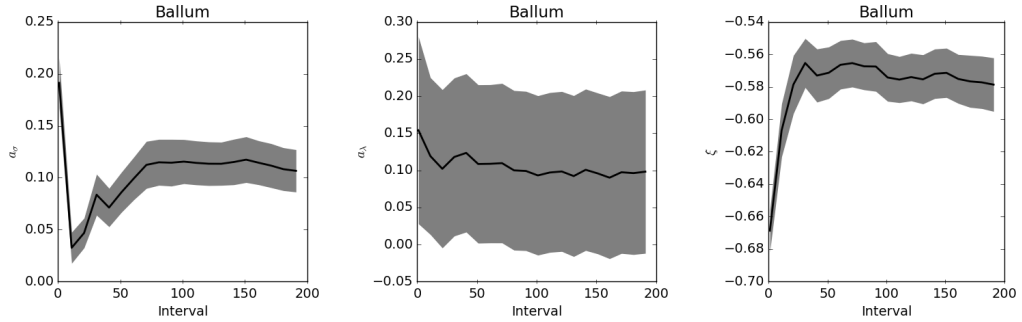

(a) The MLE of α_σ plotted against the boxcarsize.

(b) The MLE of α_λ plotted against the boxcarsize.

(c) The MLE of ξ plotted against the boxcarsize.

Figure 3.4: Three MLE's and their dependency on the boxcarsize.

size does not influence the MLEs. A choice of 5 days is made as it seems from figure 3.5b that this is past the initial drop.

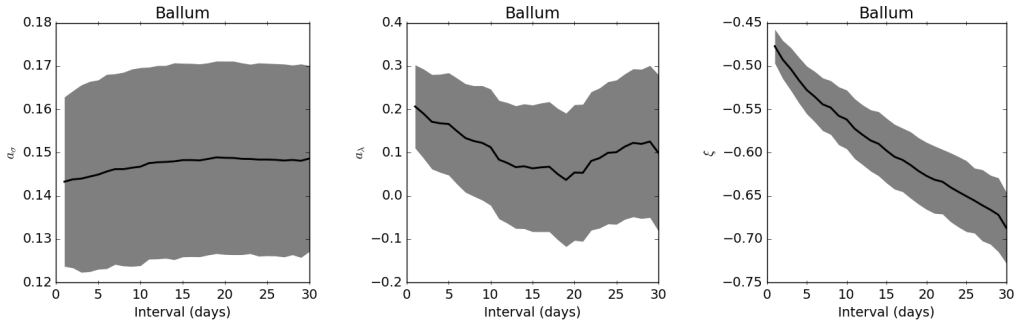

(a) The MLE of a_σ plotted against the minimum distance between extremes.

(b) The MLE of a_λ plotted against the minimum distance between extremes.

(c) The MLE of ξ plotted against the minimum distance between extremes.

Figure 3.5: Three MLE's and their dependency on the minimum distance between extremes.

3.2.5 Threshold selection

Selecting the threshold μ in (2.1) is a balancing act between bias and variance. A too low threshold will break the asymptotic behaviour of the GPD introducing a bias into the estimate and a too high threshold will lead to incorrectly discarding data and thus decreasing the quality of the estimate, leading to a higher variance. One way do decide on a suitable threshold value is to look at the Mean Residual Life (MRL) which will be a linear function of μ when a suitable threshold has been chosen[Coles, 2001]. This has been done in 3.6a, where it can be seen that for $\mu \in (50; 90)$ there is an approximate linear behaviour. Then the line breaks but gets linear again for $\mu \geq 110$.
Another complimentary method is to plot the MLE of ξ against the threshold. When a suitable threshold has been chosen, ξ should be independent of the choice of threshold[Coles, 2001]. This has been done in figures 3.6b and 3.6c. Since the model is non-stationary the free variables a_ξ and b_ξ has been plotted instead, but they are expected to be constant as well. As can be seen there is a clear dependence in b_ξ up until $\mu \simeq 65$, where after the value is approximately stable until it drops off at around $\mu = 120$. For a_ξ the value varies steadily around 0 and does not show any dependency.
In figure 3.6d the percentage of events exceeding a given threshold. As can be seen the amount of exceedances for $\mu = 65$ is 7% which is approximately twice per month, comparatively if $\mu = 110$ the amount of exceedences is 2%, which is approximately once every second month. It seems a bit excessive to have two extreme events per month and therefore a threshold of 110 has been chosen in this paper.
A similar analysis has been done for the other stations and predictors and the value is stable across the datasets.

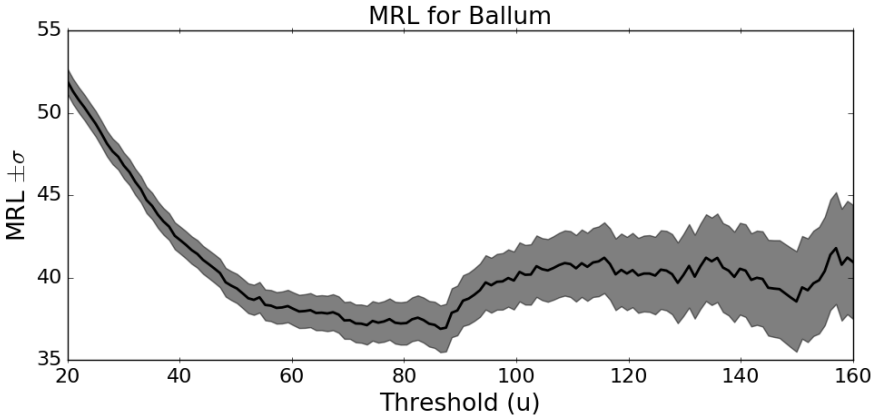

(a) (u,MRL)-plot Ballum. The shaded areas are the σ uncertainties.

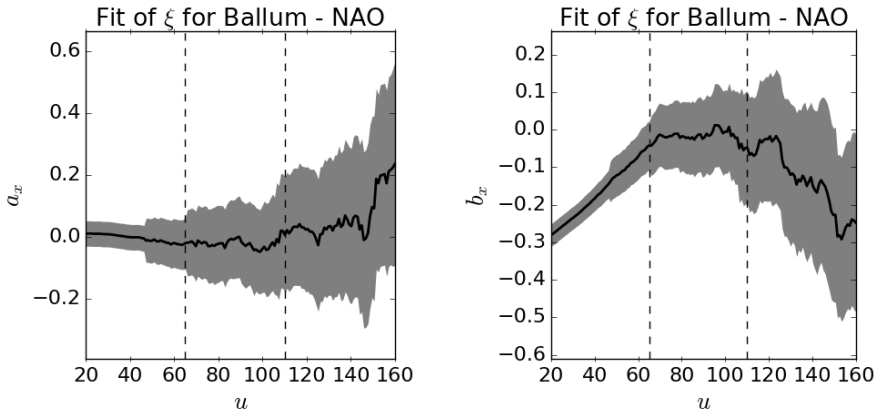

(b) The fitted a_x against the threshold for the Ballum dataset against the NAO-indicator. The vertical lines show $u = 65$ and $u = 110$.

(c) The fitted b_x against the threshold for the Ballum dataset against the NAO-indicator. The vertical lines show $u = 65$ and $u = 110$.

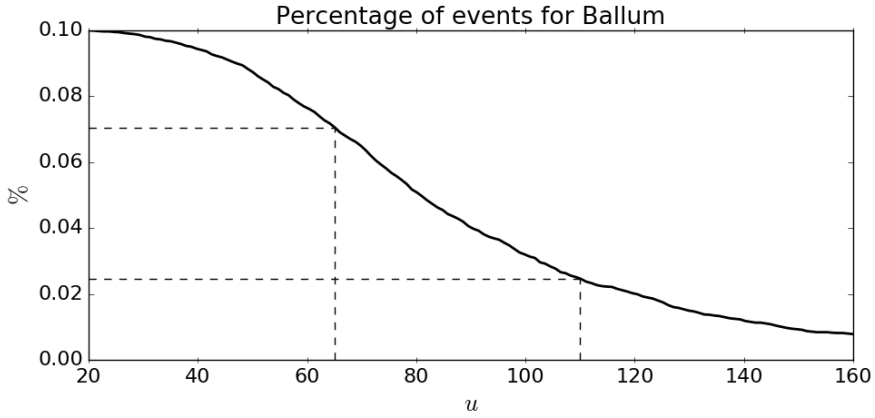

(d) The percentage of events above a certain threshold. The vertical lines show $u = 65$ and $u = 110$.

Figure 3.6: Threshold plots.

3.3 Predictors

This section introduces the different predictors used.

3.3.1 Single time series

The single time series are used directly as predictors in the MCMC-simulations with only minor adjustments such as rescaling, or converting to anomalies.

North Atlantic Oscillation

Period	01/1950 - 08/2016
Range	-2.6 to 2.6
Source	National Oceanic and Atmospheric Administration

The NAO as defined in this paper is the one used by the NOAA. It is calculated as the first Rotated Empirical Orthogonal Function of the SLP of the Northern Hemisphere north of N20° as defined in Barnston and Livezey [1987]. The Rotated EOF is a more advanced form of the EOFs used in this paper, where the constraint of orthogonality between the EOFs is loosened. The NAO is expected to be one of the best predictors as it has been shown to have a strong influence on the weather in Europe especially during the winter. One problem is that it is difficult to predict. Therefore no extrapolation of the results will be based on the NAO.

Global Mean Temperature

Period	01/1850 - 07/2016
Range/$C°$	-0.9 to 1.2
Source	Met Office Hadley Center

The Global Mean Temperature is derived from the SST-field described in Section 3.3.2, by taking the mean of the field. Since the grid is not evenly spaced each grid point is weighted with the cosine of its latitude. This indicator is the only predictor based on a global dataset. Thus it is expected to represent a wide range of weather patterns both relevant and irrelevant, and therefore it will be used as our baseline predictor, for when we compare predictors in Section 5.2.

Noise from Global Mean Temperature

Period	01/1850 - 07/2016
Range/$C°$	-0.9 to 0.9
Source	Met Office Hadley Center

This predictor is used as a check for the validity of the predictors. The noise is generated so that it has the same power spectrum as the Global Mean Temperature, as described in Ebisuzaki [1997]. Only one time series of noise is generated, so it is not going to be used as a tool of estimating statistical significance, but only as an indicator of the validity of the other predictors when comparing the Bayes factors.

EOFs

The EOFs that are identified as relevant through the teleconnections, are transformed into single time series in three different ways.
For the SST a simple weighted mean of the area is done for the reconstructed

field.
For the SLP a weighted average of the area is done for the reconstructed field and it is then it is transformed into monthly anomalies with base in 1961-1991.
For the third EOF of the SST the PC of the EOF is used as the predictor with no further modifications.

3.3.2 Field of time series

The fields of time series are used for making teleconnection analyses to explore potential teleconnection patterns. On the basis of the teleconnection some sub-fields can be subjected to EOF-analysis and through either a mean of the EOFs or simply by the PCs, the fields are reduced to a single time series, used for the MCMC-simulations.

Sea Level Pressure

The HadSLP2r(Low variance) is a global field of monthly mean sea level pressure (MSLP) released by the Hadley Center in the UK. It is based on historical data dating back to 1850 up till 2012, and the grid is made complete by interpolation of the data. A near-real-time product exists, HadSLP2r, which goes up until today, but it is not consistent in mean and variance with the pre-2012 data and therefore not suitable for the long-term analysis of this paper. Further information can be found in Allan and Ansell [2006].

Grid	A 5° latitude by 5° longitude grid covering the entire globe.
Period	01/1850 - 12/2012
Range/hPa	-26.1 to 34.9
Name	HadSLP2rLowVar
Source	Met Office Hadley Center

Sea Surface Temperature Anomaly

The HadCRUT4 is a field of temperature anomalies with basis in 1961-1991 monthly medians. The original dataset is a 100 member ensemble based on observations with different bias-adjustments, here the median of the ensemble is used. The version used in this paper is one where interpolation has been used to fill in time slices with no observations as described in Cowtan and Way [2014]. Further information about the original HadCRUT4 can be found in Morice et al. [2012].

Grid	A 5° latitude by 5° longitude grid covering the entire globe.
Period	01/1850 - 07/2016
Range/C°	-20.4 to 17.2
Name	HadCRUT4
Source	Met Office Hadley Center

Chapter 4

Results

In this chapter the results of the different fits and comparisons is presented. There are 3 different kinds of models that are used in this chapter:

Full: In this model all the parameters are dependent on the predictor H. The relations are as follows:

$$\begin{aligned} \sigma(t) &= a_\sigma \cdot H(t) + b_\sigma \\ \xi(t) &= e^{a_\xi \cdot H(t) + b_\xi} \\ \lambda(t) &= e^{a_\lambda \cdot H(t) + b_\lambda} \end{aligned}$$

Fixed X: Here X is one of the parameters. This is equivalent to the Full model except that the parameter X is kept as a constant.

