UNIVERSITY OF TROMSØ UIT

FACULTY OF SCIENCE AND TECHNOLOGY
DEPARTMENT OF PHYSICS AND TECHNOLOGY

Long-range memory in Earth surface temperatures: Spatial scale dependence and land-sea differences

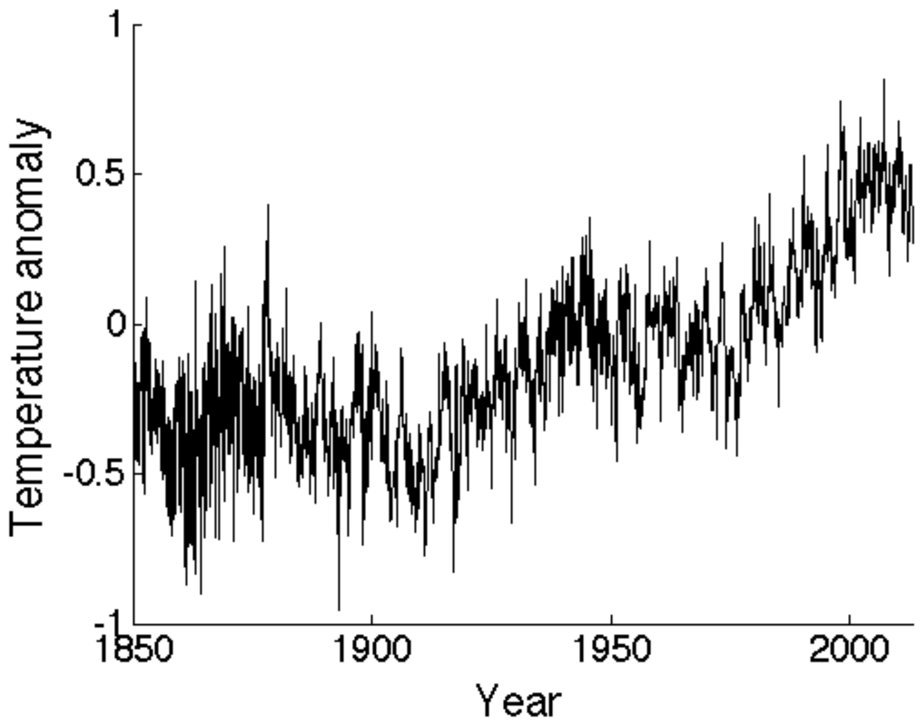

Hege-Beate Fredriksen

FYS-3900 Master's Thesis in Physics

May 2013

Abstract

The use of long-range memory models as a description of the noise in Earth surface temperatures has increased the recent years, and as a measure of the persistence for such time series we have the Hurst exponent. It is known that sea surface temperatures are more persistent than land temperatures, and that global temperatures are more persistent than local temperatures. We also know that the persistence is higher for lower latitudes than for higher latitudes. My results confirm these observations, and in addition they reveal what the Hurst exponents are for spatial scales between local and global. This is done by performing spatial averaging over gridded temperature data to obtain new time series in more coarse-grained grid boxes.

To find an explanation for the increase in Hurst exponent that is seen when increasing the spatial scale, I have studied how the autocovariance function for a large grid box depends on the spatial cross-covariances within the box. If these are strong compared to the autocovariances in that area they will have an impact on the Hurst exponent.

Scale free long-range memory models are found to give a good description for global temperature and many of the local temperatures on time scales from a few months to ten years. The largest deviations are observed in the eastern equatorial Pacific where ENSO is a very dominating process.

Acknowledgement

First of all I would like to thank my supervisor Kristoffer Rypdal for helping me find an interesting project and for good help and advices during this year. I am also grateful for getting the opportunity to go to EGU, and some smaller seminars and workshops organized by my supervisor. Meeting other scientists and presenting my work for others has been good inspiration for a further career in research.

I would also like to thank my family, friends and boyfriend for good support. A special thank to my fellow student and friend Tine for the good company at the office, and to Ola and Martin for providing user friendly functions in Mathematica for generating synthetic noise.

Contents

Chapter 1

Introduction

By studying a plot of the global temperature (like that on the front page), we can clearly see that it has been increasing since the recordings started. If we study a local temperature record, a possible trend in the temperature is not observed that clearly. It looks like it is hidden in the large fluctuations in the temperature, even though the largest temperature fluctuation of them all, the annual variation, is subtracted. So what does actually happen with the temperature when it is averaged globally?

It is also well known that the coastal climate differs from the climate in interior land areas. Near the coast we expect to have milder winters and colder summers than if we go further away from the coast. Or expressed in another way, we observe that the seasonal variations are weaker near the coast. From that we would expect the temperature time series to appear much more random if we just subtract the seasonal variation from all of them. But statistical analyses show that there are still some geographical differences remaining. What is it about the sea and the land that causes this difference, and how large is it actually? These questions are what I seek to answer in this thesis.

1.1 Earlier work

The questions discussed are related to how strong the variations in temperature on different time scales are. Recent analyses of temperature shows evidence that the strength of variation has a power law dependence on time scale (or frequency) studied, and this property is described by only one parameter, the exponent of the power law. Some work has been done to study this exponent at different time scales, and results show that we can divide the time scales in a few regimes with different power laws. It has also been observed before that this exponent is different for land and sea, that it depends on latitude, and that it is higher for globally averaged temperatures than local temperatures. But to my knowledge nobody has made a systematic study of how it is increasing with increasing spatial scale, or tried to explain why it is increasing.

1.2 Structure of my work

I will start by defining and exploring some of the properties of long-range memory (LRM) processes and how they can be modelled in chapter 2. This is then followed by a description of how to estimate the Hurst exponent that describes the LRM in chapter 3. In chapter 4 I will test these methods on synthetic data to see how well they perform and if they can be trusted. Chapter 5 is a discussion of the trends we see in the temperature, and how to distinguish between trend and LRM-noise.

Now that most of the theory and methods used are described chapter 6 gives a short description of the dataset that I will analyze. Chapter 7 describes how to treat missing data points in the time series, and how the spatial averaging is performed. It also has a discussion and an analysis of what happens to the autocovariance function when we coarse grain time series by spatial averaging over areas of increasing size. Chapters 8, 9, 10 and 11 show results of the analyses applied to the real time series. First I analyze globally and hemispherically averaged time series, and then I go to gridded time series of increasing spatial scale. The analyses are divided in only sea surface temperatures, only land temperatures, and combined surface temperatures.

Chapter 12 discusses the validity of the stochastic model we assume in different regions, and I also analyze one regional dataset to study only land temperatures in more detail. In chapter 13 the concluding discussions are given.

Chapter 2

Long-range memory processes

2.1 Time series

A time series is a sequence of random variables $\{X_t\}$, $t = 0, \pm 1, \pm 2, ...$, which could for examle be temperature measurements over some time. Two important properties of a time series are: (1) what kind of distribution we assume that the terms are realizations of, and (2) how much the values X_t at different times depend on each other, that is, the persistence of the time series. As a measure of linear dependence we have the autocovariance function:

$$\gamma(s,t) = \operatorname{cov}(x_s, x_t) = \mathbb{E}[(x_s - \mu)(x_t - \mu)] \tag{2.1}$$

where μ is the mean value of x. By analyzing temperature measurements over a long time we can see that there exists dependence on almost all time scales, and that it is strong for the shortest time scales and becomes smaller the longer time scales we go to. The decrease of the dependence with time scale is quite smooth, with no time scale much more important than the others nearby. As there is no characteristic scale is this dependence we call it *scale free*. A stocastic model that has this property is a long-range memory process, and this is the model I will use as my null-hypothesis for the internal variations that we consider as noise. A *long-range memory* process is defined as a process where the autocovariance function sums to infinity.

There exists several models that can be used to simulate time series that exhibits different degrees of memory, or persistence. Examples of models that can be used to model long-range persistence are fractional Gaussian noises and FARIMA models. In the literature temperature has also been studied assuming simpler stochastic models exhibiting short-range persistence, such as autoregressive (AR) models or white noise.

2.2 Stationarity

In a *strictly stationary* time series the values $x_{t_1}, ..., x_{t_k}$ has the same joint cumulative distribution function as a time shifted set of values

$$P\{x_{t_1} \leq c_1, ..., x_{t_k} \leq c_k\} = P\{x_{t_1+h} \leq c_1, ..., x_{t_k+h} \leq c_k\}$$

for all $k = 1, 2, ...,$ all time points $t_1, t_2, ...,$ all numbers $c_1, c_2, ...$ and all time shifts $h = 0, \pm 1, \pm 2, ...$

A finite variance time series is called *weakly stationary* if:

(i) The mean value μ_t is constant and does not depend on time.

(ii) The autocovariance function $\gamma(s, t) = \mathbb{E}[(x_s - \mu_s)(x_t - \mu_t)]$ does only depend on s and t through their difference $|s - t|$.

In most real time series it is impossible to determine if we have strict stationarity. So the most widely used term is the weak stationarity, which is also the one I will refer to when I use the term *stationary*.

2.3 Stochastic noise models

In order to get a feeling for the stochastic noise models I will start by defining a *Gaussian white noise*, and a *Brownian motion*. In a Gaussian white noise $\{w_t\}_{t=1}^{\infty}$ each random variable w_t has a normal distribution with mean 0 and variance σ_w^2, and there is no dependence between the variables. An example of such a time series where $\sigma_w^2 = 1$ is shown in figure 2.1 (a). In a Brownian motion the value at each time step is given by the value at the previous time step plus a random number picked from a normal distribution with zero mean. In other words, a Brownian motion is the cumulative sum of a Gaussian white noise, $W(t) = \sum_{i=1}^{t} w_i$. An example of a Brownian motion is shown in figure 2.1 (b).

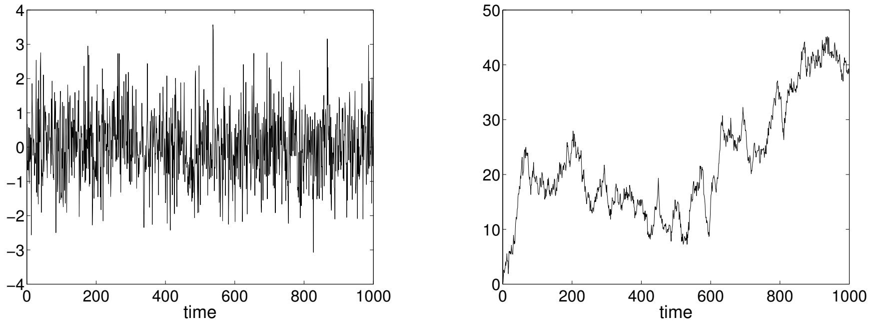

(a) White noise (b) Brownian motion

Figure 2.1

In the two time series in figure 2.1 we have an example of one time series with no dependence and one with strong dependence. Many real time series look like they are something in between these two cases, and Harold Edwin Hurst described examples of such time series empirically by rescaled-range analysis. He is famous especially for his studies of the hydrology in the Nile River, but he also studied other time series such as tree rings thicknesses, sunspot numbers, varve thickness, atmospheric temperature and pressure, etc. He found that the strength of persistence in the time series he studied could be described by only one parameter, which has been named the *Hurst exponent* after him. Later, [*Mandelbrot and Van Ness* (1968)] and several others have developed a class of time series called fractional Brownian motions and fractional Gaussian noises, which can also be described by the Hurst exponent. These models are extentions of the Brownian motions and white Gaussian noises, and are the models I will use when studying temperatures. We can easily see that a Brownian motion is not stationary because its values are drifting away, and so the variance will increase with time. The nonstationarity of the variance is what distinguishes the motions from the noises, although the differences between the most persistent noises and the least persistent motions are marginal. So in practice it is sometimes hard to determine if we have a noise or a motion.

2.3.1 Self-similar processes

If a nontrivial process $Y(t)$ is *H-selfsimilar*, then for all $a > 0$ there exists an $H \geq 0$ such that $Y(at) \overset{d}{=} a^H Y(t)$, where $\overset{d}{=}$ means that they are equally distributed. In practice this will be a useful property of a time series, because it means that we can measure the same statistical properties of a time series no matter what time scale we study, as long as we are in the range of time scales where the assumption of selfsimilarity is valid. If we, for instance, assume that this is valid for a temperature time series on time scales ranging from less than a month to decades, we can expect to see the same statistical properties when we study monthly resolved and yearly resolved time series.

The q-th order *structure function* is defined as:

$$S_q(t) = \mathbb{E}[|Y(t)|^q] \tag{2.2}$$

For a self-similiar process the structure functions are power laws:

$$S_q(t) = \mathbb{E}[|Y(1)t^H|^q] = \mathbb{E}[|Y(1)|^q]t^{Hq} \tag{2.3}$$

If $Y(t)$ is the cumulative sum of a stationary process $X(t)$, that is $Y(t) = \sum_{i=1}^{t} X(i)$, then we have that

$$\begin{aligned}\mathbb{E}[(Y(t) - Y(s))^2] &= \mathbb{E}\big[(\sum_{i=1}^{t} X(i) - \sum_{i=1}^{s} X(i))^2\big] \\ &= \mathbb{E}\big[(\sum_{i=s+1}^{t} X(i))^2\big] = \mathbb{E}\big[(\sum_{i'=1}^{t-s} X(i'+s))^2\big]\end{aligned}$$

Since X is stationary the expected value of the product of X at two time points depends only on the time difference. Hence this expression is the same as:

$$\mathbb{E}\big[(\sum_{i'=1}^{t-s} X(i'))^2\big] = \mathbb{E}\big[Y(t-s)^2\big]$$

By using this result and by expanding $(Y(t) - Y(s))^2$ we obtain that:

$$\begin{aligned}\mathbb{E}[Y(t)Y(s)] &= \frac{1}{2}\mathbb{E}\big[Y(t)^2 + Y(s)^2 - (Y(t) - Y(s))^2\big] \\ &= \frac{1}{2}\mathbb{E}\big[Y(t)^2 + Y(s)^2 - Y(t-s)^2\big]\end{aligned}$$

2.3.2 Fractional Gaussian noises and fractional Brownian motions

If we assume that $Y(t)$ is a self-similar process and use equation (2.3) with $q = 2$ we get that the covariance is:

$$\gamma(t,s) = \mathbb{E}[Y(t)Y(s)] = \frac{1}{2}\mathbb{E}[Y(1)^2]\big(t^{2H} + s^{2H} - |t-s|^{2H}\big) \tag{2.4}$$

This covariance function, together with the criteria that the process starts at 0 and has expectation value 0, is what characterizes a fractional Brownian motion (fBm). The increments $X(t) = Y(t+1) - Y(t)$ of the process $Y(t)$ will then be the fractional Gaussian noise (fGn), and the parameter $H \in (0,1)$ is what we call the Hurst exponent for this process. It should be noted that other definitions of fGns and fBms exist that are similar to this definition, see for instance [*Qian* (2003)]. By using eq. (2.4) we can easily derive the autocovariance function for $X(t)$:

$$\begin{aligned}\gamma(t,s) &= \mathbb{E}\left[(Y(t+1) - Y(t))(Y(s+1) - Y(s))\right]\\ &= \mathbb{E}[Y(t+1)Y(s+1)] - \mathbb{E}[Y(t+1)Y(s)] - \mathbb{E}[Y(t)Y(s+1)] + \mathbb{E}[Y(t)Y(s)]\\ &= \frac{1}{2}\mathbb{E}[Y(1)^2]\big(|t-s+1|^{2H} - 2|t-s|^{2H} + |t-1-s|^{2H}\big)\end{aligned}$$

Or expressed by the time lag $n = t - s$, since it is stationary:

$$R_{XX}(n) = \frac{1}{2}\mathbb{E}[Y(1)^2]\big(|n+1|^{2H} - 2|n|^{2H} + |n-1|^{2H}\big) \tag{2.5}$$

If we assume $n \geq 1$ we can remove the absolute values. (We know that $R_{XX}(0) = \mathbb{E}[Y(1)^2]$ and $R_{XX}(-n) = R_{XX}(n)$). Then we recognize the terms inside the parentheses as the centered difference approximation to the second derivative of $f(n) = n^{2H}$;

$$\frac{d^2n^{2H}}{dn^2} = (n+1)^{2H} - 2n^{2H} + (n-1)^{2H} - \frac{1}{12}f^{(4)}(n) - ... \tag{2.6}$$

Hence we get the power law approximation to the autocovariance:

$$R_{XX}(n) \approx \frac{\mathbb{E}[Y(1)^2]}{2}\frac{d^2n^{2H}}{dn^2} = \mathbb{E}[Y(1)^2]H(2H-1)n^{2H-2}$$

But the exact form of the autocovariance is:

$$R_{XX}(n) = \frac{\mathbb{E}[Y(1)^2]}{2}\left(\frac{d^2(n^{2H})}{dn^2} + \frac{1}{12}\frac{d^4 f}{dn^4} + \text{higher even order derivatives}\right)$$

where the k'th order derivative of $f(n) = n^{2H}$ with respect to the time lag n is:

$$\frac{d^k(n^{2H})}{dn^k} = 2H(2H-1)(2H-2)\cdot ... \cdot(2H-k+1)n^{2H-k} \tag{2.7}$$

For $H \in (0,1)$ and $k > 2$ we have that $2H - k < 0$, such that this expression is largest when n is as small as possible. Hence the power-law approximation to the covariance is not good for the smallest time scales. In figure 2.4 (b) we will also see that the deviations from a power law in a logarithmic plot is highest for the smallest Hurst exponents.

Now I will take a look at what happens if we average the time series $X(t)$, $t = 0, 1, ...$ over time segments of length m, such that we get the new time series: $X^{(m)}(k)$, $k = 0, 1, ...$ given by

$$\begin{aligned} X^{(m)}(k) &= \frac{1}{m} \sum_{t=km}^{km+m-1} X(t) = \frac{1}{m} \sum_{t=km}^{(k+1)m-1} (Y(t+1) - Y(t)) \\ &= \frac{1}{m} \left[Y((k+1)m) - Y(km) \right] \end{aligned}$$

The autocovariance of this time series will be:

$$\begin{aligned} \gamma(k,l) &= \mathbb{E}\left[X^{(m)}(k) X^{(m)}(l)\right] \\ &= \frac{1}{2m^2} \mathbb{E}\left[Y(1)^2\right] \left(|(k-l+1)m|^{2H} - 2\,|(k-l)m|^{2H} + |(k-l-1)m|^{2H} \right) \\ &= \frac{m^{2H-2}}{2} \mathbb{E}\left[Y(1)^2\right] \left(|\tau+1|^{2H} - 2\,|\tau|^{2H} + |\tau-1|^{2H} \right) \end{aligned}$$

where $\tau = k - l$ is the time lag. So we will see the same power law behavior for the averaged time series as for the original.