Fixed: In this model all of the parameters are kept as constants, equivalent to no dependence on the predictor.

The models for the Poisson-distribution and the Generalized Pareto-distribution are kept separate so the dependence for the two can be analyzed separately. Since the two models are independent this does not have any influence on the results.

4.1 Teleconnection

The results of teleconnection analysis as described in Section 2.2 is presented here in this section. As the amount of graphs is large, only a fraction will be included to highlight tendencies and other considerations, the rest can be found in Appendix A.1.
The plots are all for the Full model as described in the introduction of this chapter and only the predictor coefficients, (a_λ, a_σ and a_ξ), are presented as they are the values used later in the analysis.
Additionally the teleconnection-analysis has also been done for the seasons separately to both highlight seasonal dependence and the stability of the patterns in the teleconnection plots. The map plots of the teleconnection have been dotted with a black dot at every grid point where the MLE of the 66%-confidence interval contains 0, so as to not make highly unpredictable areas appear more convincing than they are.

4.1.1 Sea Surface Temperature

The SST is the HadCRUT4 as described in Section 3.3. They are monthly anomalies. In Figure 4.1 is the teleconnection map for the entire series of Ballum and the SST. They have a overlap from 1972 – 2015 (43 years).
To check for seasonal dependency of the parameter the fit has also been done for each of the four quarters (JFM, AMJ, JAS, OND) which can be seen in Figure 4.2.

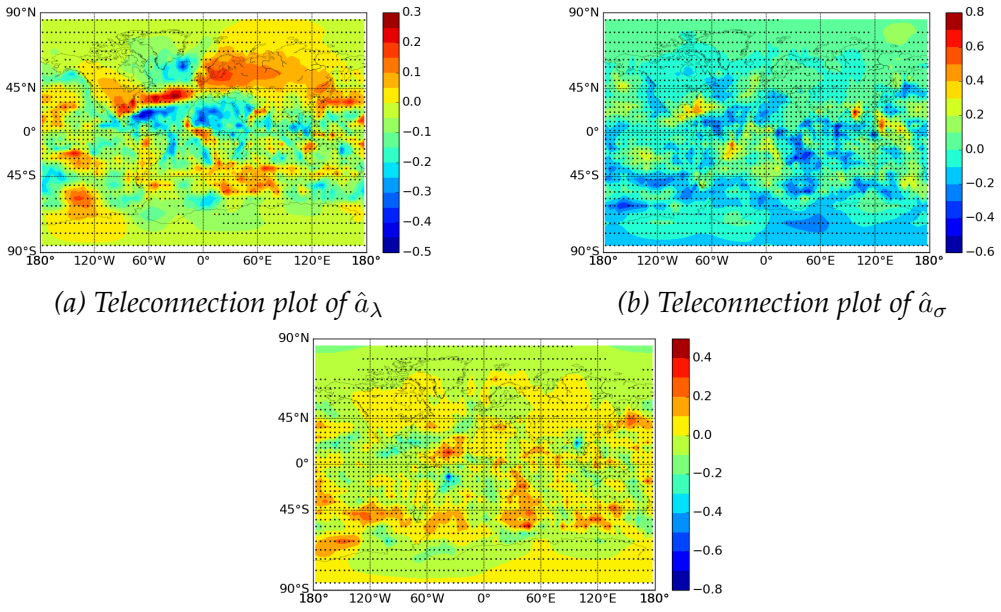

(a) Teleconnection plot of $\hat{a}_\lambda$

(b) Teleconnection plot of $\hat{a}_\sigma$

(c) Teleconnection plot of $\hat{a}_\xi$

Figure 4.1: Teleconnection plots of the MLE of the predictor coefficients for Ballum and the SST. The black dots indicate areas where the 66%-CI spans over 0, which is interpreted as the sign of the coefficient being dubious in that area.

4.1.2 Sea Level Pressure

The SLP is the HadSLP2r as described in Section 3.3. Unlike the SST, these are not monthly anomalies by default. As can be seen in In Figure 4.3 if the analysis is done with the unprocessed field the analysis is dominated by a signal that is only located over equatorial sea-zones with alternating sign between north and south. This indicates a dominant seasonal mode and therefore the SLP is changed into monthly anomalies for the rest of the teleconnection analysis.
The Figure 4.4 is the teleconnection plot for the entire series of Ballum and the SLP. They have a overlap from 1972 – 2012 (40 years).
To check for seasonal dependency of the parameter the fit has also been done for each of the four quarters (JFM, AMJ, JAS, OND) as can be seen in Figure 4.5.

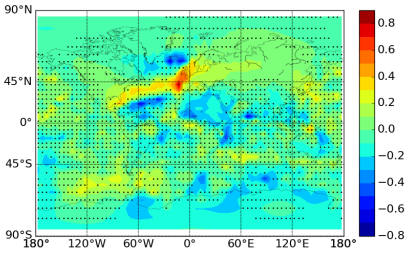

(a) Teleconnection plot of $\hat{a}_\lambda$ for JFM

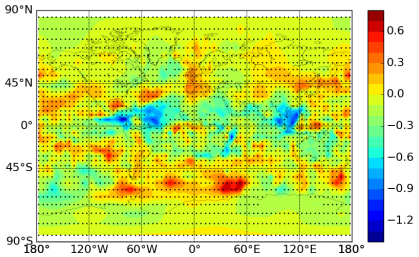

(b) Teleconnection plot of $\hat{a}_\lambda$ for AMJ

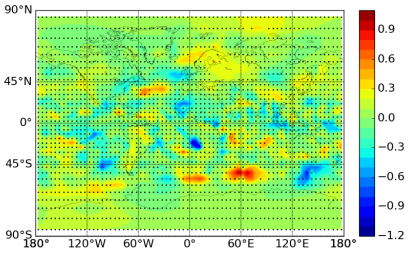

(c) Teleconnection plot of $\hat{a}_\lambda$ for JAS

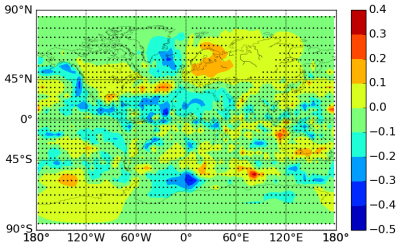

(d) Teleconnection plot of $\hat{a}_\lambda$ for OND

Figure 4.2: Seasonal teleconnection plots of $\hat{a}_\lambda$ for Ballum and the SST. The black dots indicate areas where the 66%-CI spans over 0, which is interpreted as the sign of the coefficient being dubious in that area.

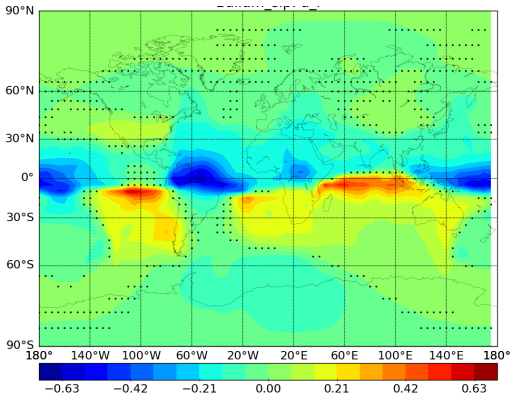

Figure 4.3: Teleconnection plot for $\hat{a}_\lambda$ for Ballum and SLP without shifting to monthly anomalies.

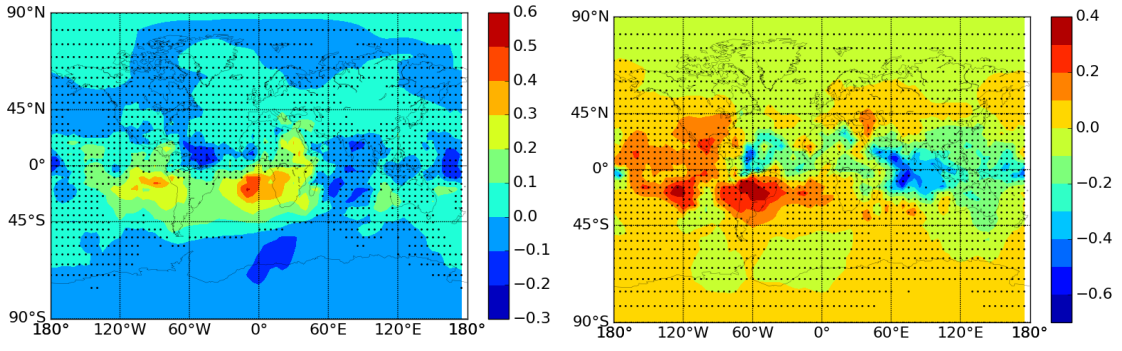

(a) Teleconnection plot of $\hat{a}_\lambda$ *(b) Teleconnection plot of $\hat{a}_\sigma$*

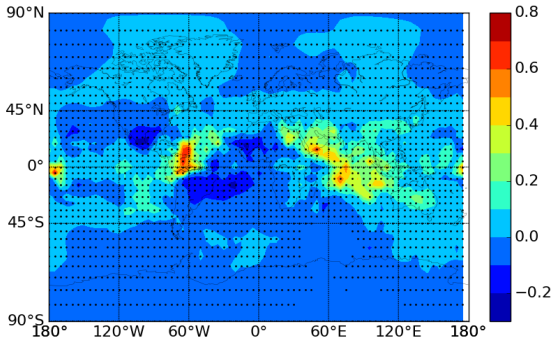

(c) Teleconnection plot of $\hat{a}_\xi$

Figure 4.4: Teleconnection plots of the MLE of the predictor coefficients for Ballum and the SLP. The black dots indicate areas where the 66%-CI spans over 0, which is interpreted as the sign of the coefficient being dubious in that area.

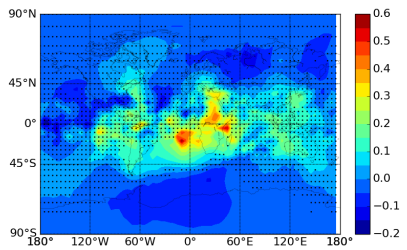

(a) Teleconnection plot of $\hat{a}_\lambda$ for JFM

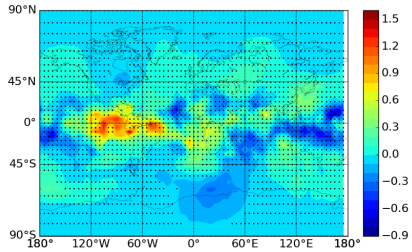

(b) Teleconnection plot of $\hat{a}_\lambda$ for AMJ

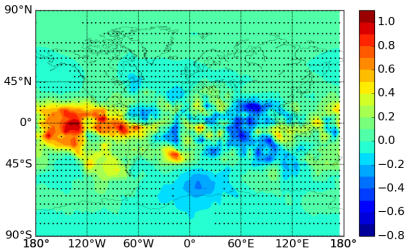

(c) Teleconnection plot of $\hat{a}_\lambda$ for JAS

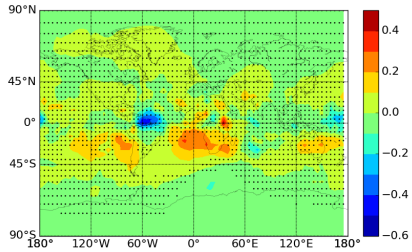

(d) Teleconnection plot of $\hat{a}_\lambda$ for OND

Figure 4.5: Seasonal teleconnection plots of $\hat{a}_\lambda$ for Ballum and the SLP. The black dots indicate areas where the 66%-CI spans over 0, which is interpreted as the sign of the coefficient being dubious in that area.

4.2 Empirical Orthogonal Functions

This section presents the 3 leading EOFs for the SLP and the SST. They are generated from the datasets described in Section 3.3.2 with the method described in Section 2.3.

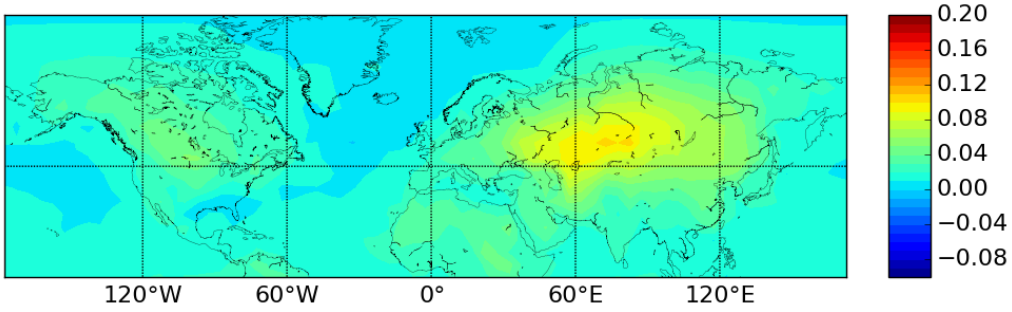

(a) First EOF of SST of Northern Hemisphere. Explains 15.7% of the total variance.

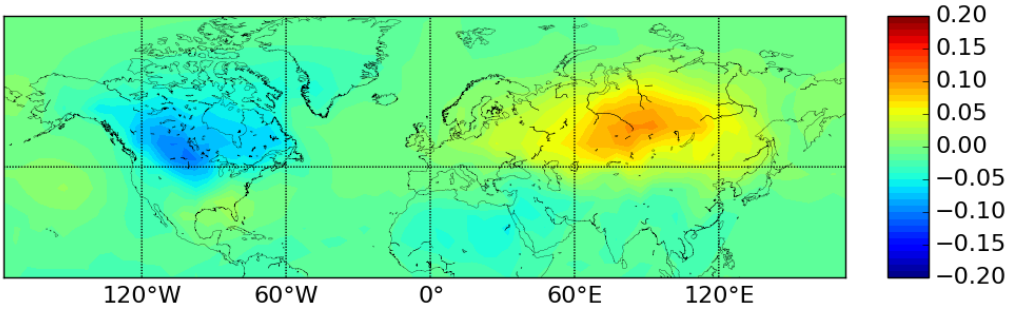

(b) Second EOF of SST of Northern Hemisphere. Explains 8.5% of the total variance.

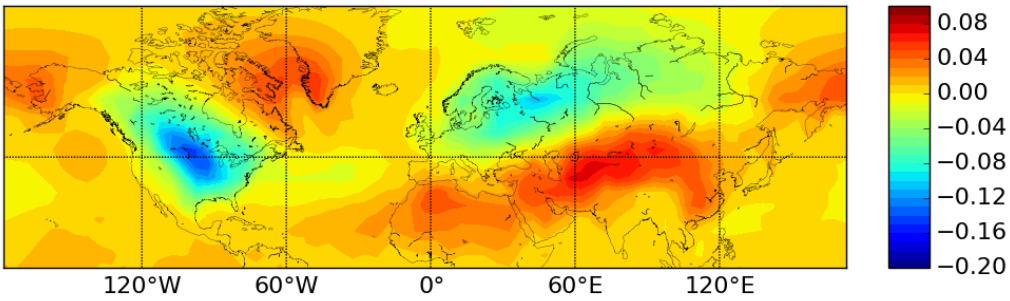

(c) Third EOF of SST of Northern Hemisphere. Explains 6.5% of the total variance.

Figure 4.6: First three EOFs of SST over Northern Hemisphere. Combined they explain 30.6% of the total variance of the field. Notice the similarities between the third EOF and the teleconnection with SST.

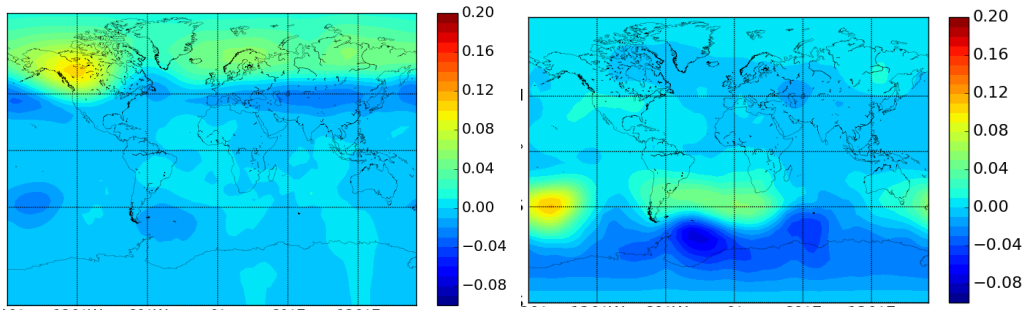

(a) First EOF of SLP. Explains 14.6% *of the total variance.*

(b) Second EOF of SLP. Explains 7.7% *of the total variance.*

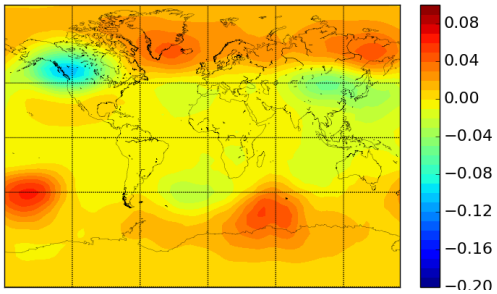

(c) Third EOF of SLP. Explains 6.8% *of the total variance.*

Figure 4.7: First three EOFs of the SLP for the entire world. Combined they explain 29.1% *of the total variance of the field.*

4.3 Bayes factors

This section presents the results of the model and predictor comparisons as described in Section 2.1.4. The percentages shown are the fraction of belief in that model compared to the rest of the models in the row. For instance in Table 4.1 the first row shows that there is an overwhelming belief in the Full Poisson compared to the Fixed. The ratio between the percentages is > 100 indicating very strong evidence in favor of the Full model according to Table 2.2. For the Generalized Pareto the belief is strongest in the Fixed model. Compared to the Full model there is a factor 10, which corresponds to a positive evidence, but compared to the Fixed ξ there is only a factor 2 corresponding to very little evidence.
It should be kept in mind that there needs to be a factor of three or more between the Bayes factors before it can be considered positive evidence.
In this table there are predictors based on EOFs which have not yet been presented. They are defined in Section 5.1.

	Poisson		Generalized Pareto			
	Full	Fixed	Full	Fixed σ	Fixed ξ	Fixed
Ballum						
NAO	100%	0%	5%	22%	23%	50%
Global Temperature	30%	70%	7%	17%	26%	50%
3rd PC of SST	100%	0%	1%	29%	10%	60%
EOF of SLP – NA	25%	75%	6%	16%	10%	68%
EOF of SST – NA	27%	73%	10%	19%	11%	60%
Esbjerg						
NAO	100%	0%	18%	7%	36%	39%
Global Temperature	1%	99%	12%	14%	21%	53%
3rd PC of SST	100%	0%	5%	44%	12%	39%
EOF of SLP – NA	10%	90%	7%	27%	19%	47%
EOF of SST – NA	22%	78%	23%	6%	30%	41%
Havneby						
NAO	99%	1%	52%	15%	5%	28%
Global Temperature	7%	93%	5%	1%	18%	76%
3rd PC of SST	100%	0%	12%	19%	40%	29%
EOF of SLP – NA	36%	64%	6%	5%	18%	71%
EOF of SST – NA	24%	76%	4%	21%	21%	54%
Højer						
NAO	38%	62%	20%	10%	43%	27%
Global Temperature	38%	62%	10%	15%	21%	54%
3rd PC of SST	100%	0%	2%	13%	45%	40%
EOF of SLP – NA	14%	86%	11%	29%	15%	45%
EOF of SST – NA	28%	72%	9%	18%	15%	58%

Table 4.1: Relative Bayes factors for the different stations for the different models. The green cell is the model with the highest Bayes factor.

	NAO	Global Temperature	3rd PC of SST	EOF of SLP – NA	EOF of SST –NA	Noise
Ballum						
Poisson	0%	0%	100%	0%	0%	0%
GPD	12%	13%	5%	14%	25%	31%
Esbjerg						
Poisson	0%	0%	100%	0%	0%	0%
GPD	25%	18%	9%	3%	35%	10%
Havneby						
Poisson	0%	0%	100%	0%	0%	0%
GPD	69%	2%	13%	2%	3%	11%
Højer						
Poisson	0%	0%	100%	0%	0%	0%
GPD	51%	11%	3%	17%	8%	10%

Table 4.2: Relative Bayes factors for the different stations for the different predictors. For all predictors the "Full Model" is used for both the GPD and the POI. There is a slight bias in favor of SLP. The green cell is the predictor with the highest Bayes factor.

4.4 Parameter estimates

In Table 4.3 the predictor coefficient-estimates from the MCMC-simulations are presented. They are calculated as the mean of the ensemble. Only the results from the Full models are included, since the other models are submodels, and thus included if the coefficients is 0. The coefficients where the 66%-CI does not span 0 are highlighted in colors to indicate that there is significant evidence in favor of the sign of the coefficient.
In Figure 4.8 a cornerplot of the parameters are shown. This shows the combined Markov Chains and their distributions projected unto 2D-spaces pairwise. Only two plots are shown out of the 40 possible, since they represent the two categories that the plots fall into: Strongly correlated parameters and weakly correlated parameters.
If the parameters were independent we would assume them to be normally distributed around their mean and thus generate circular contour plots. If the contour plots are eliptic it means that there is a non-zero covariance betweenn the two parameters. There is a general tendency in the plots that the more belief there is in a certain station-predictor combination, the less covariance there is between the parameters. This is clear in the referenced figure where there is approximately double the belief in NAO compared to the Global Temperatue and the contours for the Global Temperatue is more elliptic.

	a_λ	a_σ	a_ξ
Ballum			
NAO	2.80×10^{-1}	7.76×10^{-2}	1.12×10^{-3}
Global Temperature	7.47×10^{-2}	-8.08×10^{-2}	-2.69×10^{-2}
3rd PC of SST	-4.21×10^{-2}	3.54×10^{-3}	4.67×10^{-3}
EOF of SLP – NA	9.73×10^{-2}	-1.10×10^{-1}	1.82×10^{-1}
EOF of SST – NA	-1.81×10^{-2}	1.77×10^{-2}	-2.39×10^{-2}
Esbjerg			
NAO	2.78×10^{-1}	9.81×10^{-2}	1.84×10^{-2}
Global Temperature	-7.57×10^{-2}	5.86×10^{-1}	-3.91×10^{-1}
3rd PC of SST	-4.17×10^{-2}	5.88×10^{-3}	-3.33×10^{-3}
EOF of SLP – NA	2.96×10^{-1}	6.94×10^{-3}	9.82×10^{-3}
EOF of SST – NA	-1.43×10^{-1}	3.72×10^{-1}	-2.08×10^{-1}
Havneby			
NAO	2.56×10^{-1}	2.78×10^{-2}	-1.04×10^{-1}
Global Temperature	7.72×10^{-2}	1.88×10^{-1}	-1.89×10^{-1}
3rd PC of SST	-4.20×10^{-2}	-9.95×10^{-3}	1.22×10^{-3}
EOF of SLP – NA	2.27×10^{-1}	-2.52×10^{-1}	2.83×10^{-1}
EOF of SST – NA	-6.71×10^{-2}	1.71×10^{-2}	-1.25×10^{-2}
Højer			
NAO	2.13×10^{-1}	1.96×10^{-1}	-4.83×10^{-2}
Global Temperature	2.64×10^{-1}	1.30×10^{-1}	-1.34×10^{-1}
3rd PC of SST	-3.75×10^{-2}	-9.70×10^{-3}	1.76×10^{-3}
EOF of SLP – NA	1.15×10^{-1}	4.77×10^{-2}	1.70×10^{-1}
EOF of SST – NA	8.25×10^{-2}	1.27×10^{-1}	-8.24×10^{-2}

Table 4.3: The mean parameter values of the MCMC-chains. Only the time dependent parameters for the "Full" models are shown. Cells marked with a color indicates that the 66%-confidence interval did not span over 0, strengthening the belief in the sign of the parameter. Green indicates a positive parameter while red indicates a negative.

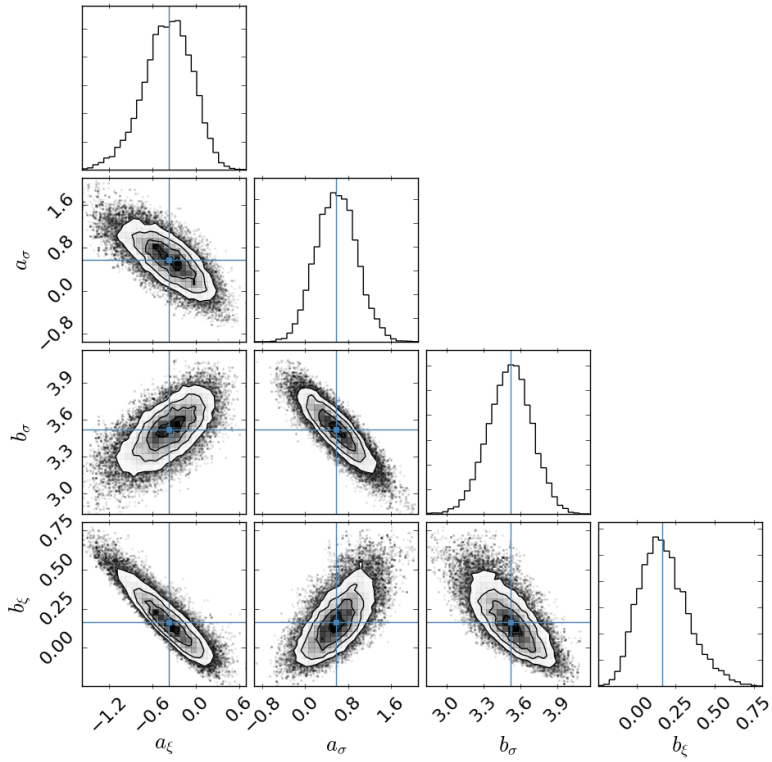

(a) The corner plot for Esbjerg and the Global Temperature.