2.3.3 FARIMA models

Another class of long-range memory processes are FARIMA (also called ARFIMA). They have many of the same properties as fGns, but they are not self-similar. To understand these models we start by defining an ARMA(p,q) process:

$$x_t = \phi_1 x_{t-1} + ... + \phi_q x_{t-p} + w_t + \theta_1 w_{t-1} + ... + \theta_q w_{t-q}$$

This process is stationary and is a combination of a moving average and an autoregressive process with coefficients $\phi_1, ...\phi_q, \theta_1, ...\theta_p$. An extenstion of this class are ARIMA(p,d,q) models, which are integrated ARMA models. They are defined such that if we difference the process d times, $\nabla^d x_t = (1 - B)^d x_t$, then we get an ARMA(p,q) process. For ARIMA-models d is an integer, but if we extended this definition even more and say that d is allowed to be non-integer values as well, then we get a FARIMA-process. They are

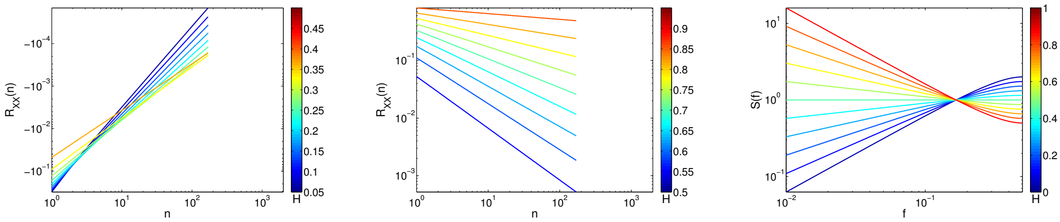

(a) ACF (H<1/2) (b) ACF (H>1/2) (c) Exact PSD

Figure 2.2: This figure shows a plot of the analytical expressions for the ACF and power spectrum of FARIMA processes. For the ACF the time scales only go to 171 because Matlab could not handle the very large values that we get for the gamma functions at longer time scales.

fractionally integrated ARMA(p,q) processes, and more specifically, these processes are defined for $-0.5 < d < 0.5$.

The ACF for FARIMA can be shown to be:

$$R_{XX}(n) = \frac{\Gamma(n+d)\Gamma(1-d)}{\Gamma(n-d+1)\Gamma(d)} \tag{2.8}$$

where Γ is the gamma function and n is the time lag. For $d > 0$ the sum over this function for $-\infty < n < \infty$ is infinite, and hence this is a long-range memory process. It also has the property that it is asymptotically a power law: $R_{XX}(n) \sim n^{2d-1}$ for large values of n. If we compare this to the power law that we have for fGns, we see that we can relate d to the Hurst exponent H by the relation: $d = H - \frac{1}{2}$. It can also be shown that the power spectrum for FARIMA is:

$$S(f) = [4\sin^2(\pi f)]^{-d}\sigma_w^2 \tag{2.9}$$

where σ_w^2 is the variance of the white noise process $\{w_t\}$. Figure 2.2 shows that the ACF and PSD are quite good approximations to power laws asymptotically. As $f \to 0$ we can see from equation (2.9) that $S(f) \sim f^{-2d}$, and hence we get a slope of $-\beta = -2d$ in the logarithmic plot of the PSD for the smallest frequencies.

2.4 The power spectral density

The power spectral density (PSD) for a time series X(t) is defined by

$$S(f) = \lim_{T\to\infty} \frac{E\big[|X_T(f)|^2\big]}{2T} = \lim_{T\to\infty} E\Big[\frac{1}{2T}\Big|\int_{-T}^{T} X(t)e^{-i2\pi ft}dt\Big|^2\Big] \tag{2.10}$$

where 2T is the length of the time series. This tells us what the expected energy per time unit is, as a function of frequency. The PSD can be related to the autocovariance by the Wiener—Khinchin theorem, which says that the power spectral density is the fourier transform of the autocovariance;

$$S(f) = \sum_{n=-\infty}^{\infty} R_{XX}(n)\, e^{-i2\pi fn} \tag{2.11}$$

provided that the autocovariance is weakly stationary and absolutely integrable. But from the definition of a long-range memory process we know that

$$\sum_{n=-\infty}^{\infty} R_{xx}(n) \to \infty$$

for $H > 1/2$, which means that it will not be absolutely integrable in this case.
So if we have an estimate for the autocovariance and the necessary conditions are fulfilled we can also make an estimate for the power spectral density, called the Blackman-Tukey estimator:

$$\widehat{S}^{(BT)}(f) = \delta \sum_{n=-M}^{M} \widehat{R}_{XX}(n)e^{-i2\pi fn\delta} \tag{2.12}$$

where δ is the length of the time step. For a time series with an ACF following a power law asymptotically, the power spectral density also has a power law dependence on frequency in the asymptotic limit:

$$S(f) \sim f^{-\beta} \tag{2.13}$$

The value β is a measure of persistence in the time series, and is related to the Hurst exponent by

$$\beta = 2H - 1 \tag{2.14}$$

See for instance the book of [*Beran* (1994)].

2.4.1 Extension of H

For $H \in (0,1)$, the time series $X(t)$ will be stationary and β will be in the interval $(-1,1)$. If we try to extend β to higher values than 1, then we actually study self-similar time series with self-similarity exponent $H \in (0,1)$, which is equal to the Hurst exponent for the differentiated time series. [*Malamud and Turcotte* (1999)] shows that if we take the derivative of a time series with PSD $S(f)$, then the corresponding PSD for the new time series will be

$$S'(f) = (4\pi^2 f^2)S(f)$$

If S(f) has a power law $S(f) \sim f^{-\beta}$ it means that the new PSD will have the power law $S'(f) \sim f^{-(\beta-2)}$.

Hence time series with β and H larger than 1 could be defined as the cumulative sum of a time series with $\beta - 2$ and $H - 1$. Likewise, time series with $\beta < -1$, $H < 0$ can be defined to be the difference of time series with $\beta + 2$, $H + 1$. By using the extension of β and H just described we can assign a spectral index and Hurst exponent to all the time series which can be obtained by differentiation or integration of a self-similar time series. For some of the analysis methods it is a requirement that the Hurst exponent is in a certain interval though, and then we need to difference or take the cumulative sum of the time series. For other methods we can do the analysis directly on the original time series.

2.5 Persistence in fGns and fBms

As mentioned earlier, the strength of persistence is determined by the Hurst exponent, and figure 2.3 shows some examples of what the time series with some given Hurst exponents will look like. When $H < 1/2$ the time series is antipersistent, which means that the correlations are negative and the value at one time point is likely to have the opposite sign of the preceding value. In the frequency domain this means that the contribution from the highest frequencies will be larger than the contribution from the lower frequencies, since $\beta < 0$. For $H > 1/2$ the time series is persistent and has positive correlations. Hence the value at one time point is likely to be close to the preceding value, and the higher frequencies will contribute less than the lower frequencies in the power spectrum, since $\beta > 0$.

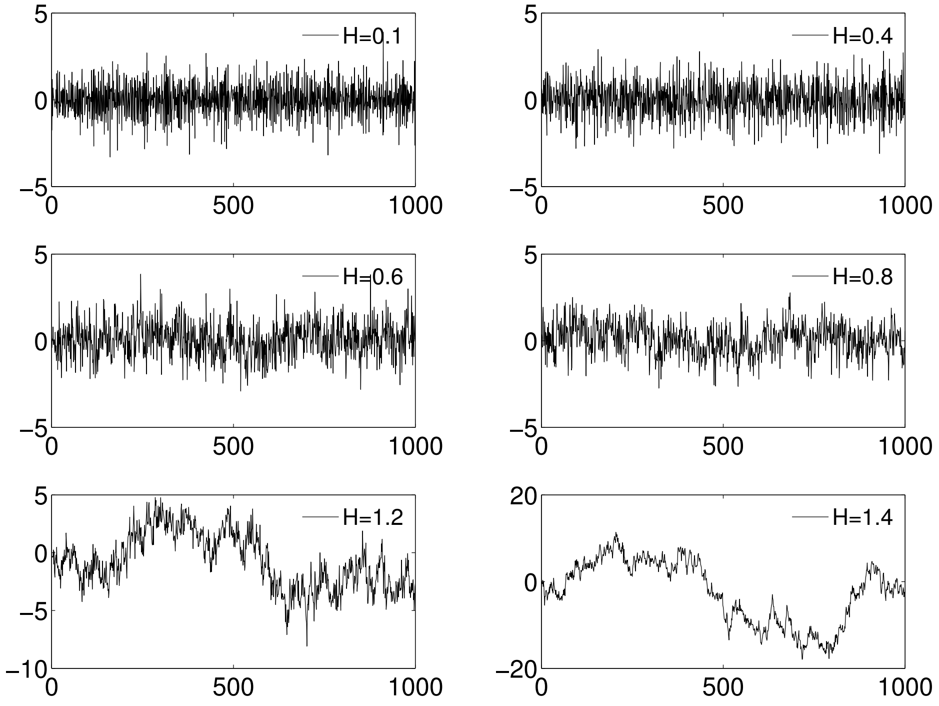

Figure 2.3: Some examples of time series.

2.5.1 Antipersistent fGns

In figure 2.4 (a) and (b) the exact autocovariance found in equation (2.5) is plotted for different values of the Hurst exponent, and we see that it is a good power law only asymptotically. By using the Blackman-Tukey estimator given in equation (2.12), the power spectrum is estimated and plotted in figure 2.4 (c). Here we see that also in the power spectrum the power law is a bad approximation for the shortest time scales/the highest frequencies. But as we go to longer time scales it will approach a power law asymptotically. In figure 2.4 (d) and (e) we get an idea of how large time scales we have to use if we want to see the power law when we have a discrete process.

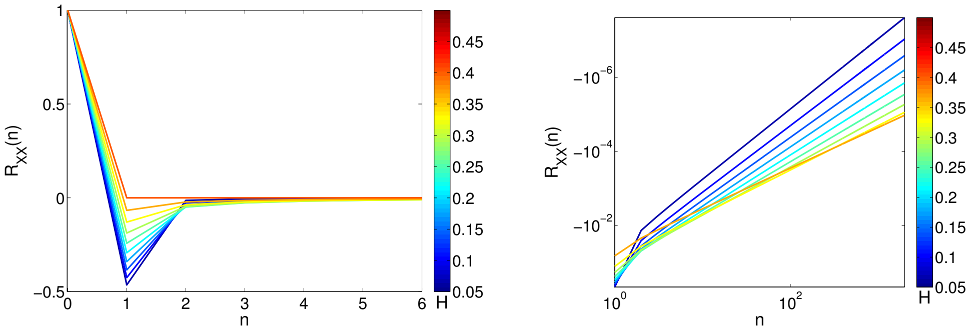

(a) Exact autocovariance

(b) Logarithmic plot of exact autocovariance for $n \geq 1$

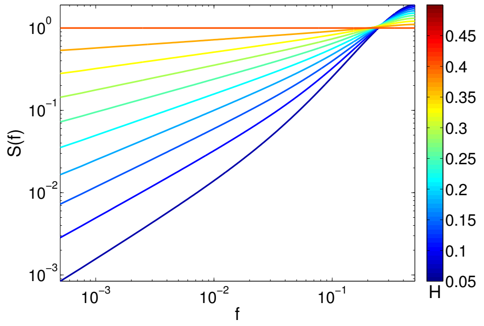

(c) PSD estimated from exact autocovariance

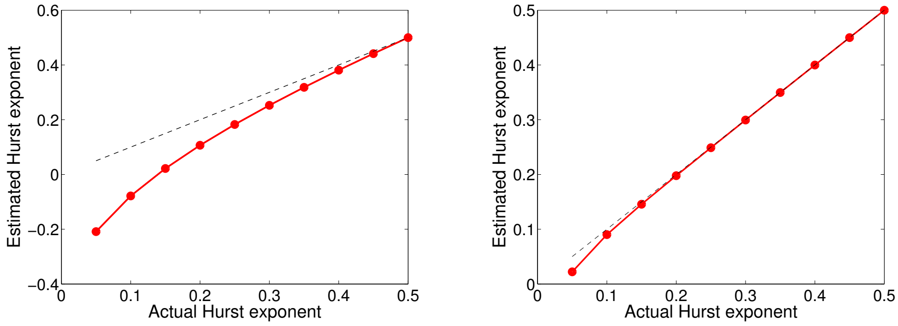

(d) Hurst exponents estimated in the frequency range (1/128 , 1/3)

(e) Hurst exponents estimated in the frequency range (1/1000,1/100)

Figure 2.4: In this figure we see that we need to use only the lowest frequencies in the PSD to get a good estimate of the Hurst exponent for the lowest Hurst exponents. Unfortunately these frequencies are also the ones that are most influenced by trends in the signal, as will be shown later.

2.5.2 Persistent fGns

In figure 2.5 we see that the power law for the autocovariance is valid for smaller time scales when we have a persistent fGn than what it was when we had an antipersistent fGn. So now we can expect it to be much easier to get good estimates for the Hurst exponent. As the Blackman-Tukey estimator fails to exist for persistent fGns I have not estimated the PSD and Hurst exponents in this case. But we will see later from other estimators for the PSD used on synthetic time series that we actually can get good estimates for the Hurst exponent for time series with Hurst exponents between 0.5 and 1.

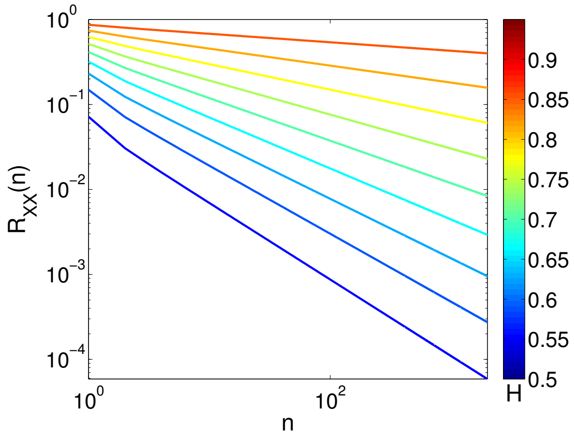

Figure 2.5: This figure shows a logarithmic plot of the exact autocovariance for persistent stationary time series.

Chapter 3

Methods used to estimate the Hurst exponent

There exists several different ways to estimate the Hurst exponent, and they can be divided into two groups: Those who works in the time domain (e.g. semivariogram, wavelet variance, MLE and DFA), and those who works in the frequency domain (e.g. periodogram). Common for the methods is that they in some way give us information about how much a time series varies at different time scales, and that the amount of variance has a power law dependence of the time scale.
For a general power law $y = ax^b$ where b is what we want to estimate, we can just take the logarithm of both sides such that

$$\log(y) = b \cdot \log(x) + \log(a) \tag{3.1}$$

So in a logarithmic plot b can be estimated by linear regression. Usually x is a vector with equal distance between the values. In a logarithmic plot this means that there will be a large distance between the smallest values of x, and a very small distance between the larger values. An example of such a plot is the periodogram (that we soon will discuss in more detail) in figure 3.1. If we used the points in the plot directly when doing the linear regression the regions where the points are most dense will be weighted more than those in the less dense regions. To avoid this problem one can use logbinning before the linear regression. This method gives us points that are evenly spaced in a logarithmic plot by averaging over all points that are inside bins of equal length in the logarithmic plot. These values will be represented by red circles in my plots, and the red lines are

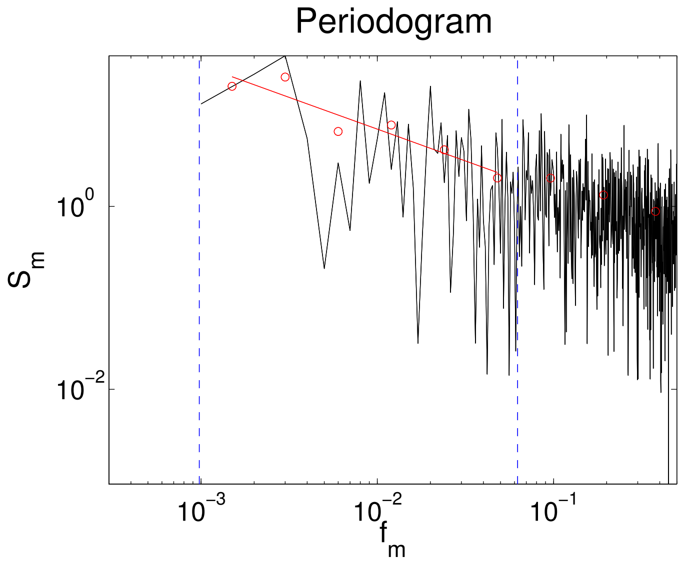

Figure 3.1: Example of a periodogram, where the process used has Hurst exponent 0.8, and the blue lines define the region that the straight line is fitted in.

linear fits to these points.

3.1 The periodogram

By discretizing the expression for the power spectral density (2.10) and ignoring the expectation operator and the fact that we don't have an infinite time series, we obtain the simplest estimator for the PSD; the periodogram:

$$\widehat{S}(f) = \frac{1}{N\delta t}\left|\delta t \sum_{n=0}^{N-1} y[n] e^{-\frac{i2\pi mn}{N}}\right|^2 \tag{3.2}$$

Here N is the length of the time series, δt is the time step, and the discrete frequency is $f = m\,\delta f = \frac{m}{N\delta t}$, where $m = 1, ..., N$. It is sufficient to use only the first half of the frequecies, since the second half would be just a mirror image of the first half. By defining the periodogram such that all of the power is in the first half of the frequencies, we can rewrite the periodogram to the form used by [*Malamud and Turcotte* (1999)]:

$$S_m = \frac{2\,|Y_m|^2}{N\delta}, \quad m = 1, 2, ..., N/2 \tag{3.3}$$

The power law approximation to the periodogram (2.13) is assumed to be valid for all values of β, but we need to be a little careful in the choise of time scales in the estimation.