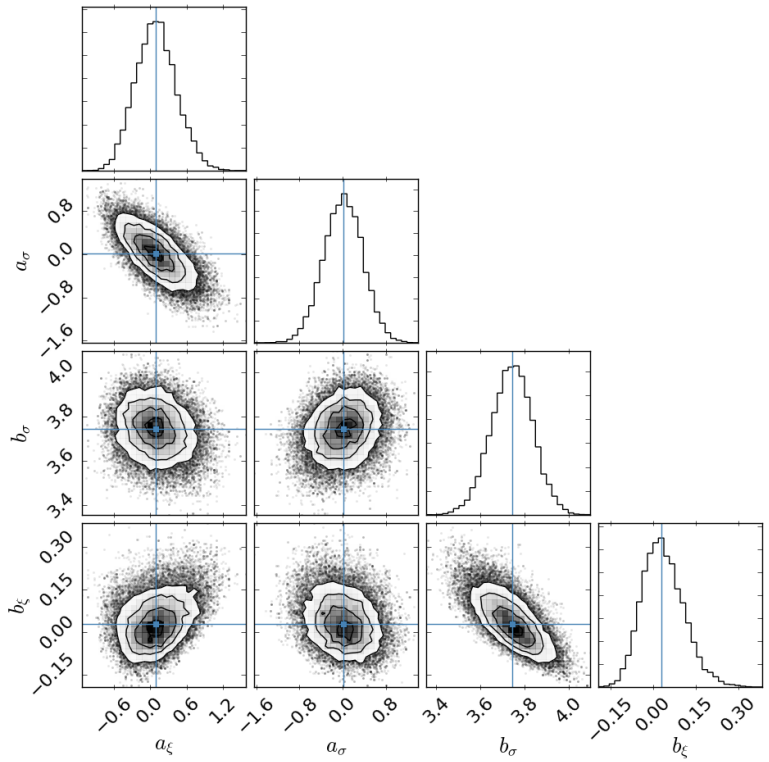

(b) The corner plot for Esbjerg and NAO.

Figure 4.8: Corner plots illustrating the distribution and covariance of the GPD-parameters for Esbjerg and the two predictors. The general trend is that the stronger the predictor the smaller the covariance of the parameters.

Chapter 5

Analysis

5.1 Teleconnection and EOF

From looking at the teleconnection analysis it becomes clear that there is a correlation between a_σ and a_ξ. Their respective teleconnection plots, for instance Figures 4.4b and 4.4c, are near complete opposites, and the same pattern is seen for all other stations and for SST-teleconnection too. This indicates that as one rise, we would expect the other to drop.
Figure 5.1 is an excerpt from Figure 4.6c and Figure 4.1 that show the resemblance between the teleconnection pattern of Ballum against the SST and the 3rd EOF of the SST. As can be seen there is some striking similarities but they are inverted. A high/low over the North Americas stretching into the Atlantic over to Europe and reaches Western Russia. From the Caribbean a low/high reaches across the Atlantic to Northern Africa and stretches east to the Middle East. Similarly a low/high hovers over the North Atlantic Ocean. There is no obvious correlation outside of these regions.
For the teleconnection there is a clear low over Iceland and a high over the Azores. This corresponds with the two centers of the NAO and hints that this might also be an influential predictor, although it is not clear why it shows up in the teleconnection of the SST and not the SLP.
The patterns described above is also described in Osborn [2010] where it is identified as "the slope of the least-squares simple linear regression between monthly CRUTEM3 land/HadSST2 sea temperature anomalies and monthly NAO index values (normalised to have unit variance)." So it may just be an indirect NAO-index, which would show in later analysis if true.
In Figure 5.2 the principal component of the third EOF can be seen.

From Figure 4.4 it can be seen that both $\hat{a}_\lambda$ and $\hat{a}_\sigma$ has has some correlation with the South Atlantic/South America/South Eastern Pacific and some over the Northern Indian Ocean.
The correlation in the Atlantic and South America can be seen to be similar to the pattern shown in Figure 4.3, and is thus assumed to just be seasonal residual.
The link to the Indian ocean is not intuitively meaningful, but as shown in Li et al. [2010] there is a correlation between the warming of the Indian Ocean and an increase in the Arctic Oscillation which is closely related to the NAO. Due to the missing physical link these areas are considered to be proxies for the NAO in

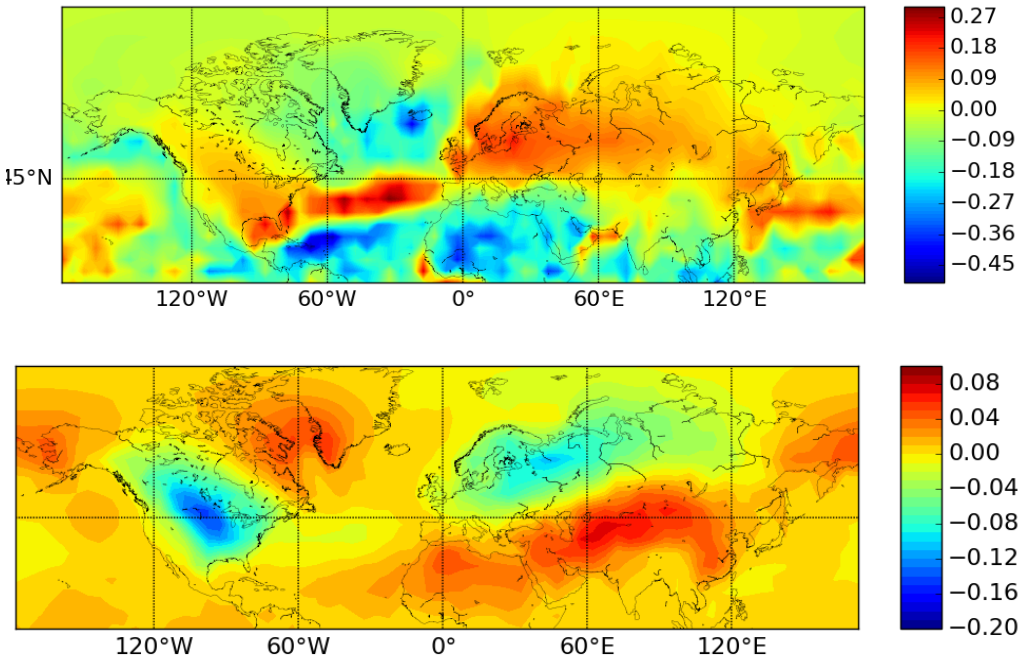

Figure 5.1: Comparison of the teleconnection for the SST and the 3rd EOF of the SST over the Northern Hemisphere.

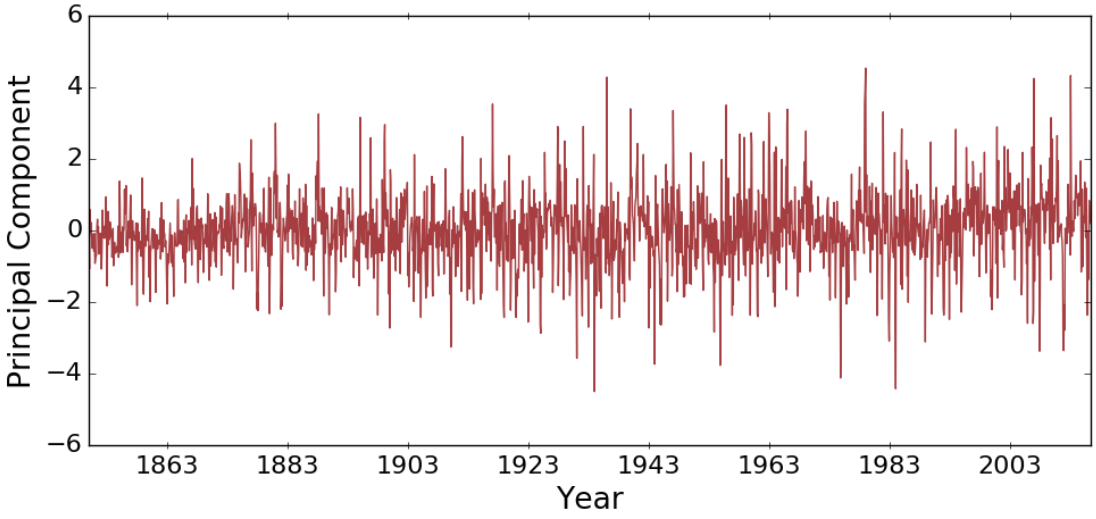

Figure 5.2: Principal component of the 3rd EOF of the SST.

this study.
There are no indications of correlation with the North Atlantic SLP, which is the source of the NAO. As the teleconnection of the SST indicates that the NAO has a significant influence, this is surprising, and draws into question whether the anomalisation of the SLP was a meaningful approach. The influence of SLP might be more based on gradients than absolute values. That would explain why no significant pattern shows up on the teleconnection.
The North Atlantic is the nearest ocean and as such the climate in and around this ocean is assumed to be the most important. Therefore a mean of the first 80% EOFs of the area $N40^\circ - N70^\circ$ and $W60^\circ - E20^\circ$ is also suggested as a predictor (EOF of SST – NA and EOF of SLP – NA).

5.2 Bayes factors

In Table 4.1 the relational Bayes factors for the different station/predictor combinations and models can be seen. It is important to emphasize that the values are with a rather large uncertainty and should mostly be interpreted in relation to each other. Referring to Table 2.2 there need to be a factor of at least 3 between the percentages before any evidence for or against the models can be inferred. Never the less the table can be read as a hint of how much faith should be put into one model over another, and patterns that are consistent between the stations can also be identified.
For the Poisson distribution there is a clear pattern where the NAO and the 3rd PC of SST (3PC) has very strong evidence in support of the Full model. One exception is for Højer where there is minor evidence against the Full model for NAO. It is not immediately clear why Højer stands out, though it might be caused by the move of the tide gauge as described in Section 3.1 but it seems odd that the other Poisson Bayes factors fall in line with the rest if that was the case. The rest of the values for Højer does not deviate significantly. For the Global temperature as a predictor there is about a factor of 4–10 in favor of the fixed model. Using Table 2.2 as reference this means a positive belief in the stationary model. As such the global temperature is not considered to be a good predictor for the frequency of events, at least given the exponential relation used. The factors for the two EOFs are around 2–4 in favor of the Fixed model. This is too little to claim any evidence in favor of either model.
For the Generalized Pareto almost all of the combination have evidence in favor of a Fixed model over the Full model, though it is minor for NAO.
The 3rd PC of SST does have minor evidence for partially Fixed models over the fully Fixed for 3 of the stations indicating some dependence on this predictor though not strong.
This reinforces the belief that the NAO is an important driver for the storm surges. Since the 3rd PC of SST was identified as a proxy for the NAO it makes good sense that these two predictors are linked together and both show similar predictive skills.
In Table 4.2 the belief in each predictor is compared. It is the Full model that has been used for the comparison, which should not induce a bias as the Fixed models are nested in the Full models.

The factors have all been calculated against the Global Temperature before converting them into fractions. This resulted in a problem with the SLP because it ends 3 years before the water level measurements, and thus the Global Temperature had to be cut off there too to make a fair comparison, thus decreasing the number of fitted events. This likely introduce a bias in favor of the SLP-based model but as it does not dominate any of the comparisons this does not make a difference to the conclusions.
For the Poisson-distribution there is very strong evidence in favor of the 3rd PC of SST over all other predictors, indicating a very strong correlation. Interestingly enough it is also heavily favored over the NAO of which it was assumed to be a proxy. This could indicate that it should be viewed as an independent pattern.
For the GPD the NAO is favored by Havneby and Højer, while Esbjerg have no strong evidence in favor of any one predictor. In Ballum the evidence is in favor of the noise which highly reduces the credibility of the predictors for this station. The 3rd PC of SST scores consequently low in this comparison compared to the NAO further enforcing the belief that it might be considered a separate pattern.
In summary of this section it is concluded that the only predictors that have evidence in favor of a non-stationary model are the NAO and the 3rd PC of SST, the rest have minor to strong evidence in favor of a Fixed model. When comparing the Bayes factors of the different predictors there is very strong evidence in favor of the 3rd PC of SST for the Poisson distribution, while there is no clear preferred predictor for the GPD although a positive evidence in favor of the NAO can be seen for Havneby and Højer.

5.3 Parameter estimates

In Table 4.3 the parameter estimates of the predictor coefficients can be seen. The magnitude of the parameters should not be compared in between the models as they have not been normalized in any way and as such of different magnitudes themselves, see Section 3.3.1 for the range of the predictors.
Therefore it is the sign of the parameters that are highlighted in the table. The estimates where the sign is consistent within the 66%-CI is colored with green for positive and red for negative to clearly present the important patterns.
The only two estimates that retain their color over all 4 stations is a_λ for the NAO and the 3rd PC of SST. This is not surprising as it was shown in the previous section that there was significant evidence for these two predictors for the Poisson distribution. It should also be noted that they have opposite sign of each other.
The same two predictors are also the only ones to have the same sign between two stations for the a_σ, and also here are they opposites. For a_ξ there is no station/predictor-combination that has a consistent sign between stations. This could be interpreted as though there is a general tendency towards a ξ that is not dependent on the predictors. Referring back to Table 4.1 this is emphasized by the Fixed ξ model being the second most "popular" after the Fixed model.
It should also be noticed that for all estimates where the sign of a_σ an/or a_ξ is significant, the corresponding a_ξ/a_σ is of opposite sign. This indicates that the two parameters are not independent of each other. This is further emphasized by Figure 4.8, where the covariance between a_σ and a_ξ is clearly visualized in the

elliptic contour plots.
This is could be due to the constraint $\mu \leq x \leq \mu - \frac{\sigma}{\xi}$ for $\xi < 0$ which makes some areas of the parameter space invalid. If this were the case we would expect that there was a stronger covariance for predictors where $\xi < 0$ most of the time since they would be more constrained. Referring back to Figure 4.8 again it is clear that the estimates are much more interdependent for the Global Temperature that for the NAO, but the calculating the percentages of $\xi < 0$ it is 20% for the Global Temperature and 50% for the NAO, making it unlikely that it is the constraint that is the cause of the covariance. It is unclear what the cause could otherwise be.
In summary of this section there is a clear evidence for the sign of a_λ for NAO and the 3rd PC of SST, and to a lesser degree for a_σ for the same two predictors. None of the other predictors showed consistency in the sign of the estimates. This reinforces the conclusion from the last section that the NAO and the 3rd PC of SST are the best predictors for the Poisson.
It was also found that the a_σ and a_ξ are likely to have opposite signs, though the cause of this interdependence is unknown.

5.4 Predictions

Both Section 5.2 and 5.3 present arguments for NAO and the 3rd PC of SST as the best predictors. There is strong evidence for a non-stationary model for the Poisson-distribution while there is weak evidence for a stationary model for the GPD.
In other words there is found no evidence in favor of the distribution of extreme storm surges changing over the last four decades. There is strong evidence for a change in the rate at which these surges happen.
The goal of this section is to extrapolate these changes to make inference about how the statistics of storm surges might change in the near future. The procedure is explained for the 3rd PC of SST.

As seen in Figure 5.3 there is a minor positive linear trend in the predictor and the variance rises dramatically at the turn of the 20th century. If we assume that there is a positive trend what would that mean for the distribution extreme events?
The coefficient a_λ is negative meaning that the rate at which the events happen would fall as the predictor rises. As the rate falls so would the return level described in Section 2.1.5.
There is some evidence that the sign of a_σ is negative. This would mean σ would fall as the predictor rises. Referring back to Figure 2.1 and the constraints of the GPD this would make the tail of the distribution shorter, shifting the dis-

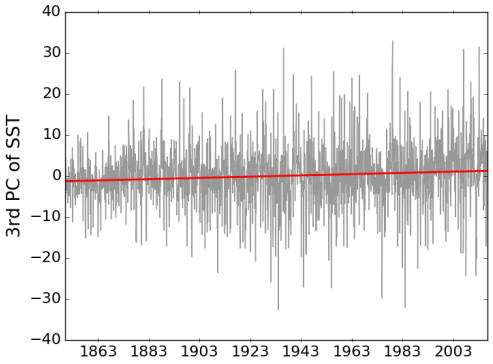

Figure 5.3: The 3rd PC of SST and a fitted linear trend with $a = 5 \times 10^{-4}\,\mathrm{yr}^{-1}$

tribution towards the threshold. This would decrease the return level for high return periods, as these events become more rare.
There is no evidence for the sign of a_ξ, but as pointed out in Section 5.3 there is a tendency towards it have the opposite sign of a_σ, which would be positive. This would mean a rising ξ opposing the effects of a falling σ. From this analysis it is not clear what will happen as the predictor changes, except that the rate of exceedances should fall.

5.4.1 Extrapolation

To do inference about the future the CMIP5-ensemble[1] is used. Each model in the ensemble runs from 1850 to 2100 with a temporal resolution of 1 month. They each consist of a $5°$ by $5°$-grid which can be projected unto the EOFs of the original grid and then used to extrapolate, as briefly described in Section 2.3. On the figure to the right the 3rd PC of the SST is extrapolated using the CMIP5 ensemble. Notice that the line has been smoothed as only trends on a larger scale are desired in the prediction.

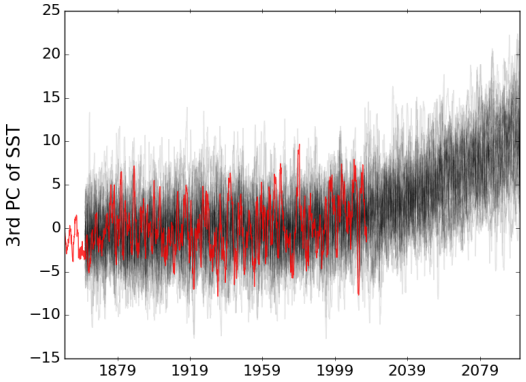

Figure 5.4: The 3rd PC of SST extrapolated with the CMIP5-ensemble for the RCP8.5 scenario.

The CMIP5 works with 4 different scenarios: The RCP2.6, RCP4.5, RCP6 and RCP8.5. Each corresponding to a different scenarios the concentration of greenhouse gases in our atmosphere, the latter being the one with the highest concentration and vice versa.
The extrapolation is calculated like this:

- Pick a scenario
- For each member of the ensemble of models corresponding to the chosen scenario calculate a smoothed time series by projecting the member unto the EOFs, and manipulating the resulting PCs
- For each member calculate a time series of ensembles of return levels for a given return period and convert them into ratios with basis in the return level January 2000
- Pool the ratio together and calculate the median, quartiles and 66%-CI.

The reason that ratios are used instead of absolute values are the inherent uncertainty in the climate models. By reducing the predictions to ratios it is easier to compare the extremes of the models as they are scaled to the same magnitude.
The results can be seen in Figure **??** and **??**.
As can be seen there is a general trend towards a fall in the return level for both

[1]The CMIP5 is a ensemble of climate models from different climate modelling groups around the world.

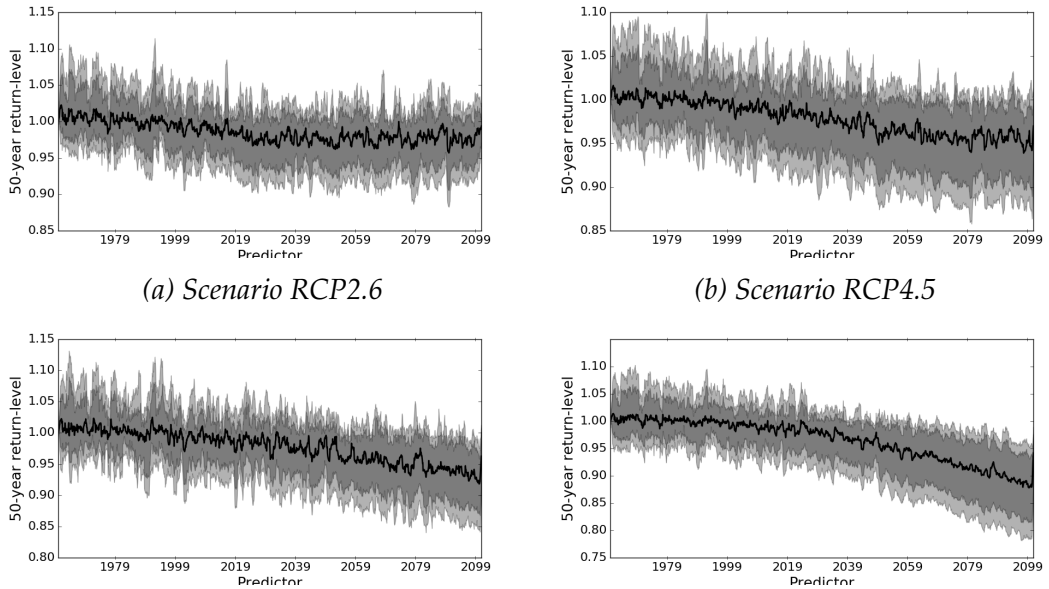

(a) Scenario RCP2.6 *(b) Scenario RCP4.5*

(c) Scenario RCP6 *(d) Scenario RCP8.5*

Figure 5.5: Prediction of the 50-year return level presented as a ratio of the 50-year return level in 2000. The station is Esbjerg and the predictor is the 3rd PC of SST

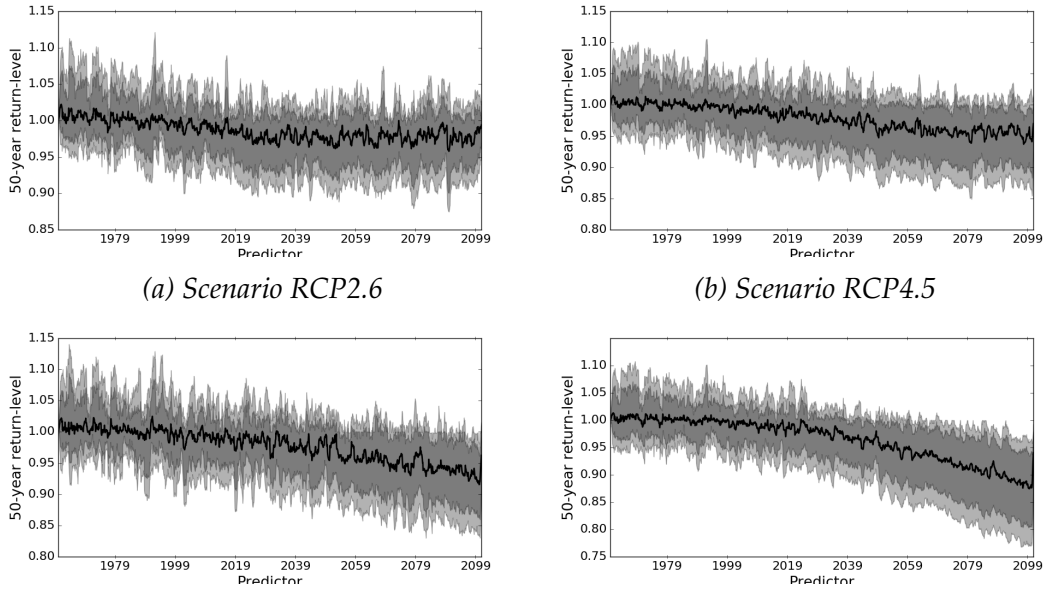

(a) Scenario RCP2.6 *(b) Scenario RCP4.5*

(c) Scenario RCP6 *(d) Scenario RCP8.5*

Figure 5.6: Prediction of the 100-year return level presented as a ratio of the 100-year return level in 2000. The station is Esbjerg and the predictor is the 3rd PC of SST

the 50-year and the 100-year return level, though the fall is greatest for the 50-year return level.The general fall in the return level is well in line with the conclusion from the previous section, where it was predicted that the overall rate of storm surges would be falling. The bigger fall in the 50-year return level could indicate that the ξ-dependence is dominant over the σ, but the difference is not large enough to conclude anything with statistical significance.
The plots are very similar for the other stations which can be found in Appendix A.2. It is problematic that the data only has a time span of 40 years, as that makes the extrapolation less credible.