Time scales

The frequencies we study are in the range

$$f \in \left[\frac{1}{N\delta t}, \frac{1}{2\delta t}\right],$$

So for a time series with monthly resolution, we have $\delta t = 1/12$ year and the time scales we study will be

$$T = \frac{1}{f} \in [2\delta t, N\delta t] = \left[\frac{1}{6}, \frac{N}{12}\right] \text{ years.}$$

where the shortest time scale is twice the time step according to Shannons sampling theorem, and the longest time scale is the length of the time series.

3.1.1 Windowing

The periodogram has quite poor statistical properties, so it is common to use different windowing techniques to improve this estimator. There exists lots of literature describing these techniques, for instance [*Jokinen et al.* (2000)]. In figure 3.2 I have studied two examples of windowing techniques, one with a Hamming window applied to the periodogram and one where I have used weighted overlapped segment averaging. As the Hurst exponent estimated from these plots and the periodogram in figure 3.1 does not differ much, I will assume that a periodogram without windowing will be sufficient for this purpose.

Another problem with this estimator is that if the time series studied also contains trends, the trends will affect the scaling properties on the longest time scales, such that we may get wrong estimates for the Hurst exponents.

3.2 Semivariogram

For a discrete time series the semivariogram is given by

$$\gamma_k = \frac{1}{2(N-k)} \sum_{n=1}^{N-k} (y_{n+k} - y_k)^2 \qquad (3.4)$$

where N is the length of the time series y, and k is a time lag. This tells us something about the average size of the fluctuations in a time series at different time scales, and

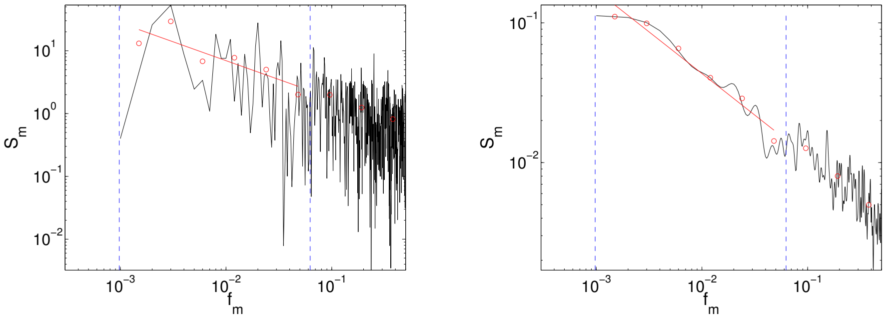

(a) Periodogram with Hamming window

(b) Periodogram with WOSA

Figure 3.2: Example of periodograms with windowing techniches, where the process used has Hurst exponent 0.8. In (a) a Hamming window is used in the Periodogram, and in (b) a Weighted Overlapped Segment Averaging (WOSA) technique is used, with 50 % overlapping, segment lengths of 128 and a Hamming window. The Hurst exponents estimated from the 3 different ways of using the periodogram are quite similar, but the peridogram with the Hamming window gives the best estimate in this case.

hence it is also a measure of the persistence in a time series. For a stationary time series [*Malamud and Turcotte* (1999)] showed that we have the relation

$$\gamma_k = V(1 - r_k) \tag{3.5}$$

where V is the variance, and

$$r_k = \frac{\frac{1}{N-k}\sum_{n=1}^{N-k}(y_{n+k} - \bar{y})(y_n - \bar{y})}{\frac{1}{N}\sum_{n=1}^{N}(y_n - \bar{y})^2} \tag{3.6}$$

is an estimator for the autocorrelation function. For a non-stationary process the semi-variogram is an estimator for the second order structure function. If the process is self-similar it should follow a power law

$$\gamma_k \sim k^{2H}$$

where H is the self-similarity exponent. The estimated exponent will then be an estimate for the Hurst exponent for the differenced process. Hence, for a stationary time series we construct the cumulative sum, and compute the self-similarity exponent to obtain the Hurst exponent H.

For a perfect fGn this method gives a very good estimate of the Hurst exponent, but it is very sensitive to the mean value used in the estimation though. When taking the cumulative sum of a process with a wrong estimate for the mean, the error we do will be accumulated. The most accurate result is obtained when we know the real expectation value for the mean, but in real time series we do not know this, and use the sample mean instead. This leads to a negative bias in the estimated H when H is close to 1 [*Rypdal et al.* (2013)]. An even larger problem with the semivariogram is its sensitivity to trends in time series.

3.3 Detrended fluctuation analysis

To deal with problems of trends influencing our estimates of the Hurst exponents, methods with detrending abilities have been constructed. A well established method for this purpose is the detrended fluctuation analysis. A description of this method can for instance be found in [*Kantelhardt et al.* (2001)] and [*Hu et al.* (2001)]. The former article describes the procedure in four steps, which I will summarize here:
(1) Calculate the profile (the cumulative sum):

$$Y(i) = \sum_{k=1}^{i} \big(x_k - \langle x \rangle\big)$$

where $\langle x \rangle$ is the sample mean of the time series $\{x_i\}_{i=1}^{N}$.
(2) Divide the profile into N_s segments of length s. If we start from the beginning of the time series, some values will probably remain at the end, so therefore we also repeat this starting from the end of the time series to be sure all points will be used. In total we then get $2N_s$ segments.
(3) Do an n'th order polynomial fit $p_\nu(i)$ to each segment ν, and subtract it from the original time series, such that for each segment s we get a time series

$$Y_s(i) = Y(i) - p_\nu(i)$$

For this n'th order fit the method will eliminate trends of polynomial order $n-1$ in the original time series.
(4) For each of the $2N_s$ segments, estimate the variance:

$$F_s^2(\nu) = \langle Y_s^2(i) \rangle = \frac{1}{s} \sum_{i=1}^{s} Y_s^2[(\nu-1)s + i]$$

Then the fluctuation function is given by:

$$F(s) = \left[\frac{1}{2N_s} \sum_{\nu=1}^{2N_s} F_s^2(\nu) \right]^{1/2}$$

The n'th order fluctuation function $F^{(n)}(s)$ will then have a power law dependence on the time scales:

$$F^{(n)}(s) \propto s^{\alpha}$$

where α will be an estimate of the Hurst exponent. This function will then be valid for scales $s \geq n+2$, but one should not use very large values of s because of poor statistics. The largest value should be significantly smaller than the length of the time series, and $s \leq N/10$ is a common requirement.

3.4 Wavelet analysis

Another method that helps us eliminate the effect of trends in the signals we study is wavelet analysis. The generalized form of a wavelet transform is, as described by [*Malamud and Turcotte* (1999)]:

$$W(t,a) = \frac{1}{a^{1/2}} \int_{-\infty}^{\infty} g\left(\frac{t'-t}{a}\right) f(t') dt' \tag{3.7}$$

This is a more general transform than the fourier transform, because in this you can not only get a picture of the frequency distribution in the signal, but also how the frequencies varies with time. In this expression a is a scale parameter. By choosing a small value for a one can study how much the signal varies on short time scales (high frequencies), and with larger values one can study how much the signal varies on longer time scales (low frequencies). It is common to study a set of scales where a is increased by a factor two each time.

The function $g(t')$ is known as the mother wavelet. It can have several different forms. One that is commonly used is the Mexican hat:

$$g(t') = \left(\frac{1}{2\pi}\right)^{1/2} (1 - t'^2) e^{-t'^2/2} \tag{3.8}$$

This is the second derivative of a gaussian, but its derivative of other orders are also used. The n'th order derivative can detrend a polynomial of order up to $n-1$.

If we calculate the variance of W and plot it as a function of a for a time series with scale free spectrum and spectral index β, we see that it is well approximated by the relation

$$V_W \sim a^{\beta} \tag{3.9}$$

By taking the logarithm of both sides we get that:

$$\log(V_W) = \beta \cdot \log(a) + \text{constant}$$

and hence we will see a straight line in a loglog plot, with slope β. From this we get an estimate of the Hurst exponent by the relation

$$H = \frac{\beta + 1}{2}$$

The power-law approximation is assumed to be valid for all values of β, but just as in the periodogram we need to be a little careful in the choice of time scales. From [*Torrence and Compo* (1998)] we know that the scale parameter a for an m'th derivative of a gaussian distribution is related to the fourier period T by

$$T = \frac{2\pi a}{\sqrt{m + \frac{1}{2}}} \tag{3.10}$$

Since the mexican hat is the 2nd derivative of a gaussian we get that

$$T \approx 4a$$

The smallest resolvable timescale in a time series is twice the sampling period ($2\delta t$), so that the smallest scale parameter one should use will be

$$\min(a) \approx \frac{\min(T)}{4} = \frac{\delta t}{2}$$

The largest meaningful timescale we can have is the length of the time series, so the largest scale parameter we should use is

$$\max(a) \approx \frac{\max(T)}{4}$$

When calculating the wavelet coefficients we move the mother wavelet along the time series. For the smallest and largest times, part of the mother wavelet is outside the endpoints and does not give us meaningful coefficients. So when estimating the wavelet variance these coefficients should not be included. If we do include them though, the largest errors are found on the largest time scales. In my analyses of temperature data I have only studied shorter time scales, such that the effect of this error will not be that large, and hence I have not removed these coefficients in the estimation.

3.5 Other methods that have not been used

- Maximum Likelihood Estimation (MLE). When we use this method we have to assume a model for the noise and trends, and then compute the most likely coefficients in this model. Since I don't know yet if all of my time series will have a good approximations to the stochastic model I have assumed, I will focus more on the methods that give me more information about how good the fit is. But for further work when I have more knowledge of the most likely models, this can be a good method to use.

- The autocovariance function (ACF), because the estimator for the ACF has very large errors for large time lags, which is inconvenient when analyzing long-range memory models where the long time scales are what we want to study. These errors can also give rise to some negative covariances that we cannot study in a logarithmic plot.

- Rescaled range analysis, because it is too similar to the semivariogram, and does not work well for time series with trends.

Chapter 4

Analysis of synthetic fGns

To explore the properties of the fGns and fBms further, I will generate synthetic data. The routine that I will use in the generation is developed and distributed as an R package by [*McLeod et al.* (2007)]. This routine is based on the exact expression for the autocovariance of a discrete fGn, and should therefore produce data that has a good power law approximation for the variations only asymptotically.

This chapter will give us some information about how well the methods described in the previous chapter works, given that discrete fGns/fBms are good models for the time series that we want to study. But as we will see later, the Hurst exponents estimated from the temperature data does not get a bias in the cases where they are believed to become nonstationary, as we see for the models.

4.1 Antipersistent fGns

Figure 4.1 shows what the average of 100 periodograms, wavelet variances, semivariograms and DFA2s look like for different Hurst exponents between 0 and 0.5. This will also give us an idea of what the analysis of fBms with Hurst exponents between 1 and 1.5 look like. The synthetic data the analysis is performed on are time series with 2000 data points. It should be noted that figure 4.1 (a) look very much like figure 2.4 (c), which could give us some extra trust in the periodogram as a sufficiently good estimator for the PSD for this purpose. In all these figures except semivariogram it is clear that the shortest scales must be avoided when fitting a straight line. But as we go to longer time scales the errorbars in the estimation of the Hurst exponent becomes larger. The method that gives us the best estimates in this case is DFA2 used on the longest time scales. But since the temperatures are not modelled accurately enough by these discrete synthetic data on the shortest time scales, we don't have to worry that much about the strong bias that we see for wavelet variance and periodogram.

The main reason for the errorbars in the estimation is the short length of the time series, which will be seen clearly in section 4.6. The size of the error does also differ between the different methods and with the time scales we study. When estimating the Hurst exponent by the different methods the range of time scales used for fitting straight lines in log-log plots should be comparable with each other. Then we don't get varying error bars due to varying number of points used in the regression.

4.2 Persistent fGns

Figure 4.2 shows the same as figure 4.1, but for persistent fGns. Also here we should not use the very smallest time scales, but we can fit a straight line to much smaller time scales than for antipersitent fGns. DFA2 seems to be the best method because it has the smallest error bars. The error bars for DFA2 are slightly smaller for the smallest time scales than for the largest, but for the smallest time scales they have a little positive bias. Perhaps is this tiny bias just a result of too few values used in the average, and not a real effect. If the estimation of the wavelet variance was modified to not include the endpoint effects, the errorbars of this would be smaller, probably close to what we see for DFA.

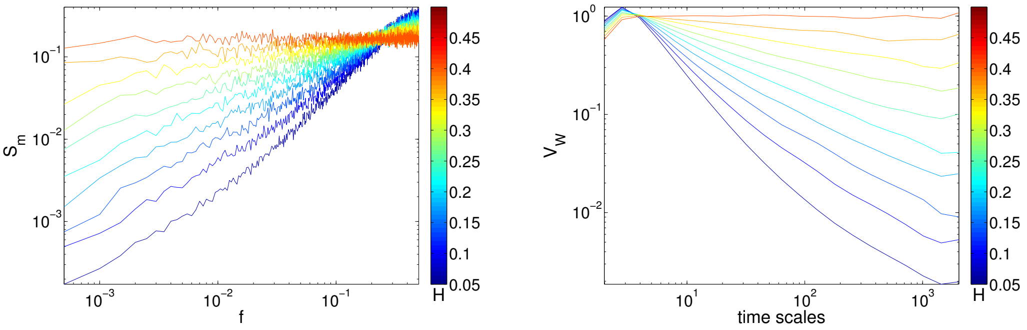

(a) Periodograms (b) Wavelet variance

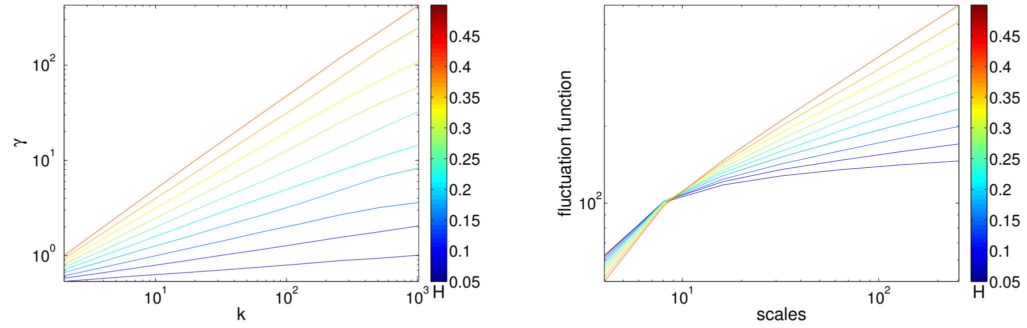

(c) Semivariogram (d) DFA2

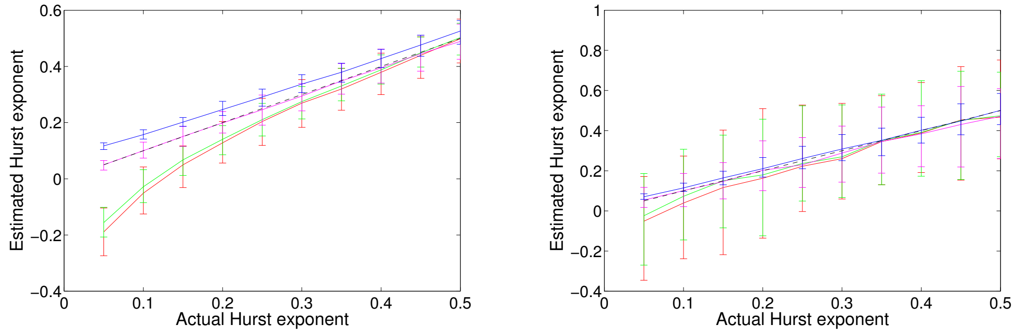

(e) $\widehat{H}$ on short time scales (f) $\widehat{H}$ on long time scales

Figure 4.1: In panel (a) to (d) the plots show the average of 100 antipersistent synthetic fgns with 2000 data points, where the colour denotes the Hurst exponent. In (e) and (f) the Hurst exponents estimated from the plots in (a) to (d) are given. In (e) the Hurst exponents are calculated on scales from 3 to 128 time units, except for DFA2 where it is calculated on scales from 8 to 128 time units to avoid the breaking point that we see in (d). In (f) the Hurst exponents are calculated on scales from 32 to 200 time units. The wavelet variance is the green curve, periodogram the red curve, semivariogram the pink curve and DFA2 the blue curve.