Chapter 6

Conclusion

In the previous chapters the validity of a range of different models and predictors have been assessed. There has been no conclusive evidence in favor of a non-stationary model for the GPD, though some station/predictor combinations showed weak evidence against the completely stationary model. This indicates that the distribution of storm surges is not expected to change significantly in the immediate future, though it is possible that the right model/predictor combination simply has not been found.

Very strong evidence has been found in favor of the non-stationary model for the rate parameter λ. Especially the 3rd PC of SST was found to outperform the other predictors overwhelmingly. The 3rd PC of SST was initially identified as "the slope of the least-squares simple linear regression between monthly CRUTEM3 land/HadSST2 sea temperature anomalies and monthly NAO index values (normalised to have unit variance)" [Osborn, 2010], but since there was very strong evidence in favor of it over the NAO it seems likely that it contains additional unidentified climatic information with great significance for the storm surge rate. This could be worth investigating in subsequent research.

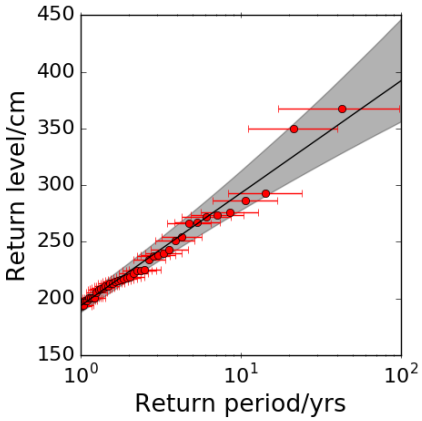

Figure 6.1: The empirical return level for Esbjerg (red dots) and a fitted stationary GPD (black line). The error bars and the shaded grey area indicates 66%-CI.

The coefficient of a_λ for the 3rd PC of SST was found to be negative with a 33% significance level. Extrapolation of the predictor shows that it is expected to rise over the next century thus resulting in an expected lower rate of storm surges.
In Figure 5.5 and 5.6 it is shown that this would result in a drop in the expected 50- and 100-year return levels, much in line with the predictions in Menéndez and Woodworth [2010].
One should be vary though as the time period used for calibration of the model is relatively short compared to the extrapolated period. It is a good indicator that none of the ratios diverge making the result more believable, but a goal for further research would be to acquire longer data series.

Scenario	MSLR
RCP2.6	$0.4\,\mathrm{m}$
RCP4.5	$0.5\,\mathrm{m}$
RCP6	$0.6\,\mathrm{m}$
RCP8.5	$0.7\,\mathrm{m}$

Table 6.1: Best guess projected relative sea level rise between 1986–2005 and 2081–2100 in the German Blight. Source: [IPCC, 2013]

Referring back to the naive prediction from the introduction, that the 100-year event of today would be equal to the 5-year event of 2100, it would be interesting to examine what influence the results of this paper would have on this prediction.
In Figure 6.1 the empirical return plot for Esbjerg with basis in 1972-2015 is presented as well as a fitted stationary model. Even though the evidence is strong in favor of a non-stationary model it looks as though the stationary approximation is reasonably good in this case.
For each return period the predicted mean ratio for the period 2081–2100 over 1986–2005 is calculated and multiplied by the corresponding fitted return levels from Figure 6.1. This is done for each of the four emission scenarios adding the corresponding predicted mean sea level rise as seen in Table 6.1. The resulting return level plot can be seen in Figure 6.2. The 100-year return event for today is roughly equal to the 45-year event of 2100 for all scenarios, a much less gruesome prediction than the naive prediction of a return period of 5 years.
The high emission scenarios seems to have slightly higher return levels for the short return periods than the low emission scenarios. Conversely the low emission scenarios have higher expected return level for the longer return periods although the uncertainty makes it impossible to draw any conclusion.
Interestingly this is in line with the findings from Figure 1.2, where the post-1994 period has more low storm surges but fewer high ones than the pre-1994 period. So according to the findings in this paper there is some evidence that, in spite of the higher projected sea level rise, the high emission scenario might be the one that results in the lowest extreme surges, making the predicted future for the West Coast of Denmark a little less bleak.

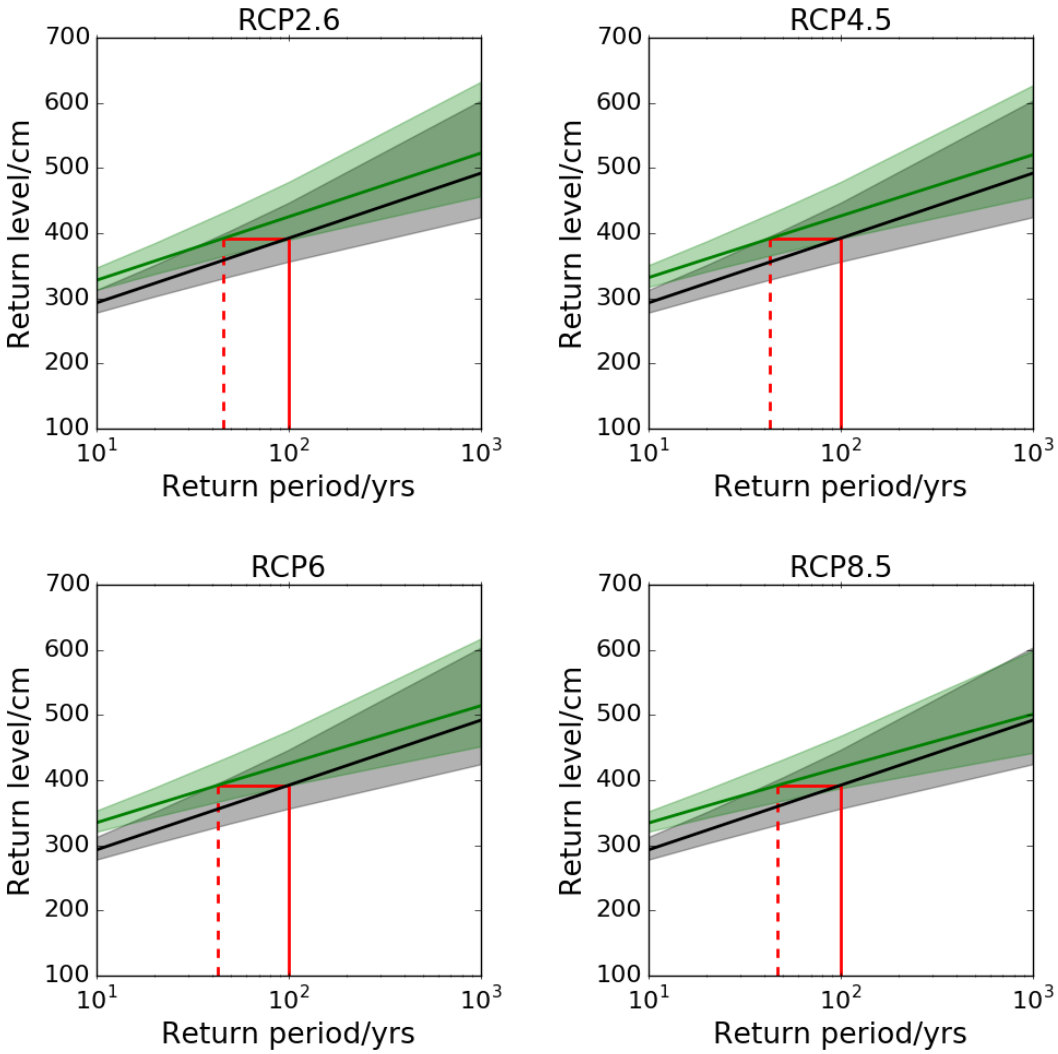

Figure 6.2: Predictions of the return level in Esbjerg for the four different scenarios. The black line is the baseline return levels from 1986–2005. The green line is predicted return levels for the period 2081–2100. The red solid line indicates the 100-year return level of today and the red dashed line indicates the expected return period for the same event in 2100.

Bibliography

Rob Allan and Tara Ansell. A new globally complete monthly historical gridded mean sea level pressure dataset (HadSLP2): 1850–2004. *Journal of Climate*, 19 (22):5816–5842, nov 2006.

Anthony G. Barnston and Robert E. Livezey. Classification, seasonality and persistence of low-frequency atmospheric circulation patterns. *Monthly Weather Review*, 115(6):1083–1126, jun 1987.

Jens Hesselbjerg Christensen, Karsten Arnbjerg-Nielsen, Aslak Grinsted, Kirsten Halsnæs, Erik Jeppesen, Henrik Madsen, Jørgen Eivind Olesen, John Roy Porter, Jens Christian Refsgaard, and Martin Olesen. *Analyse af IPCC delrapport 2 - Effekter, klimatilpasning og sårbarhed: med særligt fokus på Danmark*. Miljøministeriet, Naturstyrelsen, 2014.

Stuart Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer London, 2001.

Kevin Cowtan and Robert G. Way. Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends. *Quarterly Journal of the Royal Meteorological Society*, 140(683):1935–1944, feb 2014.

Wesley Ebisuzaki. A method to estimate the statistical significance of a correlation when the data are serially correlated. *Journal of Climate*, 10(9):2147–2153, sep 1997.

M.G.G. Foreman. Manual for tidal heights analysis and prediction, 1977.

William F. Hanlon. *The energy spectrum of ultra high energy cosmic rays measured by the High Resolution Fly's Eye observatory in stereoscopic mode*. ProQuest Dissertations Publishing, 2008.

Steacy Dopp Hicks. Unterstanding tides, 2006.

IPCC. *Summary for Policymakers*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2013.

Fred James. Minuit tutorial - function minimization, 2004.

Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, jun 1995.

Shuanglin Li, Judith Perlwitz, Martin P. Hoerling, and Xiaoting Chen. Opposite annular responses of the northern and southern hemispheres to indian ocean warming. *Journal of Climate*, 23(13):3720–3738, jul 2010.

Melisa Menéndez and Philip L. Woodworth. Changes in extreme high water levels based on a quasi-global tide-gauge data set. *J. Geophys. Res.*, 115(C10), oct 2010.

Colin P. Morice, John J. Kennedy, Nick A. Rayner, and Phil D. Jones. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set. *Journal of Geophysical Research: Atmospheres*, 117(D8), apr 2012.

Antonio Navarra and Valeria Simoncini. *A Guide to Empirical Orthogonal Functions for Climate Data Analysis*. Springer Nature, 2010.

Timothy J. Osborn. Winter 2009/2010 temperatures and a record-breaking north atlantic oscillation index. *Weather*, 66(1):19–21, dec 2010.

Stephen H. Schneider, Terry L. Root, and Michael D. Mastrandrea. *Encyclopedia of Climate and Weather*. Oxford University Press (OUP), jan 2011.

Carlo Sørensen, Madsen Holger Toxvig, and Knudsen Søren Bjerre. *Højvandsstatistik 2012*. Kystdirektoratet, 2013.

S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, mar 1938.