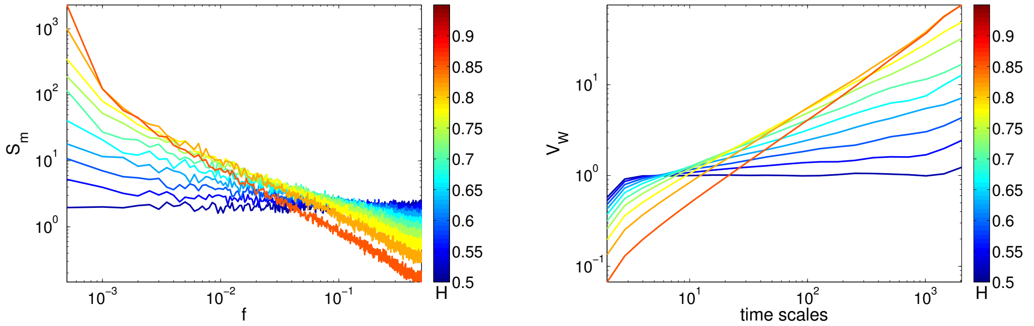

(a) Periodograms (b) Wavelet variance

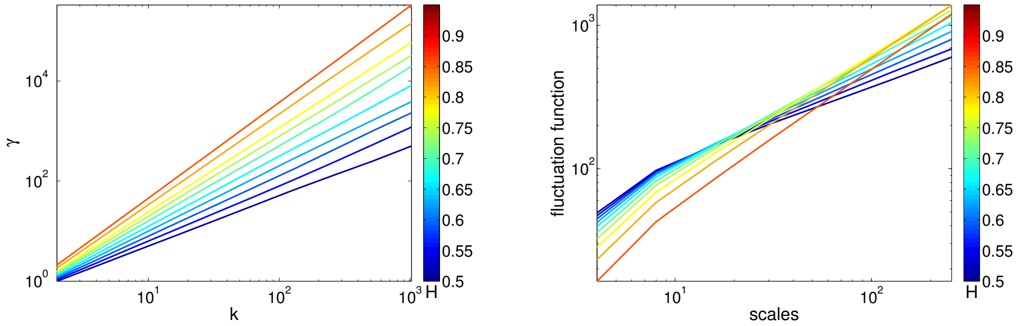

(c) Semivariogram (d) DFA2

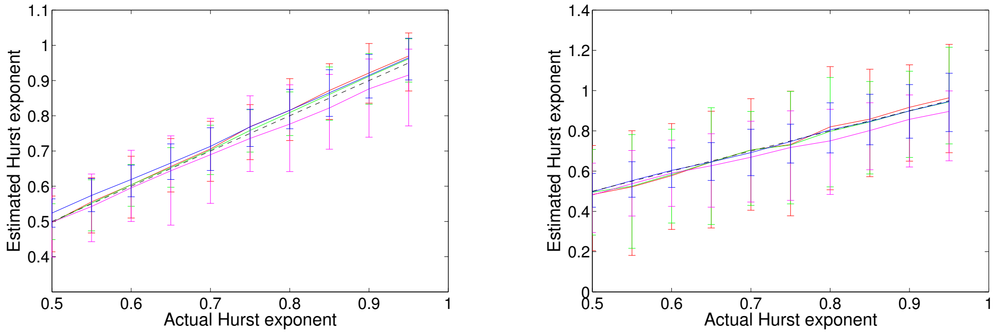

(e) $\widehat{H}$ on short time scales (f) $\widehat{H}$ on long time scales

Figure 4.2: This plot shows the same as figure 4.1, except that we have persistent fGns in this case. In (e) and (f) we have that the wavelet variance is the green curve, periodogram the red curve, semivariogram the pink curve and DFA2 the blue curve.

4.3 The effects of adding a trend to the data

In the following I will study what will happen if I try to estimate the Hurst exponent for an fGn if the time series actually is an fGn with a linear trend. Later we will see that the largest trend in global temperature records is found for land temperatures (see table 5.1). Since the largest trend is expected to give the largest errors, this is the trend I will add to the synthetic fGns/fBms that I will study. The result of adding this trend is shown in figure 4.3.

The most affected method is the semivariogram, which overestimates many of the Hurst exponents for the stationary time series. The reason for the steep slope for the Hurst exponents below 1 seen in figure 4.3 (c) is that it is the cumulative sum of those time series that are analysed. The DFA2 is as expected not affected by this trend, since it is designed to remove linear trends. The periodograms are affected by this trend on the longest time scales. The wavelet variance is supposed to remove this trend when we remove the endpoint effects, but if we don't it will be affected by the trend on the longest time scales. The shape of these does not change much if we instead add a higher order trend that is also estimated from the global land data.

For time series that we know are stationary with a trend, figure 4.3 shows that the smallest time scales are the best to use for the periodogram and the wavelet variance since the trend does not have influence there. The semivariogram gives biased and uncertain results, so it should be avoided in this case. It seems also that DFA2 has a tiny bias correction as we go to longer time scales, but then it also gets larger error bars. From this I conclude that for stationary discrete fGns with a linear trend I can continue using the time scales from 3 to 128 for periodogram and wavelet variance, and try time scales from 8 to 256 for DFA2, where 256 is chosen with a hope of a correction of the bias without getting much larger error bars. The bias of the estimated Hurst exponent on these scales and its 95% confidence intervals is shown in figure 4.4. The confidence intervals seen in that figure are interpolations of the error bars, where the step between each Hurst exponent analyzed is 0.05.

If we know that a time series is nonstationary figure 4.3 shows that the periodogram and wavelet variance has a large negative bias for the smallest time scales, and as we go to longer time scales the error bars become very large. By comparing this with the plot in figure 4.1 (f) we see that the trend is not what causes the main problems, so

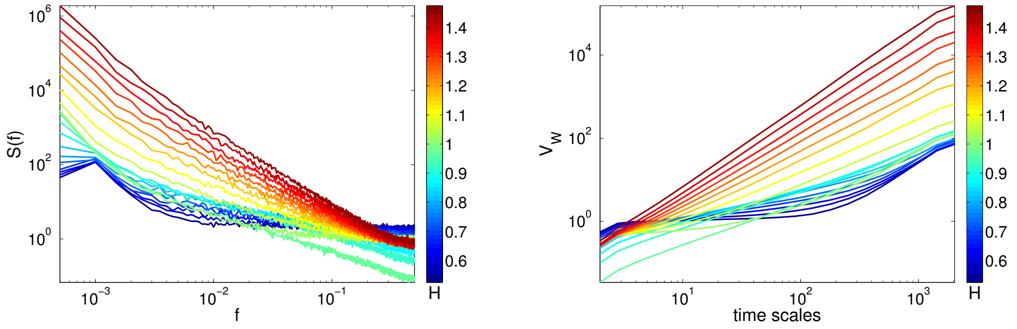

(a) Periodograms (b) Wavelet variance

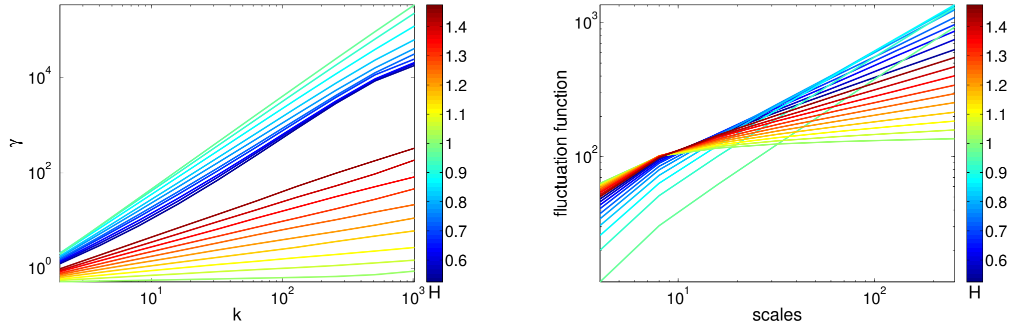

(c) Semivariogram (d) DFA2

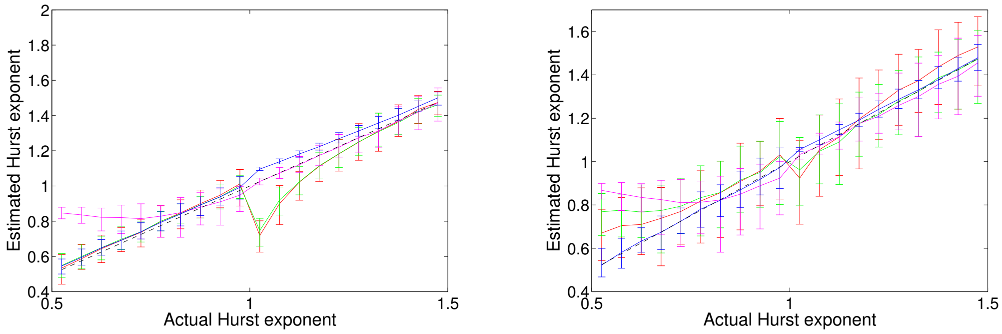

(e) $\widehat{H}$ on time scales 3 to 128 (f) $\widehat{H}$ on time scales 16 to 512

Figure 4.3: The effect that adding linear trends to the fgns/fbms will have on the shape of the periodograms, wavelet variances, semivariograms and DFA2s. In (e) and (f) we have that the wavelet variance is the green curve, periodogram the red curve, semivariogram the pink curve and DFA2 the blue curve.

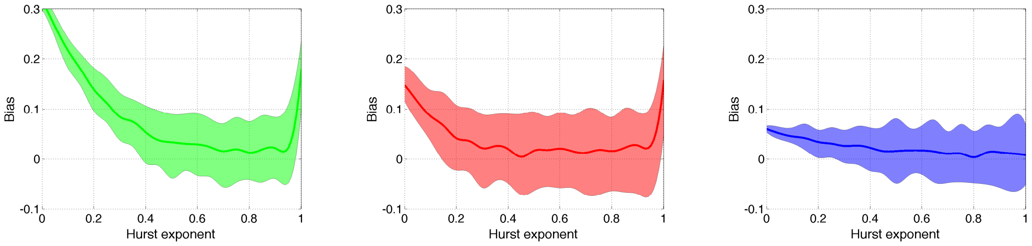

(a) Wavelet variance (b) Periodogram (c) DFA2

Figure 4.4: This figure shows the bias in the estimated Hurst exponent of an fgn with a linear trend. The region used in the estimation is 3 to 128 time units, except from in DFA2 where it starts at 8 and ends on 256.

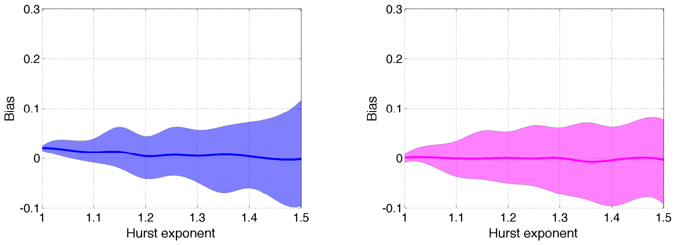

(a) DFA2 (b) Semivariogram

Figure 4.5: This figure shows the bias in the estimated Hurst exponent of an fBm with a linear trend. The region used in the estimation is 3 to 128 time units for semivariogram, and 32 to 256 time units for DFA2.

detrending would not help much on reducing the error bars.

DFA2 and semivariogram look quite trustworthy in this case. The semivariogram is almost unbiased, but has large error bars on the longest time scales. DFA2 has a small positive bias for the shortest time scales, but slightly larger error bars for the longer time scales. Hence I suggest that the DFA2 should be used in the range 32 to 256 time units even though that means a little shorter interval to fit the straight line to compared to the other methods, and semivariogram in the range 3 to 128 time units. The bias and its 95% confidence intervals is shown in figure 4.5.

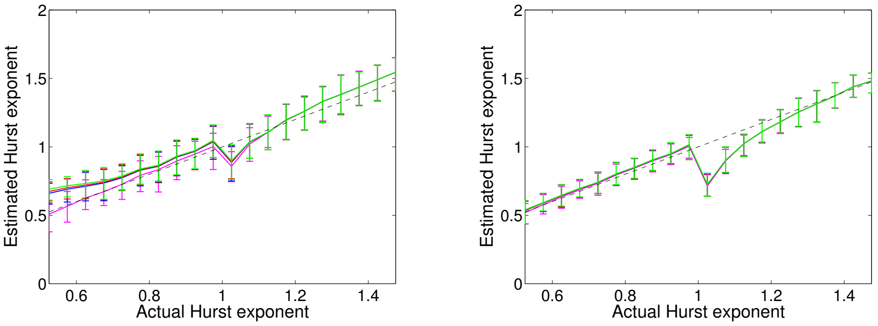

(a) Frequencies: (1/500, 1/10) (b) Frequencies: (1/128, 1/3)

Figure 4.6: The effect of linear (blue), cubic (red) and third order polynomial (green) trends when estimating the Hurst exponent with the periodogram. The pink curve is when we have no trend.

Higher order trends

Since we cannot be 100 % sure of the shape of the trend in the temperature, we should also check what the results are when we fit higher order polynomial trends to the global land temperature and add these trends to syntetic time series. The results of fitting a 1st, 2nd and 3rd order polynomial trend is shown in figure 4.6, and we see that the results are very similar on the time scales that I have chosen to use. This figure shows the results only for the periodogram, but the same is seen on the other methods.

A comment on the overestimation below $H = 1$

For the stationary time series we could see in figure 4.4 that the Hurst exponent was overestimated just below $H = 1$ for the periodogram and wavelet variance, and that should be explored a bit more. Since the method for generating synthetic time series does not work for $H = 0$ and $H = 1$, the Hurst exponents used in the generation of those are 0.001 and 0.999. The bias seen between $H = 0.95$ and $H = 0.999$ is just an interpolation, so to study exactly at what Hurst exponents we get this overestimation figure 4.7 is made. There we can see that the very large overestimation is so close to 1 that it is unlikely that it will be a problem in practice. It should also be noted that the shapes of the periodograms and wavelet variances had good power-law approximations.

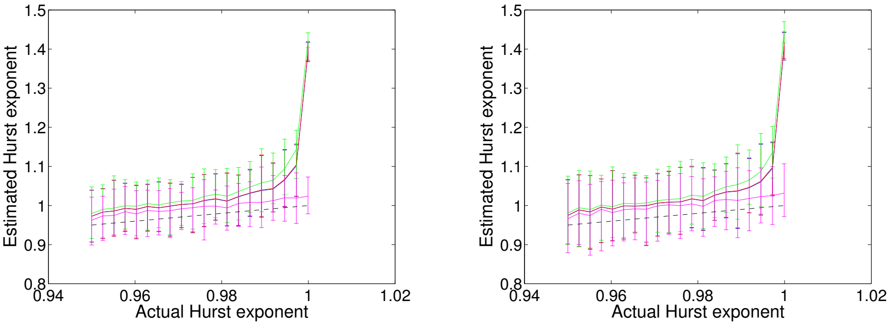

(a) Wavelet

(b) Periodogram

Figure 4.7: A closer look at the peak that we see in the periodogram and wavelet variance just below Hurst exponent 1. The extremely large bias is only in a very small interval, so this will not be a big problem in the study of real time series. The pink curve is a pure fGn, the blue curve has a linear trend, the red curve has a cubic trend, the green curve has a third order polynomial trend, and the error bars are 95 % confidence intervals.

4.4 Error estimates using Bayes theorem

In the figures above I have estimated the distribution of the estimated Hurst exponents given that we know the real Hurst exponents. But as we move on to real time series we need to turn the question around, because what we actually want to know is: What is the distribution of the actual Hurst exponent (H) given the estimated Hurst exponent ($\hat{H}$)? The answer of this question is given by Bayes theorem:

$$f_H(H|\hat{H}) = \frac{f_{\hat{H}}(\hat{H}|H) f_H(H)}{f_{\hat{H}}(\hat{H})}$$

Since the distribution of Hurst exponents, $f_H(H)$, is unknown I will just assume that it is uniform. $f_{\hat{H}}(\hat{H}|H)$ can be found from the estimates of biases done in figure 4.4 and 4.5 (if they can be assumed to be good enough), and $f_{\hat{H}}(\hat{H})$ can be found for the gridded data sets where I will calculate an ensemble of estimates for the Hurst exponent, given that all realizations of the Hurst exponents come from the same processes. Theoretically it would be nice to express the uncertainty this way, but in practice the estimates would be very uncertain since the time series are realizations of very many different processes.

4.5 Stationary or nonstationary?

In some of the methods described above, we should know in advance if the time series is stationary or nonstationary before calculating the Hurst exponent. For some methods we need to difference or take the cumulative sum of the time series before the analysis, and for some we need to study different time scales to get a good power law fit when we have a discrete fGn/fBm. In many cases it is not easy to determine if a time series is stationary or not, and if we make the wrong assumption the estimated Hurst exponent could be far away from the real Hurst exponent. So far I have not found any methods that can give a clear distinction between stationary and nonstationary time series. But if we don't care about the bias we get for periodogram and wavelet variance, then the periodogram and wavelet variance can be used to determine this. Figure 4.8 shows more clearly exactly what happens when the Hurst exponent goes from below one to above one. The physical interpretation of the periodogram would be that the process suddenly get much more energy in the fastest fluctuations if the process becomes nonstationary, which would be very strange.

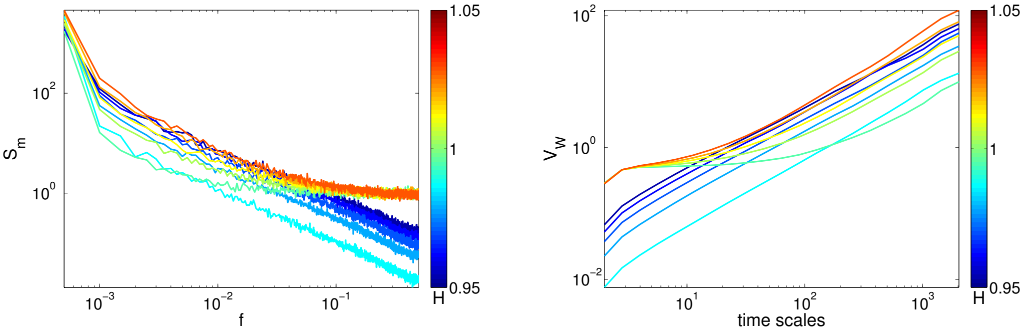

(a) Periodogram (b) Wavelet variance

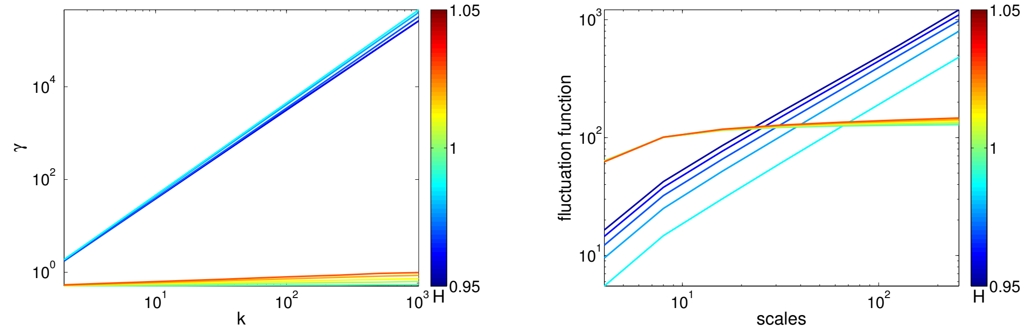

(c) Semivariogram (d) DFA2

Figure 4.8: This figure shows what happens with the shape of the plots when we go from Hurst exponents below one to Hurst exponents above 1. The slope of half of the semivariograms is very high because I have used the cumulative sum of the fGns.

4.6 Error estimates for shorter time series

The real time series that we analyze do not always have the same length, and the accuracy of the results could therefore be varying. To check the accuracy of the estimated Hurst exponent for shorter time series figure 4.9 has been made, which plots the length of the 95% confidence interval as a function of Hurst exponent and length of the time series.

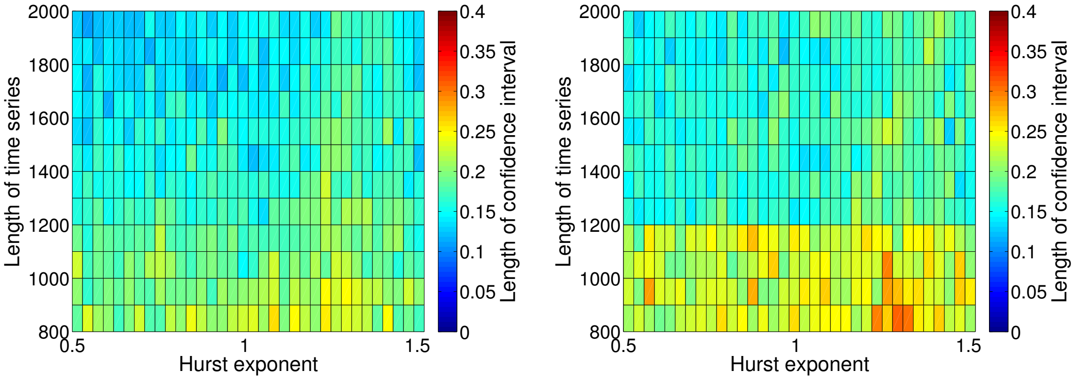

(a) Wavelet variance (b) Periodogram

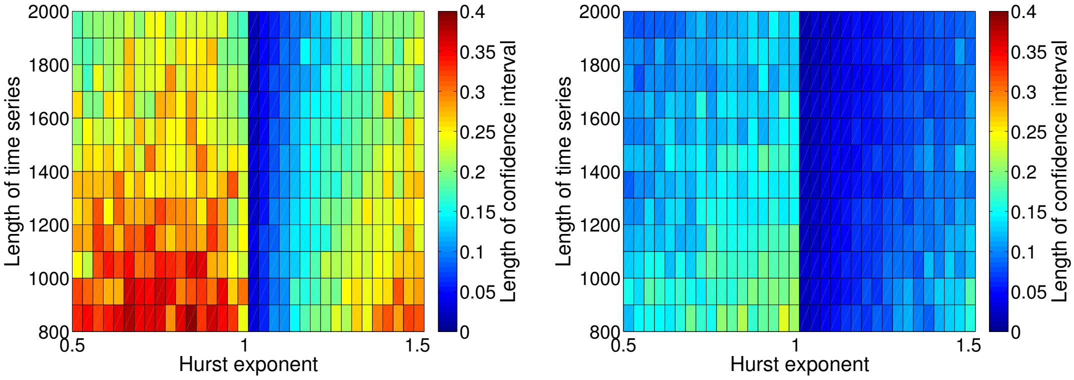

(c) Semivariogram (d) DFA2

Figure 4.9: Length of 95 % confidence intervals in the estimation of the Hurst exponents. The interval used in the estimation of the Hurst exponent is 3 to 128 time units, except for DFA2 where it starts at 8.

It is important to note that this figure does not care about the bias in the estimation, only the length of the confidence intervals. There are no trends added to the synthetic

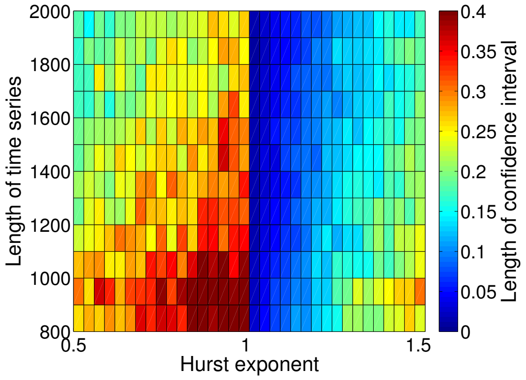

Figure 4.10: Length of 95 % confidence intervals in the estimation of the Hurst exponents using DFA2. The interval used in the estimation of the Hurst exponent is 8 to 256 time units for stationary time series and 32 to 256 for nonstationary time series.

fGns in these figures. But due to the analysis of fGns with trends above, one can expect that fGns with a linear trend added should give approximately the same results as this at the chosen time scales. For the periodogram there seems to be a clear distinction between time series shorter than and longer than 1200. This does probably have something to do with the way the logbinning is done. The DFA2 gives very small errors, but we remember from figure 4.1 that it has a bias on these time scales. As suggested earlier we should instead use longer time scales, and in figure 4.10 we can see that the error bars on these time scales are much larger.

4.7 Chapter summary

For stationary and persistent temperatures the discrete fGn could be a good model, but for nonstationary temperatures I would prefer a model that had a better power law approximation also for the smallest time scales. For that FARIMA may be better than fBms. If we assume that the temperature data do not show the deviations from power law on small time scales for the nonstationary time series, then wavelet variance, DFA2 and periodogram give us quite good results for all time series. The error is increasing though if the time series are too short, so I will choose a limit of 1200 data points when determining which time series that are good enough to use. As mentioned before, the semivariogram is a very uncertain estimator when using real data, because we don't know the real expectation value of those data, which gives rise to a bias, and because trends give rise to even greater biases.

In most of my results I will use only the wavelet variance and periodogram to estimate the Hurst exponents, mainly because it is easier to discover and interpret possible deviations from a power-law scaling with these methods than with DFA and semivariogram.

Chapter 5

Trends in temperature data

By studying a plot of the global temperature record, for example figure 5.1, one can see that the temperature has increased remarkably since the recordings started, and it is likely that humans have contributed to this increase. But exactly how much is anthropogenic and how much comes from natural variations is an ongoing discussion. My null hypothesis will be that the natural variations are fGns or fBms (at least) on the time scales that we are able to study with instrumental temperatures, but they need not necessarily have the same scaling on all time scales. In the recent book of [*Lovejoy and Schertzer* (2013)] the temperature variations are divided into three scaling regimes for local temperatures:
(1) Weather; on time scales up to around 7 days, which is the life time of the largest eddies in the atmosphere.
(2) Macroweather; on time scales longer than weather, up to around 10-100 years.
(3) Climate; on longer time scales than macroweather.
[*Pelletier and Turcotte* (1999)] also try to distinguish between different regimes when estimating slopes, where they in addition introduces another regime of white noise on the very longest time scales.

The time scales I want to study are from a few months up to around 10 years, which is within the macroweather regime. In some of the time series I will analyze, the climate regime starts at time scales just above 10 years, which is the reason why I don't go to longer time scales. This small time scale for the transition to the climate regime does probably have something to do with the trend in the temperature, and could be removed such that the macroweather regime could be extended to longer time scales for these time series. But removing only the trend caused by external forcings is not an easy task, since it is uncertain how large part of the temperature variations are a result of external forcings and how much is because of internal variations within the climate system. In analyses of detrended global temperature there is no clear distinction between the macroweather and the climate regime, which suggests that all time scales on global scales can be part of the climate regime.

The periodogram can give us an intuitive understanding of this problem and of the periodicities we have in the time series. In figure 5.2 we can see that the temperature with no detrending has more power in the lowest frequencies for that to be a part of the power law scaling seen on higher frequencies, and that some of this power can be removed by subtracting a trend from the temperature data. But we need to be careful about how much we subtract. The seventh order polynomial is a very good fit to the temperature as we see in 5.1, but if we subtract it from the signal we seem to loose too much power in the lowest frequencies. That suggests that this variation can actually be just a part of the noise, even though it looks like it could be driven by a forcing.

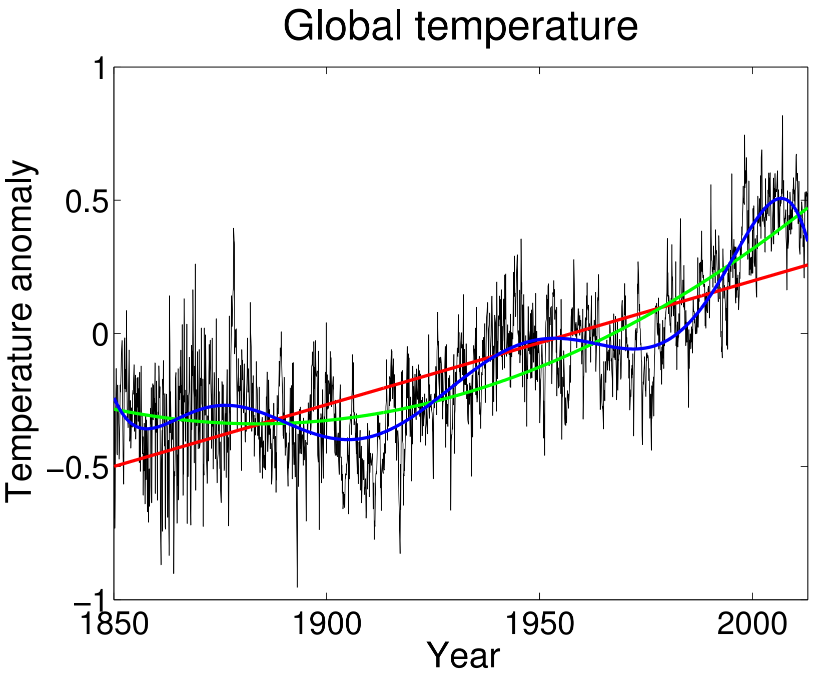

Figure 5.1: The black curve shows the monthly temperature anomaly. The red, green and blue curves are linear, cubic and seventh order polynomial fits to the temperature.

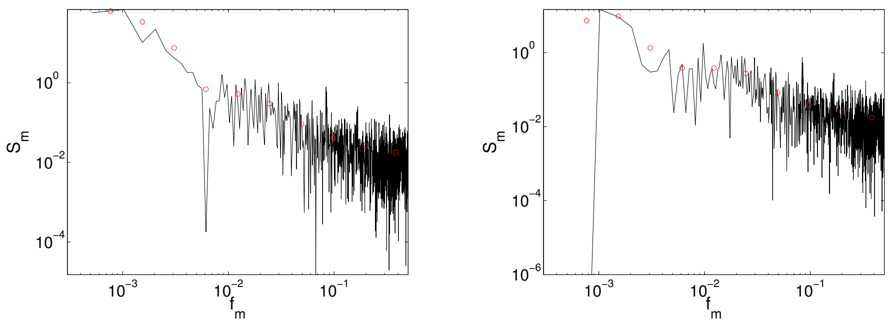

(a) No detrending (b) Linear detrending

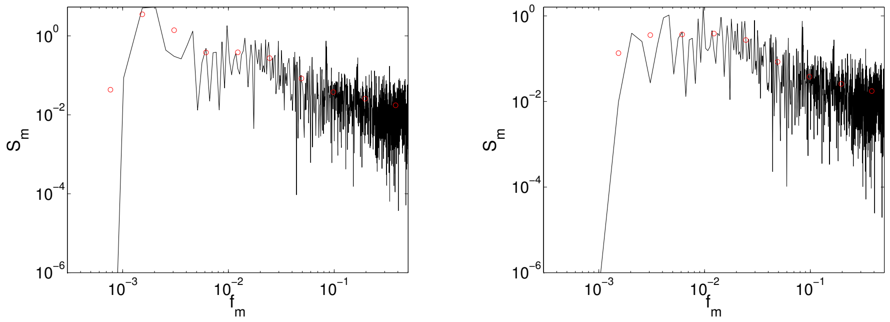

(c) Cubic detrending (d) Seventh order polynomial detrending

Figure 5.2: Periodograms of the global temperature with different orders of polynomial detrending.

5.1 AIC and BIC criteria

There exists several different ways of determining which polynomial degree that fits best to the data, and the polynomial does also depend on the chosen noise model. If we assume a white noise model a simple regression analysis will give us the trend. A linear regression on the global temperature data gives the result in table 5.1. This may be a good estimate, but it would be more accurate if we assume an fGn model for the noise and used maximum likelihood to estimate the most likely coefficients of the trend, after we have found in other ways what the most likely order of the trend is.

The AIC (Akaike's information criteria) is given by:

$$AIC = 2k - 2\log(L) + \frac{2k(k+1)}{n-k-1} \tag{5.1}$$

(see for instance [*Burnham et al.* (2010)]), where k is the number of parameters in the model, n is the length of the time series and L is the maximum value of the likelihood function. The value of k that minimizes this expression is assumed to be the number of parameters that gives us the best model. The AIC is constructed such that it is good to find a simple polynomial fit (low k) and a high value of the maximum likelihood. The last term is a correction term for the limited length of the time series. For large n this goes to zero. Another criterion that could be used to determine the number of parameters in a model is BIC (Bayesian information criteria), which differs only slightly from AIC.

	a_0	a_1
Global temperature	-0.4998	$0.3869 \cdot 10^{-3}$
Global land temperature	-0.6372	$0.5312 \cdot 10^{-3}$
Global sea temperature	-0.4522	$0.3348 \cdot 10^{-3}$

Table 5.1: A polynomial fit $a_0 + a_1 t$ is done to the temperature data, where t is time measured in months since 1850.

5.2 Normality of data

Another requirement our data needs to fulfill to fit the fGn model is that they are normal distributed. That can easily be checked by making a quantile-quantile plot and see if the points follow a straight line. Figure 5.3 shows the result of this, and we can observe that the fit is not perfect. If we do a linear detrending of the data we get the result in figure 5.4. The fit is improved compared to the previous figure. This gives us extra trust in the existence of a trend in the global temperature. For the global land temperature there is still a significant deviation, and the fit does not become better with higher order detrending. This may be a result of very cold or warm periods that we can find in land temperatures. Another explanation can be that the seasonal variations are not completely removed when subtracting the climatology.

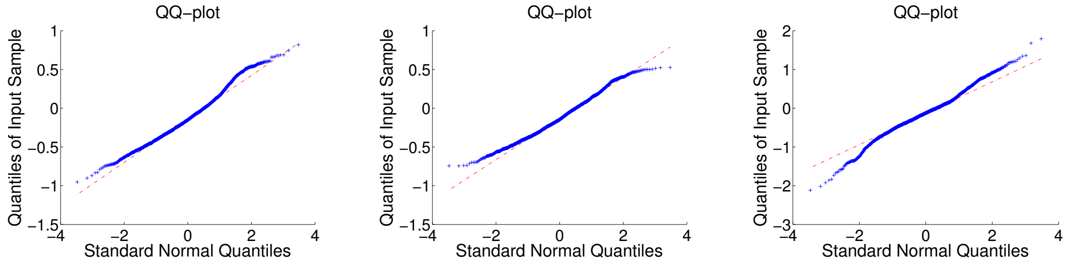

(a) (b) (c)

Figure 5.3: QQ-plot of global temperatures (a), global sea surface temperature (b) and global land temperatures (c).

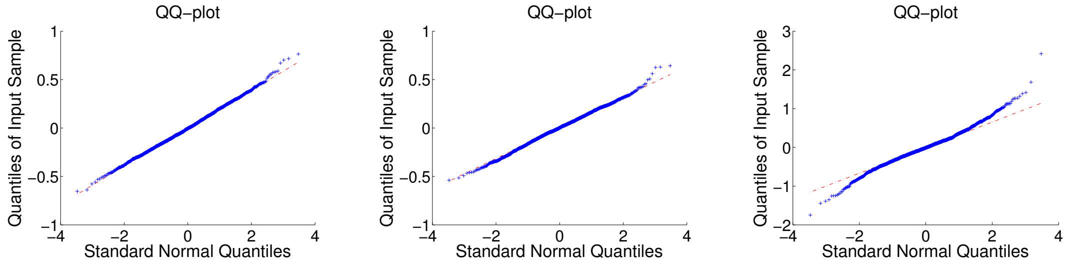

(a) (b) (c)

Figure 5.4: QQ-plot of a linear detrended version of the temperatures above.

Chapter 6

The temperature data

The temperature data used are obtained by Climatic Research Unit (CRU) at the Univeristy of East Anglia and Met Office Hadley Centre, UK. They provide several datasets for private studies and scientific research, and have recently made an update of the datasets with better estimates and bias corrections than previous versions. Some of the datasets they provide are:

- CRUTEM4: A $5° \times 5°$ gridded dataset consisting of land temperatures with monthly resolution from 1850 to today. A description of this dataset is given by [*Jones et al.* (2012)].
- HadSST3: A $5° \times 5°$ gridded dataset consisting of sea surface temperatures with monthly resolution from 1850 to today. A description of this dataset is given by [*Kennedy et al.* (2011a), (2011b)].
- HadCRUT4: A $5° \times 5°$ gridded dataset with monthly resolution from 1850 to today. The data are a blend of temperatures from CRUTEM4 and HadSST3 to give a global coverage. A description of this dataset is given by [*Morice et al.* (2012)].

These datasets and hemisperic and global means of these temperatures can be downloaded from http://www.cru.uea.ac.uk/cru/data/temperature/. The temperatures in these datasets are given as temperatures deviations from the mean value of the given month in the period 1961-1990. This mean temperature is often called the *climatology*, and when this is subtracted the seasonal variations should not be seen very clearly in the temperature data.