List of Figures

List of Tables

Appendix A

Plots

A.1 Teleconnection-plots

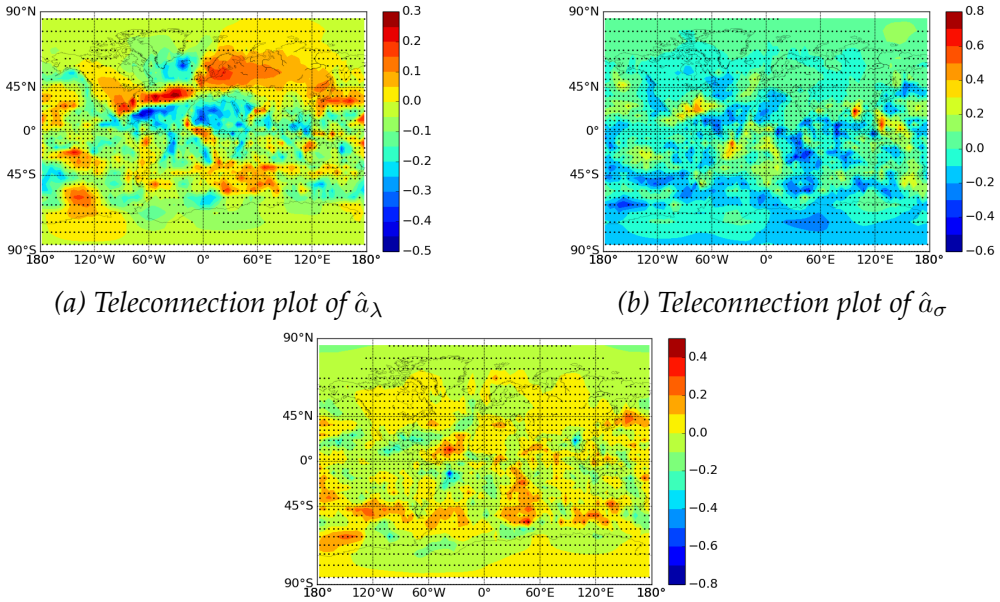

(a) Teleconnection plot of $\hat{a}_\lambda$

(b) Teleconnection plot of $\hat{a}_\sigma$

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.1: Teleconnection plots of the MLE of the predictor coefficients for Ballum and the SST

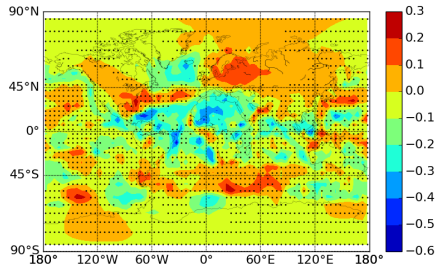

(a) Teleconnection plot of $\hat{a}_\lambda$

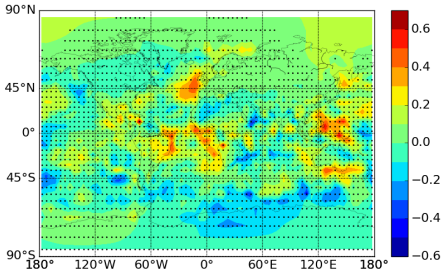

(b) Teleconnection plot of $\hat{a}_\sigma$

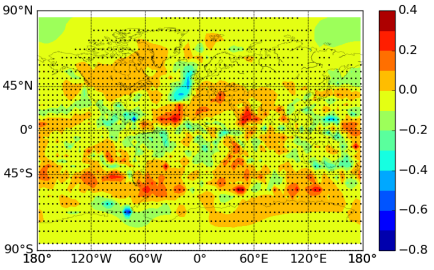

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.2: Teleconnection plots of the MLE of the predictor coefficients for Esbjerg and the SST

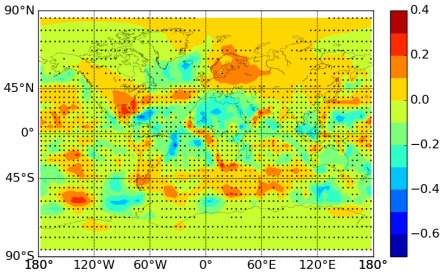

(a) Teleconnection plot of $\hat{a}_\lambda$

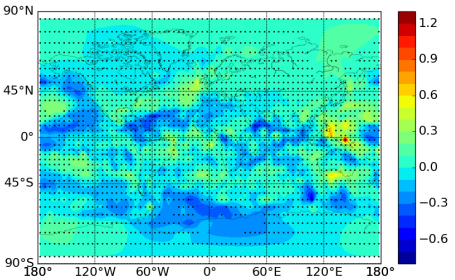

(b) Teleconnection plot of $\hat{a}_\sigma$

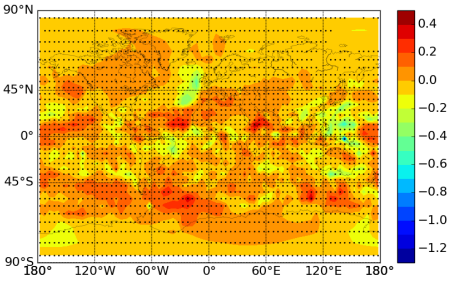

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.3: Teleconnection plots of the MLE of the predictor coefficients for Havneby and the SST

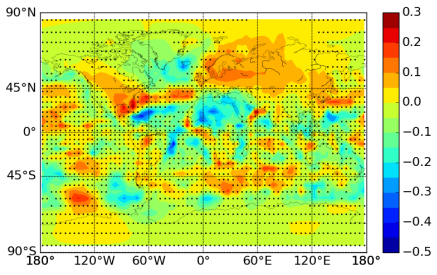

(a) Teleconnection plot of $\hat{a}_\lambda$

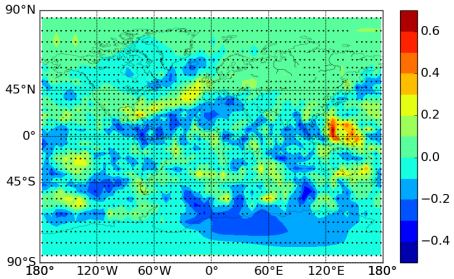

(b) Teleconnection plot of $\hat{a}_\sigma$

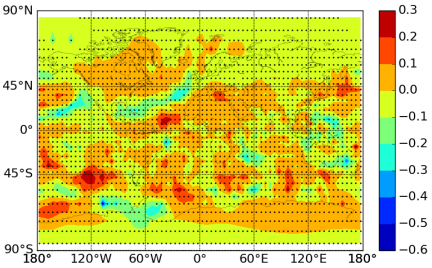

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.4: Teleconnection plots of the MLE of the predictor coefficients for Højer and the SST

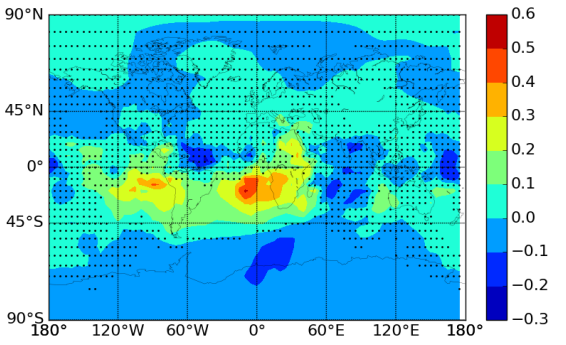

(a) Teleconnection plot of $\hat{a}_\lambda$

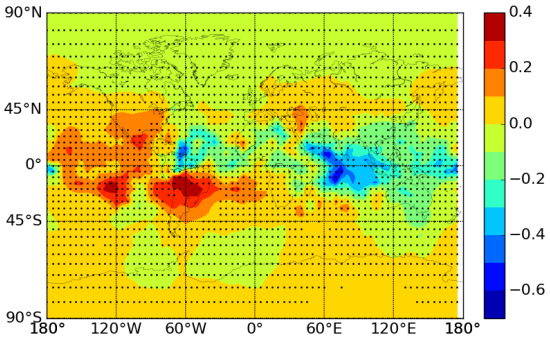

(b) Teleconnection plot of $\hat{a}_\sigma$

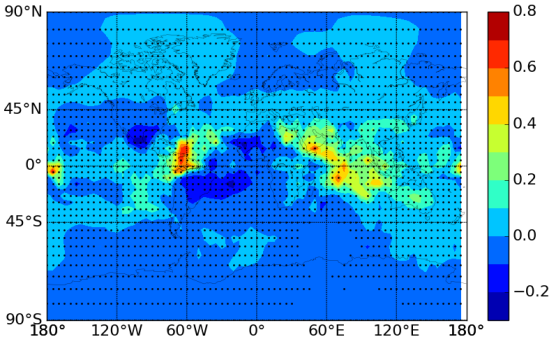

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.5: Teleconnection plots of the MLE of the predictor coefficients for Ballum and the SLP

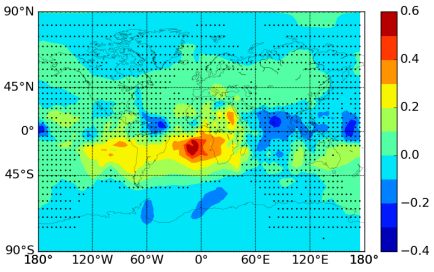

(a) Teleconnection plot of $\hat{a}_\lambda$

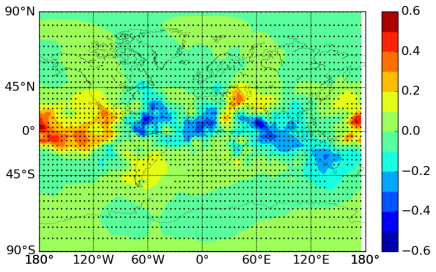

(b) Teleconnection plot of $\hat{a}_\sigma$

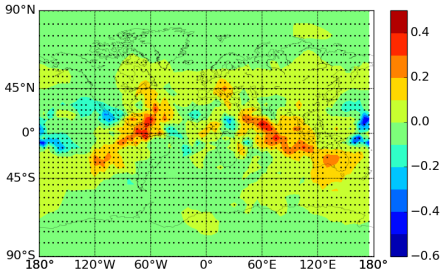

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.6: Teleconnection plots of the MLE of the predictor coefficients for Esbjerg and the SLP

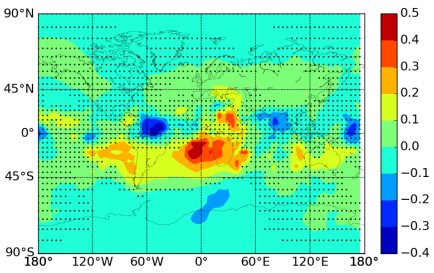

(a) Teleconnection plot of $\hat{a}_\lambda$

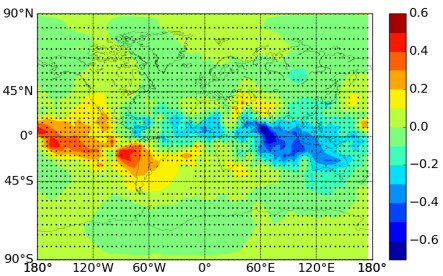

(b) Teleconnection plot of $\hat{a}_\sigma$

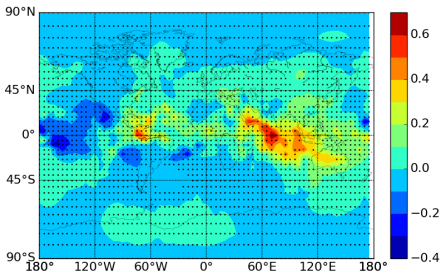

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.7: Teleconnection plots of the MLE of the predictor coefficients for Havneby and the SLP

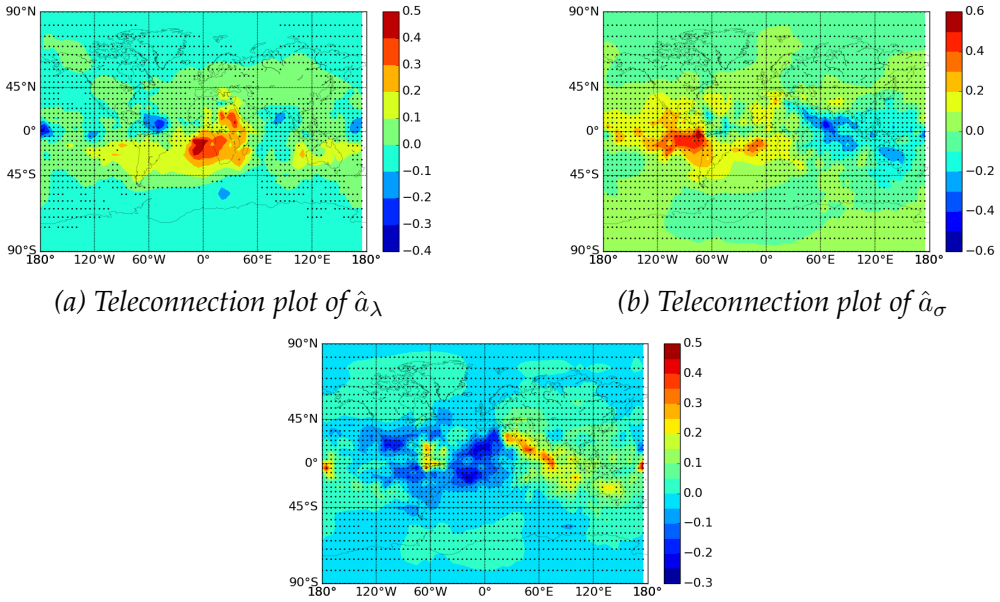

(a) Teleconnection plot of $\hat{a}_\lambda$

(b) Teleconnection plot of $\hat{a}_\sigma$

(c) Teleconnection plot of $\hat{a}_\xi$

Figure A.8: Teleconnection plots of the MLE of the predictor coefficients for Højer and the SLP