For calculation of the land temperatures more than 4800 stations are used. The majority of these stations are in Europe, North America and Japan. South America,

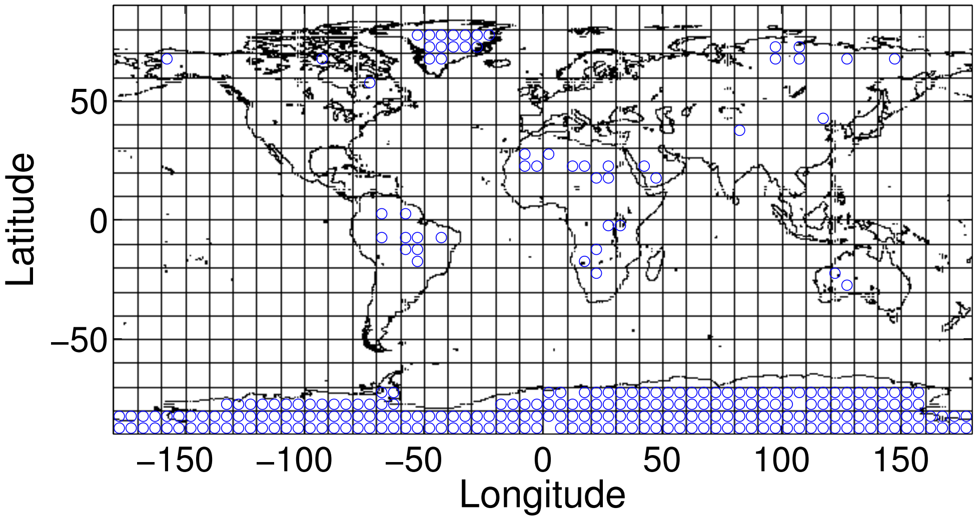

Figure 6.1: The blue circles show the grid boxes where no measurements are recorded in the in HadCRUT4 data set.

Africa and the Antarctica are not very well covered. The sea surface temperature is measured by merchant and some naval vessels, so a great part of the measurements have been made where most ships travel, and areas such as the Southern Ocan is less covered. But the coverage has become better after 1980 through the deployment of fixed and drifting buoys. Figure 6.1 shows where we don't have measurements in the combined land and sea surface temperature.

6.1 Uncertainties in the data

A detailed description of the uncertainties is given by [*Brohan et al.* (2006)] and [*Morice et al.* (2012)]. Here I will briefly summarize what the uncertainties are. The uncertainties of the gridded land data can be divided into 3 groups:

(i) Station error: The uncertainties for the individual stations.

(ii) Sampling error: The uncertainty of the average value in a grid box. It depends on the number of stations in a grid box, the positions of the stations, and on the actual variability of the climate in that box.

(iii) Bias error: Caused by changes in the measurement methods or other local effects causing bias, like urbanization or thermomenter exposure.

For the marine data we don't have only stationary points we measure from, so we cannot divide the uncertainties into the exact same groups as for land data, we rather use: (i) measurement, (ii) sampling and (iii) bias errors. Because the sea surface temperature (SST) is not measured at stationary points the errors will be spatially correlated, while the errors in land temperatures are not. The SST is measured by moving ships, and moored and drifting bouys. So where the data are recorded in a given grid box varies from month to month.

The bias in the sea surface temperature is mainly a result of a change in measurement methods. In the 19th and early 20th century the temperature was measured by ships drawing up buckets of water to their deck. In the 20th century they measured the temperature of water taken into the ships engine room, used special insulated buckets or hull contact sensors. And the recent decades buoys are used more widely. Systematic bias over long time could give rise to low-frequency components, and also give rise to wrong estimates of the trend in the data. When averaging over larger areas the station and sampling error will be smaller, but the bias error will still be the same. The bias error in marine temperatures is estimated by creating an ensemble of realizations of the temperature with different bias adjustments. Marine temperatures are in general less variable in both space and time than land temperatures, so that the sampling error, and hence the total error, will be less for marine temperatures.

It may not be obvious why we use the sea surface temperature instead of measuring the air temperature above the sea. But we will assume that it is a good estimate of the air temperature. This is often true, but in some cases it is not, like in the western boundary current in the northern hemisphere. It has been shown that that the the SST is a more useful estimate than the air temperature above the sea, because the day air temperature is largely affected by reflected solar radiation and the diurnal cycle, while the temperature in the water is more stable.

6.2 Blending marine and land data in HadCRUT4

Since the errors in HadSST3 are given by ensembles and in other ways in CRUTEM4, it is also created an ensemble of CRUTEM4 consisting of realizations of possible errors before combining the datasets. In this way it is easy to blend the datasets and calculate the combined uncertainties. When taking an ensemble of 100 realizations of land and sea surface temperature anomaly and combine them one to one, we end up with 100 realizations of the combined temperature anomaly. The combined error in a grid box is thus calculated from the ensembles, and in addition there is a sampling and measurement error added.

In the boxes where we have both land and marine data, the average temperature is calculated by weighting the land and sea surface temperatures by their fractional area. But if the land area is less than 25% it is weighted as if it was 25%, to ensure that the temperature at small islands are not swamped by the sea surface temperature.

6.3 Averaging

There are several approaches to combine the station data when calculating regional and global average temperatures. In many cases the station data are interpolated before the averages are calculated, but in HadCRUT4 this is not done. On the webpage for the land data it is stated that "*the gridding then simply averages all available station temperatures (as anomalies from 1961-90) within each grid box for each month from 1850. No account is taken of the station's elevation or location within the grid box (anomalies show little consistent dependence on altitude).*"

The results obtained from different datasets may therefore vary a little even though they are based on almost the same raw data, so [*Morice et al.* (2012)] recommends to compare results from different datasets when using the data. They have also made a comparison of global and hemispheric temperatures from the different datasets and found that they are in quite good agreement with each other.

Regional averages are calculated as a weighted average of the grid boxes, where each weight is proportional to the area of the grid boxes, and the sum of the weights is normalized to one. The global averages are calculated in slightly different ways: For HadCRUT4 it is simply the average of the Nortern and Southern hemisphere temperatures. The global average land tempereature is calculated as a weighted average of the temperature in the Nortern and Southern hemisphere, where the Northern hemisphere is weighted twice as much as the Southern hemisphere. The global average SST anomaly is a weighted sum of all available observations.

6.4 E-OBS land temperatures

The E-OBS dataset provides daily resolved land temperature data in the region 25N-75N x 40W-75E, and can be downloaded from http://eca.knmi.nl/download/ensembles/download.php. The dataset is generated by interpolating station data, and more detailed information can be found in [*Haylock et al.* (2008)]. It can be downloaded in different resolutions, and I have chosen to use the dataset with regular grid boxes of size 0.5° × 0.5°.

Chapter 7

Spatial averaging

In the study of how the Hurst exponent varies on different scales I start out with $5^\circ \times 5^\circ$ gridded data when analysing HadCRUT4, CRUTEM4 and HadSST3. To go to larger spatial scales I average over 9, 25, 49, 64, etc. grid boxes. The smallest scales are illustrated in figure 7.1, where each of the small squares denotes a $5^\circ \times 5^\circ$ grid box. Since the area of a grid box depends on latitude (see equation (A.1)), the average value of the grid boxes are weighted by the areas.

In many of the time series we find segments where we have no measurements. An example of what a time series may look like is given in the upper plot in figure 7.2. The methods I use to calculate the Hurst exponent require that the data are evenly spaced in time, so I need to fill these gaps in the time series before I do the analysis. There are several ways this could be done, and the method I have used is described below. Other ways could for example be to fill in completely random numbers, or use more advanced statistical methods to predict the most likely value given the previous values. But that

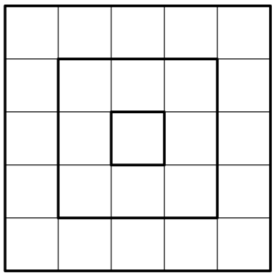

Figure 7.1: Coarse graining by box averaging.

requires that we assume a certain stochastic model in advance, which I prefer to avoid as that will make the result fit better to the model.

After filling the gaps in each time series, we meet another challenge when trying to average these time series. We cannot be sure that the surrogate data will be a good representation of the time series in an average of several time series, because spatial correlations are not taken into account in my method. If we choose to use only the raw data in the averaging we would have less data points to average over at each time point, but perhaps this is better than using data that may deviate from the real temperature.

7.1 Handling missing data points

To replace the missing data the time series are run through my "repair function". The pseudocode for this function is:

- Count how many missing measurements there are in the time series, and determine the date when the measurements start.
- Search through the time series and, at each position t in the time series, determine:
  - the length of the gap of subsequent missing measurements (in most cases this length is zero), and save the result in an array.
  - the distance from the present position in the time series to the next gap of missing points, and save the result in another array.
- Find the first position where there exists at least 6 subsequent measurements, and start the time series at this position.
- If there are no segments of the time series where there are 6 measurements in a row, then save this information.
- Search through the time series again, and for each position where we have some missing measurements, go through all the good segments with at least the same length as the missing segment. Find the segment that minimizes the expression:

$$\frac{\sqrt{(T(n)-T(j))^2+(T(n+L+1)-T(j+L+1))^2}}{2} \tag{7.1}$$

where $T(n)$ is the temperature at position n in the time series that has at least L subsequent measurements, and $T(j)$ is the temperature at the position j in the time series that has L subsequent missing values. This expression is meant to be a measure of which segment has the most similar endpoints of the interval that needs to be replaced.

After this has been done I run another test to determine if there are still gaps of missing measurements in the time series. If there are no such gaps and the repaired time series has more than 1200 data points if monthly resolved and 4000 data points if daily resolved, then I consider it good enough to be used. All other time series are discarded.

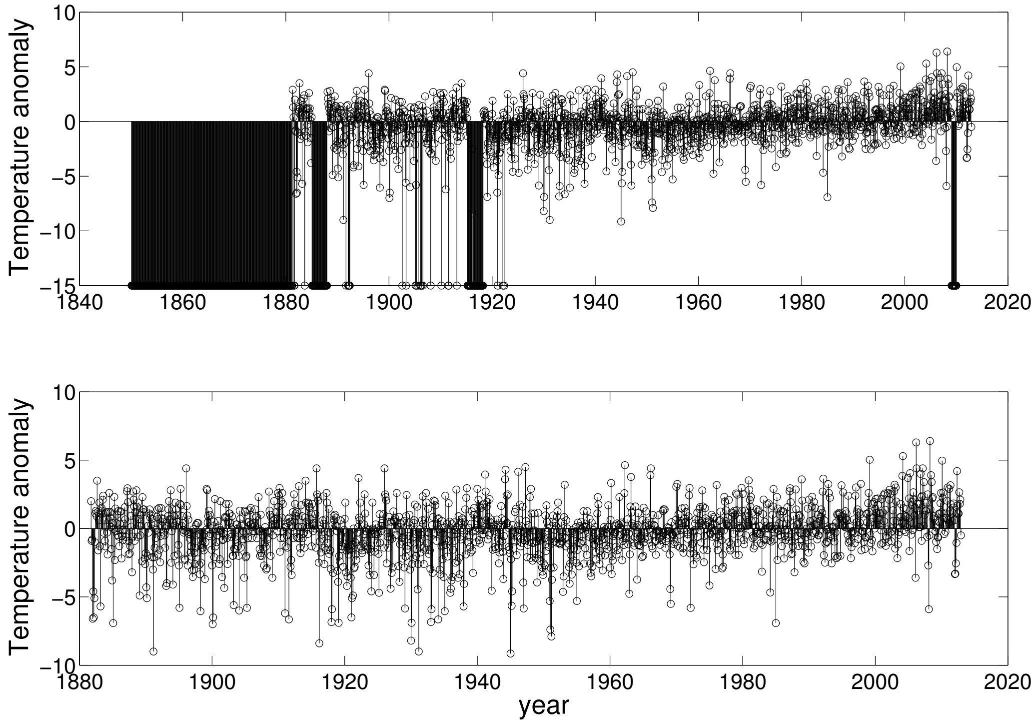

Figure 7.2: The upper figure shows the original time series for the average temperature anomaly for the grid box 40° to 45° north and 70° to 75° west. -15° means lack of measurements. The lower figure shows the time series where the bad data in the beginning are removed and the missing data later are replaced with temperatures from other parts of the time series.

7.2 Testing the method

To test the method I have generated synthetic fGns with Hurst exponents varying between -0.25 and 1.25, in steps of 0.05. At each step 100 fGns are generated, and the length of the time series is chosen to be 1955, because that is a typical length of the time series I will use. To get synthetic time series with Hurst expoents smaller than zero I have taken the difference between the elements in the fGns. And to get Hurst exponents larger than one I have taken the cumulatice sum of fGns. In chapter 3 I have studied how well the different methods can estimate the correct Hurst exponents of time series that have no missing data points that need to be replaced.

In the following a certain fraction of the synthetic time series is removed. Figure 7.3 shows the results where each of the points are removed at random, and in figure 7.4 the points that are removed is a contiguous sequence. It is obvious that if we happen to have a contiguous sequence missing in the middle of the time series that has a length greater than 1/3 of the whole time series, we can not find a similar sequence in the rest of the time series to replace it with. This missing part will then give a high contribution to the low frequencies in the spectrum, and result in an overestimated Hurst exponent. Therefore we need to check the time series when they have gone through the "repair function" if there still are some missing data. If so, these time series should not be used.

In figure 7.3 we see that the antipersistent time series are the most affected by the surrogate data, which can be explained by the introduction of more persistence in the time series when we replace the missing data with data segments that are similar to the endpoints of the interval. The real time series I will analyze are persistent, but if it one should work with antipersistent data it would be better to use another criteria for which surrogate data that fits best. The more persistent time series that we see in figure 7.3 and 7.4 can reproduce the actual Hurst exponent quite well, they just get slightly larger error bars in the estimation. The real time series will typically have a mixture of single missing data points and longer sequences of missing data, and hence we will expect the result to be something in between what we see in figure 7.3 and 7.4. Hence these analyses supports the choice of requiring 1200 measurements.

When performing coarse graining by spatial averaging one also has to consider how many measurements in the chosen area that are needed at each time to get a good estimate for the average temperature. Preferably one should have temperatures from all

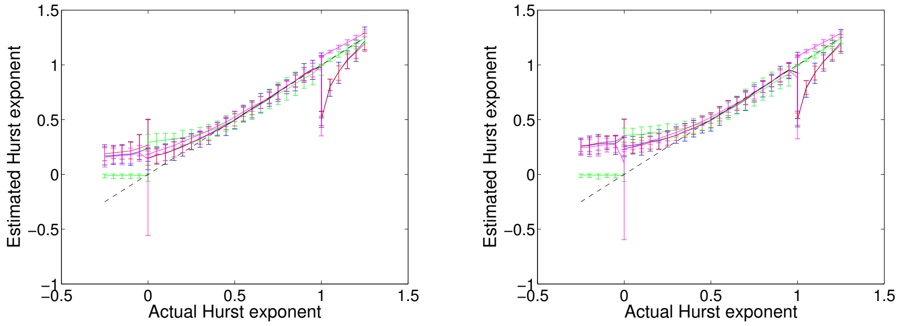

(a) 10% removed (b) 20% removed

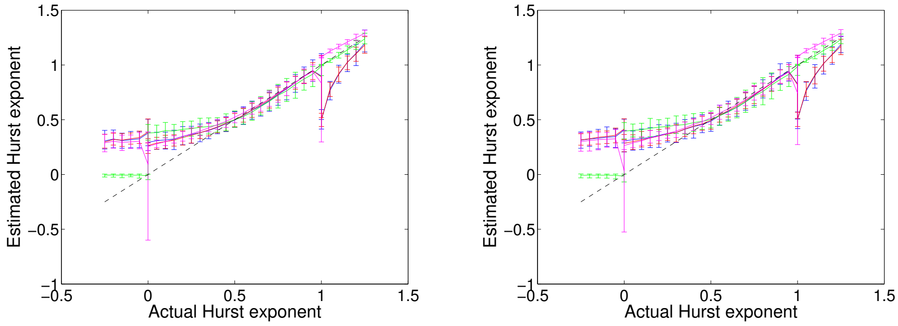

(c) 30% removed (d) 40% removed

Figure 7.3: These plots show the Hurst exponent when a part of the time series is removed at random and the missing values are replaced by other values from the time series. The coloured curves are estimated hurst exponents from periodogram (blue), wavelet variance (red), semivariogram (green) and DFA2 (pink). The error bars show the 95 % confidence intervals.

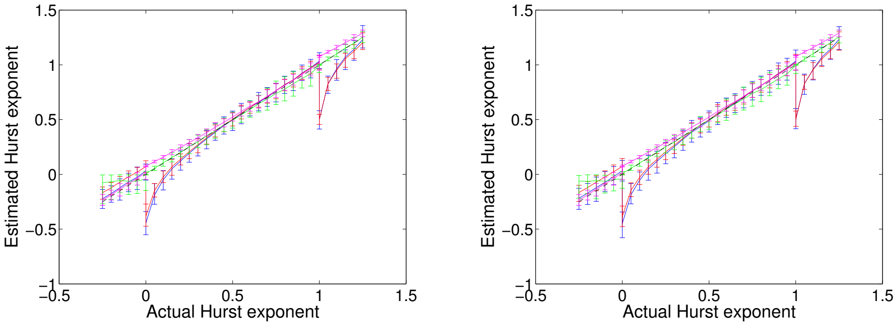

(a) 10% removed (b) 20% removed

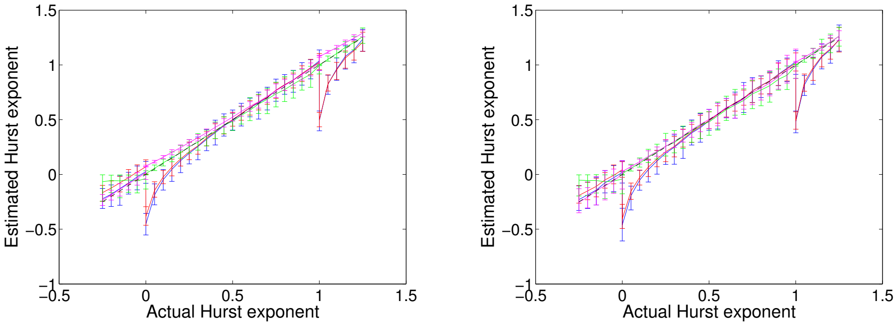

(c) 30% removed (d) 40% removed

Figure 7.4: These plots show almost the same as the last figure, except that a contiguous part of the time series is removed this time. In (d) the time series that had missing data that could not be replaced are not included in the statistics.

the grid boxes inside the area, but such a choice would discard many regions of the Earth surface and give us very poor statistics. It is also hard to test these things on synthetic independent time series, since the dependence between the real time series is what we want to study.

Coarse graining by averaging over neighboring elementary grid boxes may give misleading results if many of the elementary boxes contains missing points at a given time. A test done on large coarse-grained grid boxes of $65° \times 65°$ over oceans shows that if we require 80% of the elementary boxes to contain valid measurements, and then use the criterion of 1200 data points for repairing the resulting time series we get the Hurst exponents in figure 7.5 (a). Many of the coarse-grained series are beyond repair, and this leaves rather few coarse-grained boxes for forming the histogram. If we relax our criteria a little and require only 70% of the elementary boxes to contain valid measurements we get the result shown in (b). The relaxation of the requirement discards a smaller number of boxes and hence a more smooth histogram. In (c) we see the results we get if we let the original time series go through the "repair function" before they are averaged, and then run them through the function again after the averaging. A problem with this methodology is that the repair of individual time series from elementary boxes will reduce the spatial correlation between these time series. This correlation may be the main reason for increased Hurst exponent with increased spatial coarse graining. The results in (c) deviate quite much from (a), and the range of Hurst exponents is very large. In figure 7.6 the fits to a normal distribution for the detrended version of these temperatures are tested for the different criteria. In (a) and (b) the data look like they are normal distributed, but not in (c). Hence it seems that the criterion in (a) can be a reasonable criterion to use, even though we get less time series in the statistics. Perhaps is the criterion in (b) not that bad either, but I will choose the strict criterion in (a).

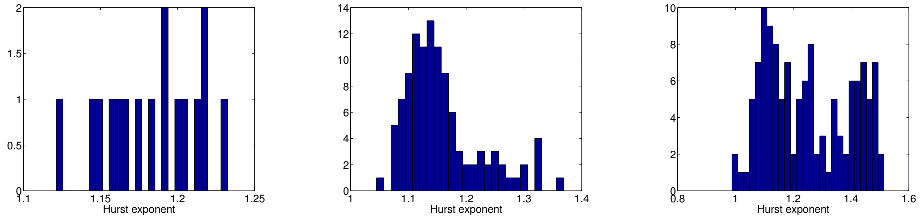

(a) Require 80%

(b) Require 70%

(c) Require 80%, but repair the time original time series first.

Figure 7.5: Histograms showing the estimated Hurst exponents from periodograms, with different criteria of which time series one should use.

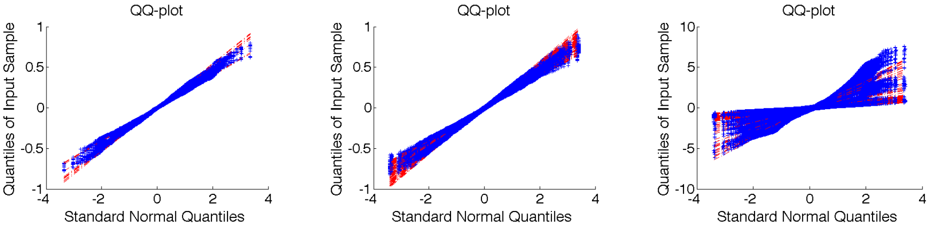

(a) Require 80%

(b) Require 70%

(c) Require 80%, but repair the time original time series first.

Figure 7.6: This figure shows QQ-plots of the detrended version of the time series in the figure above.

7.3 ACF of averaged time series

A coarse-grained temperature time series $y(t) = \frac{1}{N}\sum_i^N x_i(t)$ is the average of several other time series $x_i(t)$ of length n, which can be expected to have $\bar{x}_i = 0$ since they are temperature anomalies. The ACF for $y(t)$ will then be:

$$\begin{aligned}
\gamma_y(\tau) &= \mathbb{E}[y(t+\tau)y(t)] \\
&= \frac{1}{N^2}\mathbb{E}\big[\sum_i x_i(t+\tau)\sum_j x_j(t)\big] \\
&= \frac{1}{N^2}\mathbb{E}\big[\sum_i x_i(t+\tau)x_i(t) + \sum_i x_i(t+\tau)\sum_{j\neq i} x_j(t)\big] \\
&= \frac{1}{N^2}\sum_i \Big(\mathbb{E}[x_i(t+\tau)x_i(t)] + \sum_{j\neq i}\mathbb{E}[x_i(t+\tau)x_j(t)]\Big) \\
&= \frac{1}{N^2}\sum_i \Big(\gamma_{x_i}(\tau) + \sum_{j\neq i}\gamma_{x_i x_j}(\tau)\Big)
\end{aligned}$$

where $\gamma_{x_i x_j}(\tau)$ is the cross-covariance function of the two jointly stationary time series x_i and x_j. It can be estimated by:

$$\widehat{\gamma}_{x_i x_j}(\tau) = \frac{1}{n}\sum_{t=1}^{n-\tau} x_i(t+\tau)x_j(t) \tag{7.2}$$

and the autocovariance function $\gamma_{x_i}(\tau)$ of x_i can be estimated the same way, by just replacing the j in equation (7.2) by i. We have assumed that the ACF will have a power law dependence on time lag, given by: $\gamma_{x_i}(\tau) \propto H_i(2H_i - 1)\tau^{2H_i-2}$ where the Hurst exponents H_i are in general different.

In the simplest possible case where we assume that we average over independent time series with the same Hurst exponent, the cross-covariances will be zero, such that

$$\gamma_y(\tau) \propto \frac{1}{N^2}\sum_i \tau^{2H-2} = \frac{1}{N}\,\tau^{2H-2} \tag{7.3}$$

which has the same Hurst exponent as the original time series. As we will see later this is not the case, so this simple model is not enough. Typically the time series averaged over will have slightly different Hurst exponents, and the spatial correlations between the time series can be strong. To investigate the strength of the cross-covariance between surface temperatures I have estimated it using equation (7.2) on the $5° \times 5°$ gridded HadCRUT4 data set. In this case missing measurements are not replaced in any way, they are just

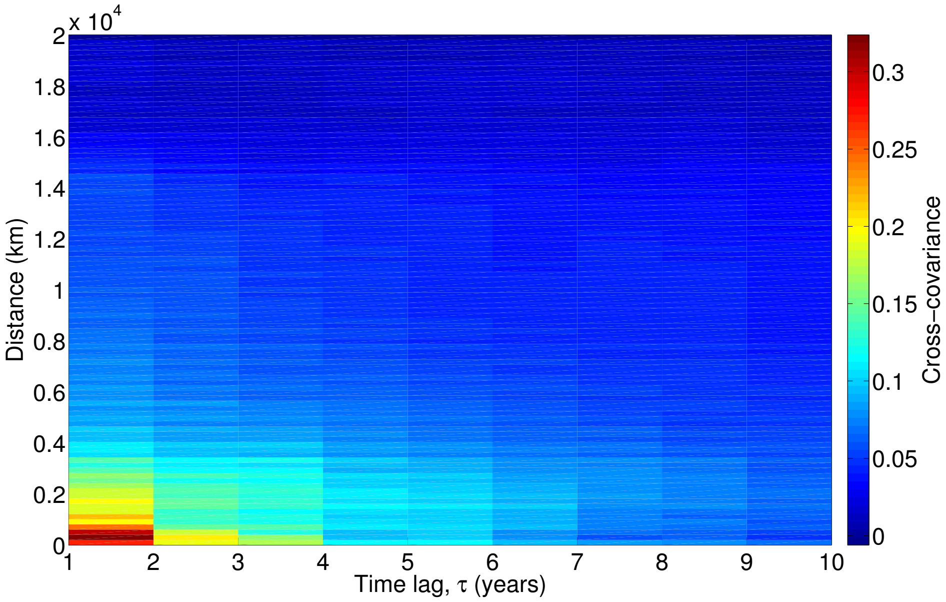

Figure 7.7: Averaged cross-covariances

ignored in the estimation. The cross-covariance is expected to depend on several factors; distance between grid boxes, if we are on land or sea, the topography, the time lag, the area of the grid boxes, etc. To simplify the results I have studied only the dependence on time lag and distance between the grid boxes. The differences in area between these small grid boxes are assumed to be small enough that we ignore them. The results we then get are shown in figures 7.8 and 7.9.

As we can see from figure 7.9 the cross-covariance seems to have a power-law dependence on time lag up to around $1.4 \cdot 10^4$ km distance, and the strength of this slope also depends on the distance between the points. Hence I will try to fit the results to a model:

$$\gamma_{x_i x_j}(\tau) = F(|x_i - x_j|)\tau^{-\alpha(|x_i - x_j|)}$$

where $d \equiv |x_i - x_j|$ is the distance between the two points and τ is the time lag. To calculate d I have used a formula that gives the distance between two points along a sphere:

$$d = r \cdot \arccos\left(\sin(\phi_s)\sin(\phi_f) + \cos(\phi_s)\cos(\phi_f)\cos(\Delta\lambda)\right) \tag{7.4}$$

(see for instance [*Gade* (2010)]) where $r = 6.371.009$ m is the value I have used for the Earth radius, ϕ_s is the latitude of the first point, ϕ_f is the latitude of the second point,

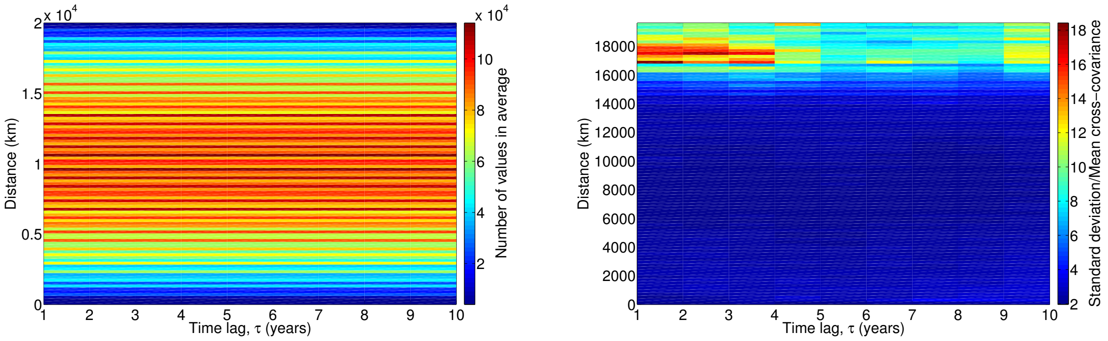

(a) Numbers of values in average (b) Relative variance of the values

Figure 7.8: Figure (a) shows that we have many time series to average over, such that we have good statistics for most of the values used in 7.8. Figure (b) tells us that the cross-covariances differ a lot.

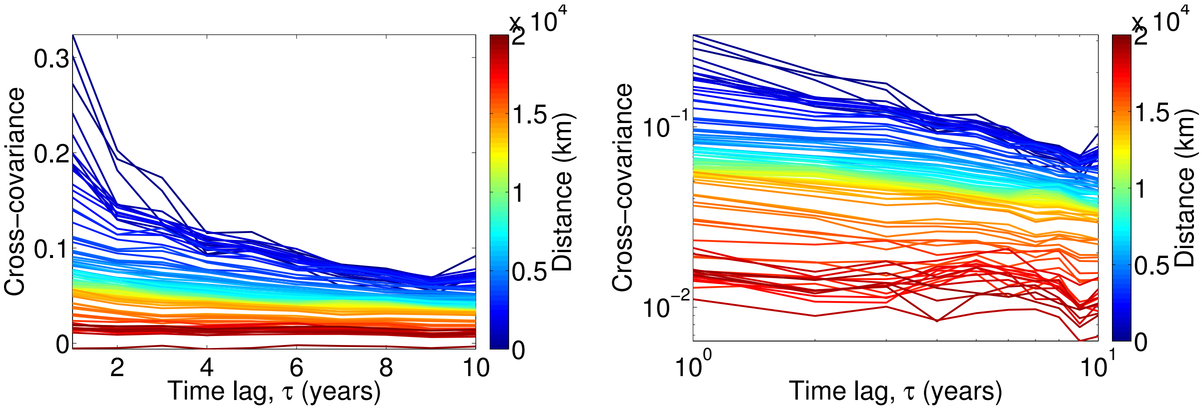

(a) (b)

Figure 7.9: Another way to visualize the results in figure 7.8. In figure (b) the two largest distances are not included because they contained negative cross-correlations.

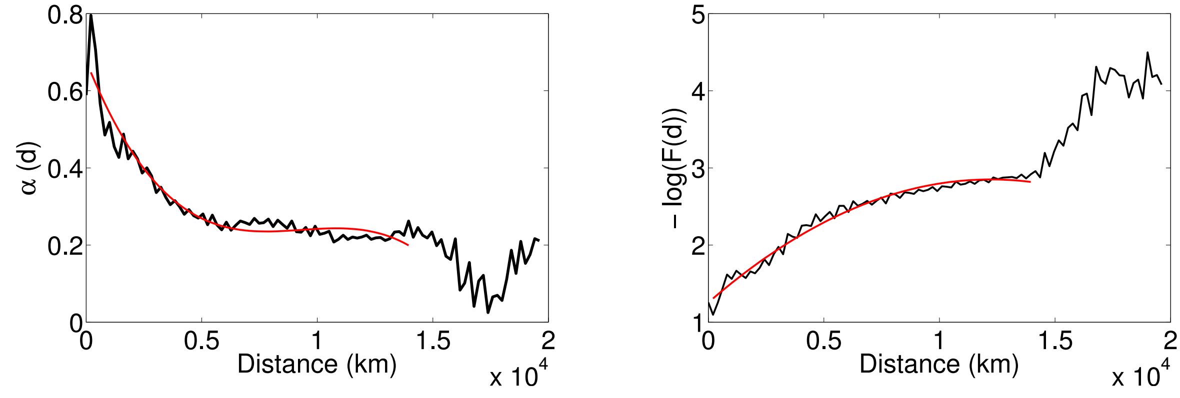

(a)

(b)

Figure 7.10: These plots show the slopes and the constant terms in the linear fits of the plots in figure 7.9 (b). The red graphs are a third order polynomial fit (a) and a second order polynomial fit (b).

and $\Delta\lambda$ is the difference in longitude between the two points.

The model parameters are calculated by doing a linear regression to the expression

$$\log(\gamma_{x_i x_j}(\tau)) = \log(F(d)) - \alpha(d) \cdot \log(\tau)$$

which is plotted in figure 7.9 (b), and the results of this regression are shown in figure 7.10.

I find that

$$-\log(F(d)) \approx p(d) = 1.2551 + 0.2608 \cdot 10^{-3} d - 0.1067 \cdot 10^{-7} d^2$$

$$\alpha(d) \approx 0.6772 - 0.1492 \cdot 10^{-3} d + 0.16468 \cdot 10^{-7} d^2 - 0.5900 \cdot 10^{-12} d^3$$

So an approximated model for the cross-covariances is:

$$\gamma_{x_i x_j}(\tau) \approx e^{-p(d)} \cdot \tau^{-\alpha(d)} \tag{7.5}$$

for $0 < d < 1.4 \cdot 10^4$ km, and $1 < \tau < 10$ years.

A model for the autocovariance function can then be:

$$\gamma_y(\tau) = \frac{1}{N^2}\left(\sum_{i=1}^{N} \gamma_{x_i}(\tau) + \sum_{i=1}^{N}\sum_{j\neq i} \gamma_{x_i x_j}(\tau)\right) \tag{7.6}$$

$$\approx \frac{1}{N^2}\left(\sum_{i=1}^{N} \gamma_{x_i}(\tau) + 2\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} e^{p(d)} \cdot \tau^{\alpha(d)}\right) \tag{7.7}$$

If instead of taking a normal average over the gridded temperatures, we take a weighted average $y(t) = \frac{1}{\sum_i A_i} \sum_{i=1}^{N} A_i x_i$, where A_i is the area of the grid i, then we get the autocovariance function:

$$\gamma_y(\tau) = \frac{1}{\left(\sum_{i=1}^{N} A_i\right)^2} \sum_{i=1}^{N} \mathbb{E}\left[A_i^2 x_i(t+\tau)x_i(t) + \sum_{j\neq i} A_i A_j x_i(t+\tau)x_j(t)\right] \tag{7.8}$$

$$= \frac{1}{\left(\sum_{i=1}^{N} A_i\right)^2} \left(\sum_{i=1}^{N} A_i^2 \gamma_{x_i}(\tau) + \sum_{i=1}^{N}\sum_{j\neq i} A_i A_j \gamma_{x_i x_j}(\tau)\right) \tag{7.9}$$

[*Huybers and Curry* (2006)] suggest that the Hurst exponent is latitude dependent because of the variations in the strength of the annual cycle between the equator and the higher latitudes. Even though we ignore the peak at the annual frequency in the power spectrum when fitting a slope, there will be some nonlinear processes transferring energy from the annual frequency to the frequencies around, which results in a smaller difference in the power at long timescales and shorter timescales, and hence a lower Hurst exponent. The model I have used here does not include latitudinal dependence, so when testing the model I will approximate $\gamma_{x_i}(\tau)$ using the mean Hurst exponents at a given latitude calculated from the HadCRUT4 dataset using wavelet variance. This dependence is shown in figure 7.11. For latitudes higher than 60° the values along the blue curve are used. Using this ACF and equation (7.5) in equation (7.9) gives the Hurst exponents in figure 7.12 when we average over zonal grid boxes. If we follow one latitude in this figure we can see that the Hurst exponent in general is increasing when increasing the area we average over. Around equator we see no clear increase, perhaps because the model is too simple. The very large standard deviations for the cross-covariance tells us that there are large spatial variations. So it could be that the cross-covariance is also latitude dependent, mainly because of ENSO, such that we could get an increase also at these latitudes. The highest Hurst exponents in figure 7.12 are only around 0.85, so we still have a long way to go to get the Hurst exponent observed for global temperature. I suspect that the cross-covariance at longer time scales than 14000 km that I have not included in my model is crucial to increase the Hurst exponent to the global one. The lack of a good power law here is probably also related to ENSO somehow.