A.2 Predictions

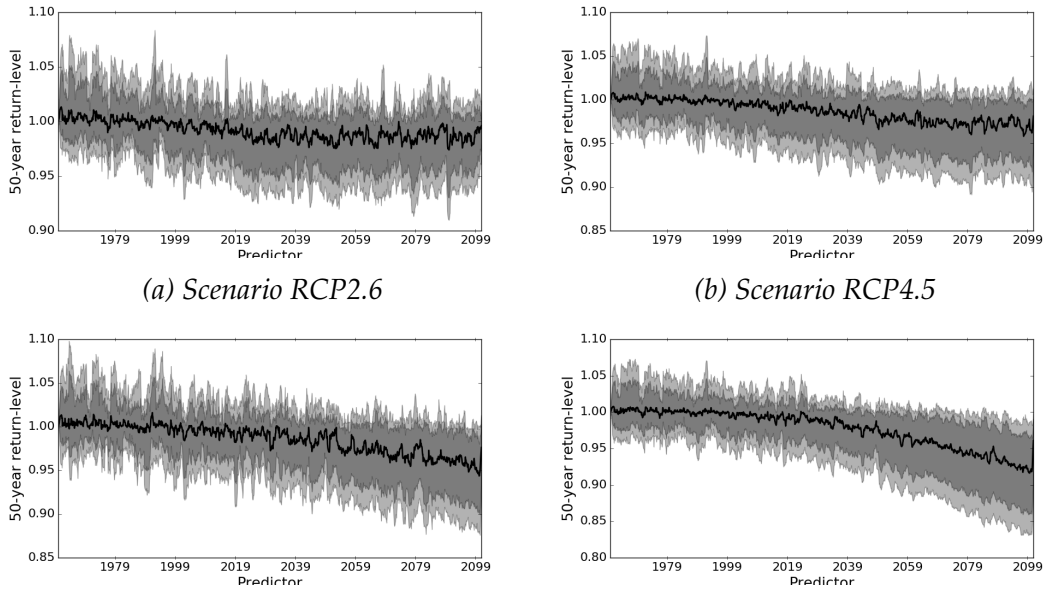

(a) Scenario RCP2.6

(b) Scenario RCP4.5

(c) Scenario RCP6

(d) Scenario RCP8.5

Figure A.9: Prediction of the 50-year return level presented as a ratio of the 50-year return level in 2000. The station is Ballum and the predictor is the 3rd PC of SST

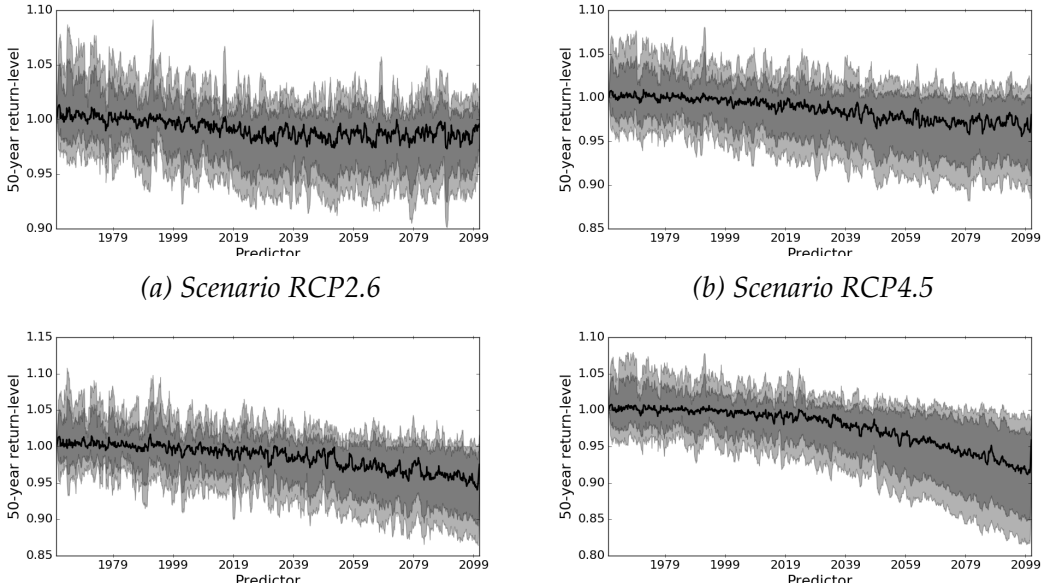

(a) Scenario RCP2.6 *(b) Scenario RCP4.5*

(c) Scenario RCP6 *(d) Scenario RCP8.5*

Figure A.10: Prediction of the 100-year return level presented as a ratio of the 100-year return level in 2000. The station is Ballum and the predictor is the 3rd PC of SST

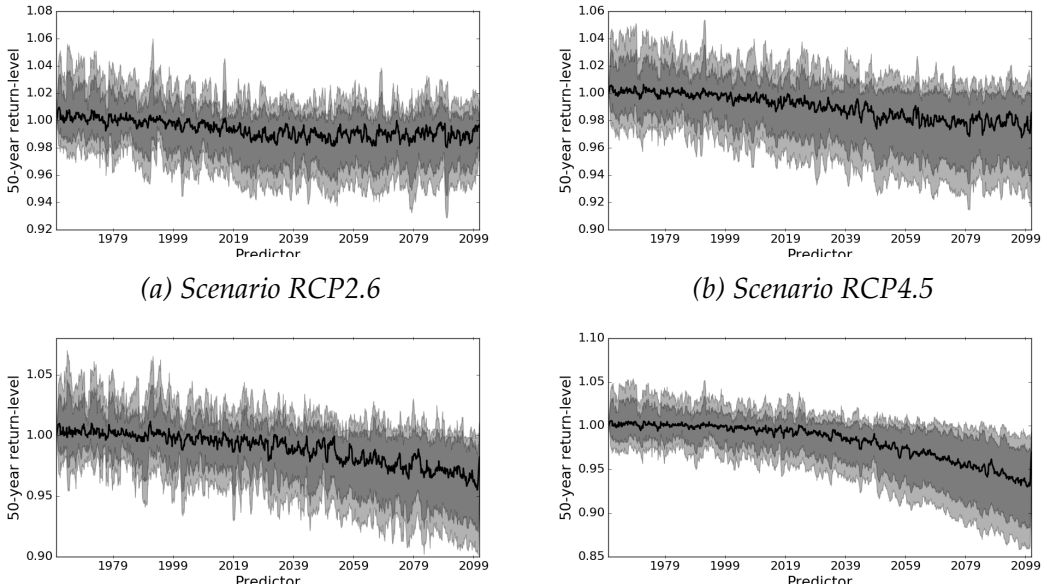

(a) Scenario RCP2.6 *(b) Scenario RCP4.5*

(c) Scenario RCP6 *(d) Scenario RCP8.5*

Figure A.11: Prediction of the 50-year return level presented as a ratio of the 50-year return level in 2000. The station is Havneby and the predictor is the 3rd PC of SST

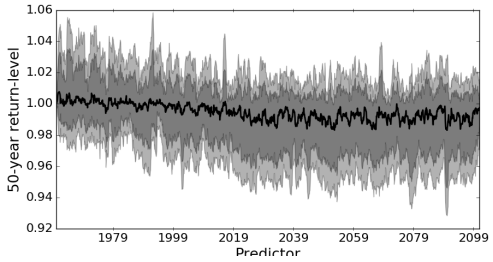

(a) Scenario RCP2.6

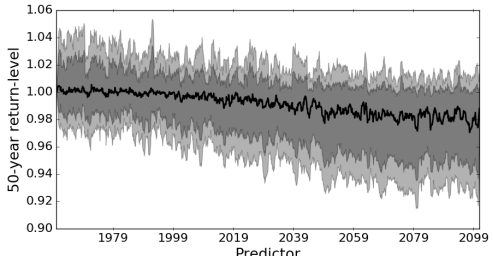

(b) Scenario RCP4.5

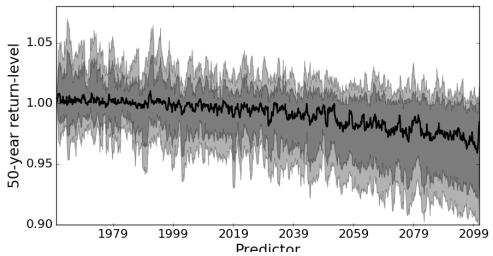

(c) Scenario RCP6

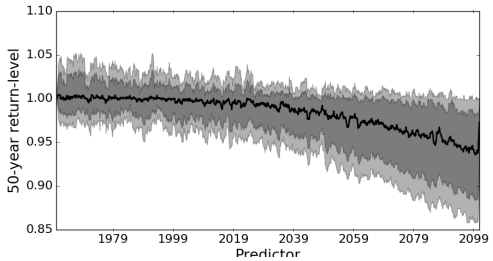

(d) Scenario RCP8.5

Figure A.12: Prediction of the 100-year return level presented as a ratio of the 100-year return level in 2000. The station is Havneby and the predictor is the 3rd PC of SST

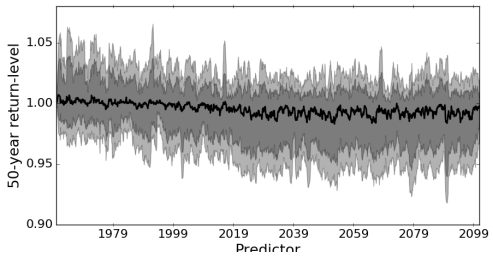

(a) Scenario RCP2.6

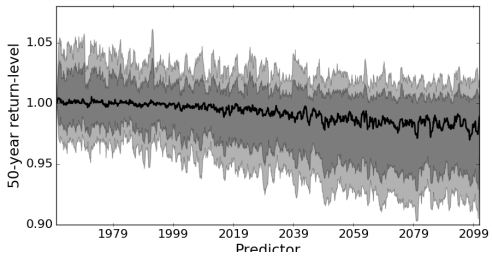

(b) Scenario RCP4.5

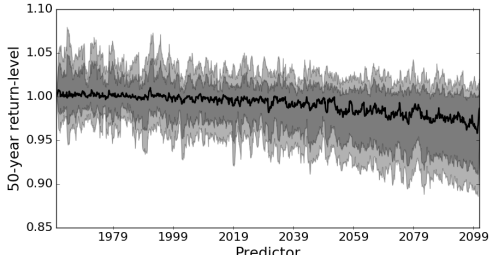

(c) Scenario RCP6

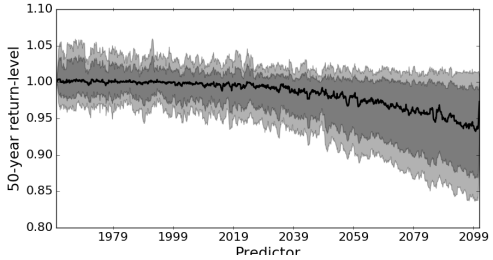

(d) Scenario RCP8.5

Figure A.13: Prediction of the 50-year return level presented as a ratio of the 50-year return level in 2000. The station is Højer and the predictor is the 3rd PC of SST

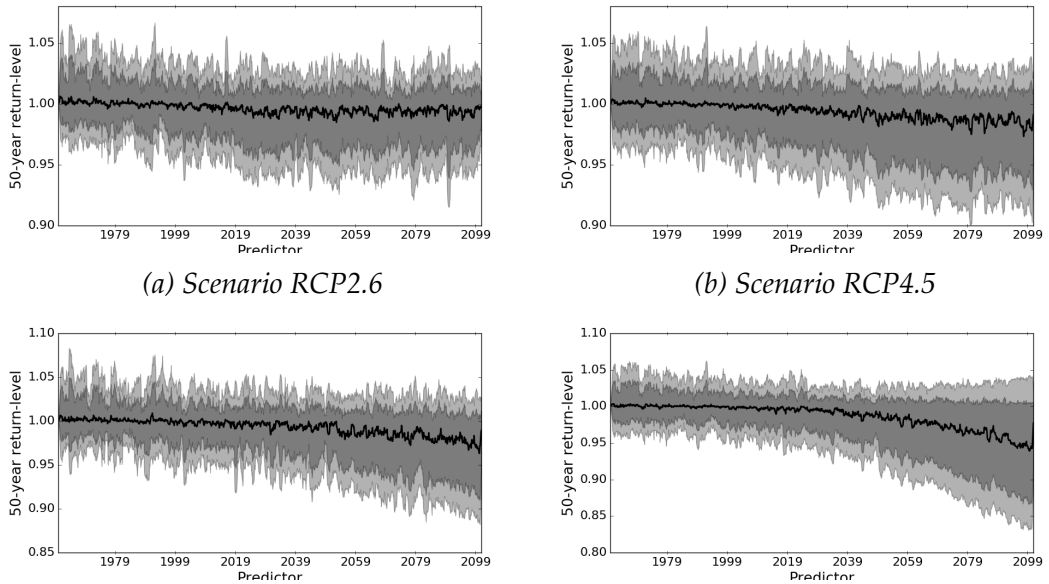

(a) Scenario RCP2.6 *(b) Scenario RCP4.5*

(c) Scenario RCP6 *(d) Scenario RCP8.5*

Figure A.14: Prediction of the 100-year return level presented as a ratio of the 100-year return level in 2000. The station is Højer and the predictor is the 3rd PC of SST