In equation (7.9) we can see that the autocovariance $\gamma_y(\tau)$ is determined by the sum of terms that also contains power laws. So the power law of the spatially averaged time series is dependent on which one of the two sums are dominant in this expression. The

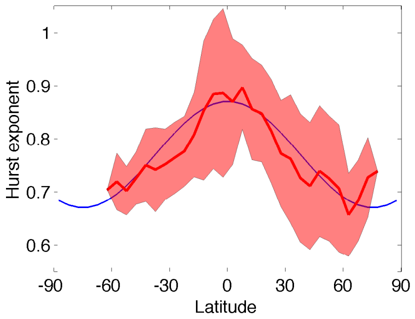

Figure 7.11: The latitudinal dependence of the Hurst exponent, calculated by averaging the values shown in figure 9.1 over all longitudes. The blue curve is a sine that is used to give us an approximation for the Hurst exponent at the highest latitudes.

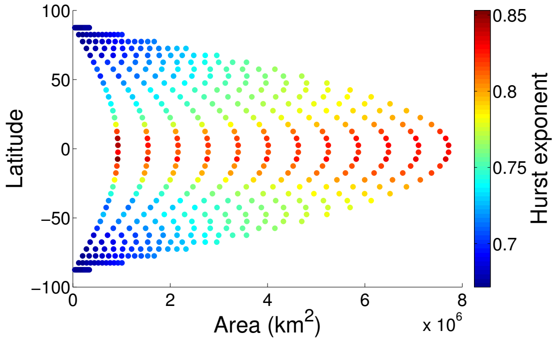

Figure 7.12: This figure shows how the Hurst exponent is changing with latitude and area of the coarse-grained grid boxes according to my model when I average zonally.

power-law exponent α for the cross-covariance can be compared to the exponent we have for the autocovariance, which is related to the Hurst exponent by $\alpha = 2 - 2H$. Then we can think of the H that we get for the cross-covariance as a Hurst exponent for spatial correlations, and see that according to this simple model the highest Hurst exponent we could get from spatial cross-covariances is $H \approx 0.9$. This is higher than most of the Hurst exponents observed for the smallest grid boxes, and could hence explain an increase in the Hurst exponent for spatially averaged temperatures if the cross-covariances are more dominant than the autocovariances. But for the global H to be higher than one, we actually need $\alpha(d)$ to be less than 0 for some values of d. In other words, we need some processes that are more strongly correlated at longer time scales than at the shorter.

7.3.1 Analysis of land temperatures only

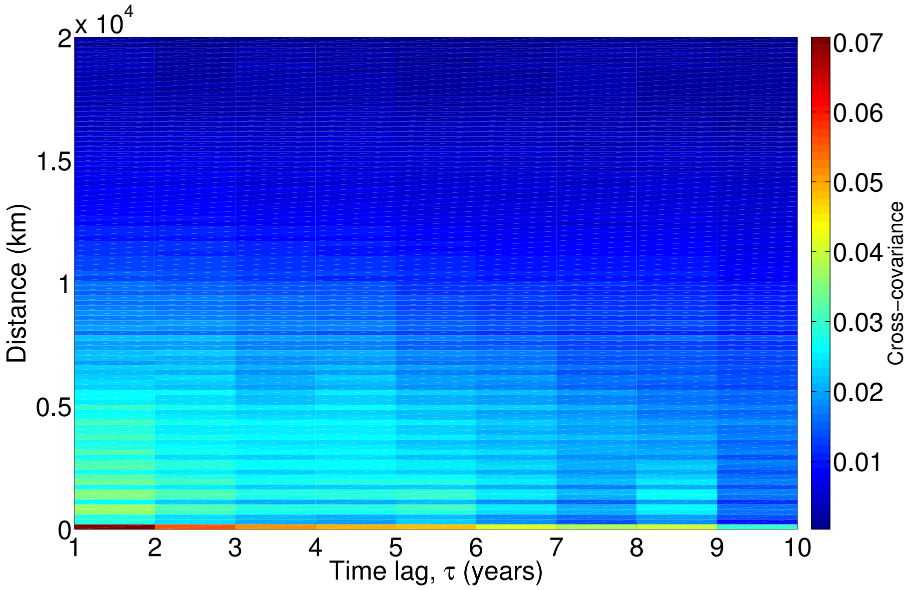

Figure 7.13: Averaged cross-covariances

When repeating the same analyzes as above for land data only I find that no nice polynomial fit can be made to figure 7.16 (a), but in (b) we have an almost perfect fit to a second order polynomial:

$$-\log(F(d)) \approx p_{\text{land}}(d) = 3.3769 + 0.2447 \cdot 10^{-5}d + 0.7225 \cdot 10^{-8}d^2$$

Perhaps are the cross-covariances in landdata very influenced by the topography in addition to the distance. Many of the points used in the calculation have a sea between them, which could also affect the cross-covariance, especially for the longest distances. Because of this I find it surprising that we get such a nice fit to a second order polynomial for $-\log(F(d))$. The value of this is very high though, implying that the factor $e^{-p(d)}$ in $\gamma_{x_i x_j}(\tau)$ will be small, and the spatial correlations are weighted less in equation (7.9). Also interesting to note is that the curve in figure 7.16 bends in the opposite direction of the curve in figure 7.10, indicating very weak cross-covariances between land locations separated by more than 10^4 km.

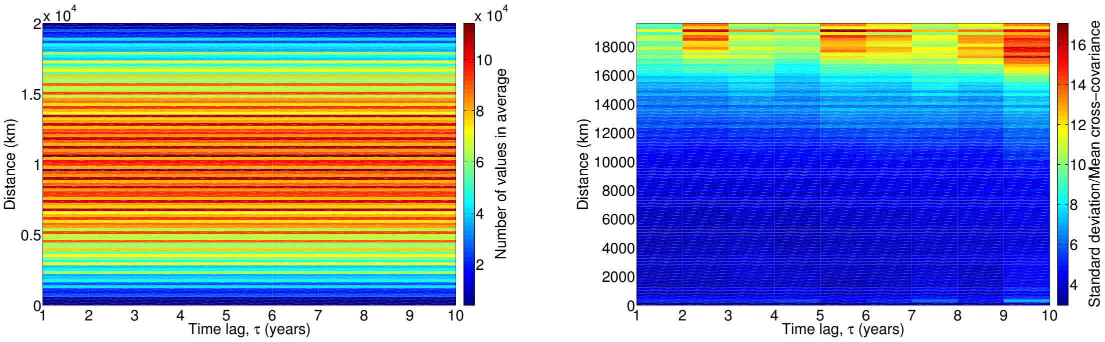

(a) Numbers of values in average (b) Relative standard deviation of the values

Figure 7.14: Figure (a) shows that we have many time series to average over, such that we have good statistics for most of the values used in 7.13. Figure (b) tells us that the cross-covariances differ a lot.

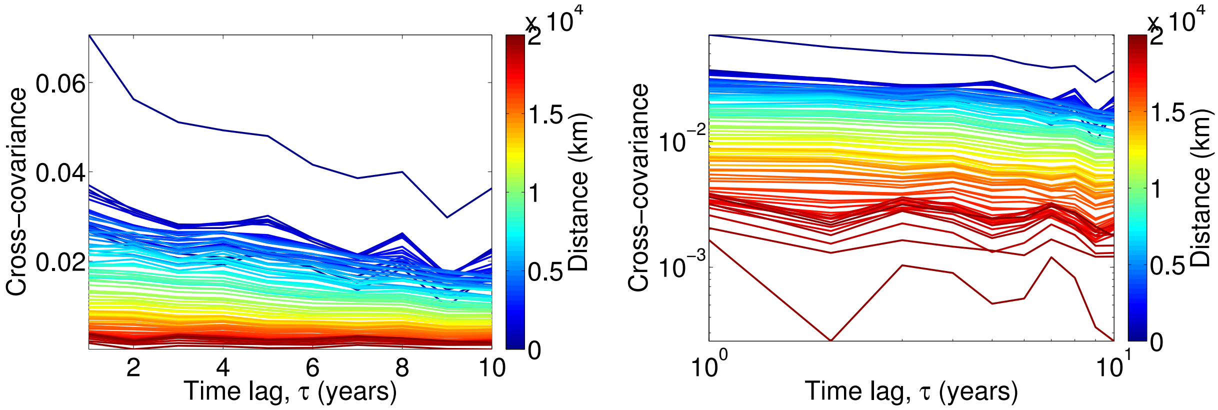

(a) (b)

Figure 7.15: Another way to visualize the results in figure 7.13.

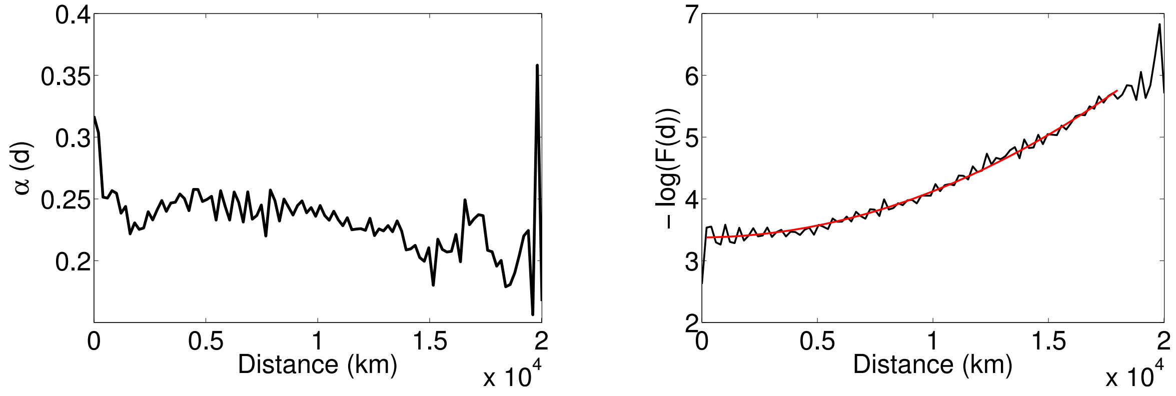

(a) (b)

Figure 7.16: These plots show the slopes and the constant terms in the linear fits of the plots in figure 7.15 (b). The red graph in (b) is a second order polynomial fit.

7.3.2 Analysis of sea surface temperatures only

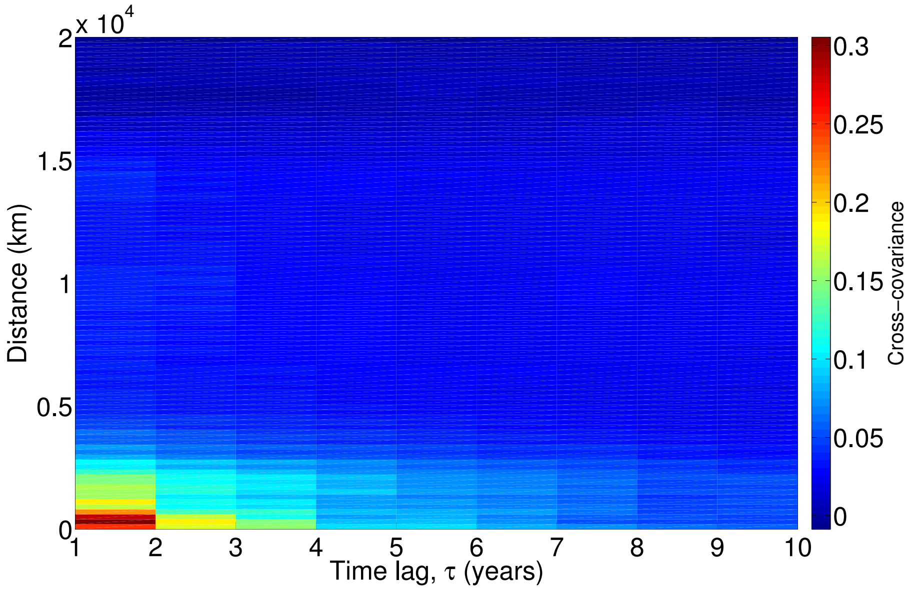

Figure 7.17: Averaged cross-covariances

The polynomial fits done in figure 7.20 are:

$$-\log(F(d)) \approx p_{\text{sea}}(d) = 0.9769 + 0.7131 \cdot 10^{-3}d - 0.7012 \cdot 10^{-7}d^2 + 0.2233 \cdot 10^{-11}d^3$$

$$\alpha_{\text{sea}}(d) \approx 0.7580 - 0.1704 \cdot 10^{-3}d + 0.1735 \cdot 10^{-7}d^2 - 0.5560 \cdot 10^{-12}d^3$$

For these analyses we can see more clearly that there are some processes resulting in an $\alpha < 0$ over the longest distances. This slope is steepest when we only study time scales up to around 6 years.

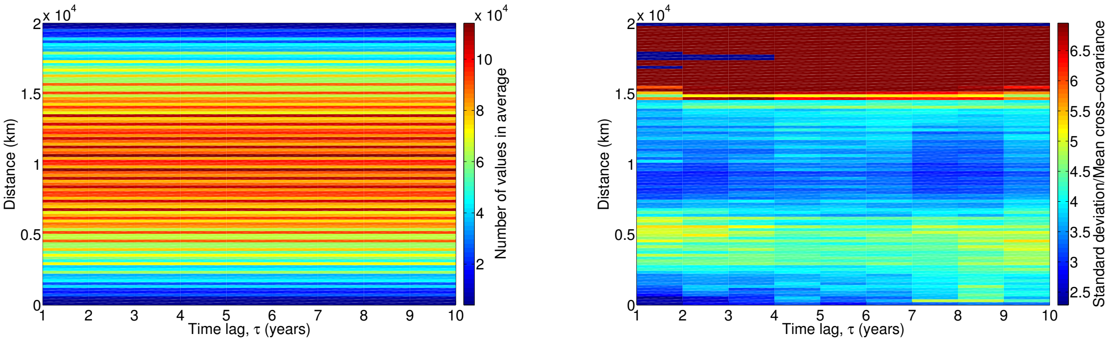

(a) Numbers of values in average (b) Relative standard deviation of the values

Figure 7.18: Figure (a) shows that we have many time series to average over, such that we have good statistics for most of the values used in 7.17. In figure (b) the some of the standard deviations have extremely large values, going far beyond the limits of the colorbar. These are represented by the darkest red and blue areas.

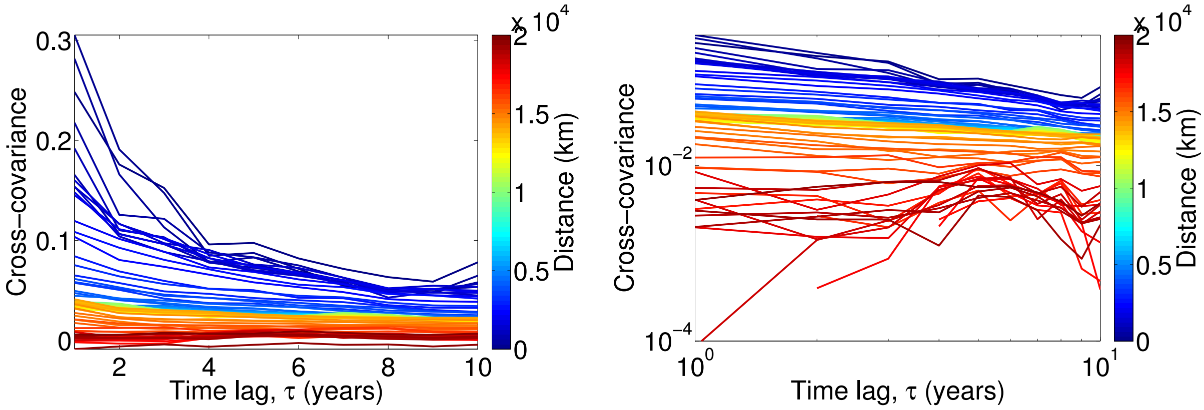

(a) (b)

Figure 7.19: Another way to visualize the results in figure 7.17. Some of the cross-covarinances are not seen in figure (b) as they are negative.

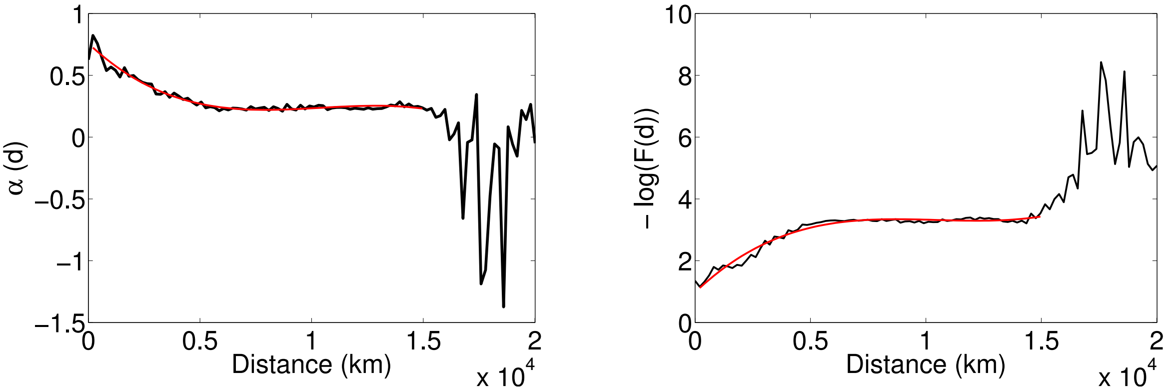

(a) (b)

Figure 7.20: These plots show the slopes and the constant terms in the linear fits of the plots in figure 7.19 (b). The red graphs are third order polynomial fits.

Chapter 8

Global and hemispheric averages

In this chapter globally averaged land, sea, and combined land and sea surface temperatures are analyzed. The straight line fit is done as close as we can get to the time scales from 3 to 128 months, but because of the logbinning these time scales may not be exact in all cases. Although the climatology is subtracted from the time series, there is still a peak in the periodogram around the frequency corresponding to the seasonal variations, especially for the land temperature. So to avoid that this will ruin the statistics all frequencies closer than 0.005 to $1/12$ or any of its harmonics are not taken into account when calculating the average values of the bins in the power in the spectrum. The steeper slopes at large scales in the wavelet variances, and for small frequencies in the periodograms are due to the trends in the data. In the wavelet variance this effect would have been eliminated by using only the part of the time series where the wavelet coefficient is computed for the central part of the time series where the entire wavelet is inside the time boundaries. This effect is particularly strong for land temperatures, which are subject to stronger trends.

8.1 Global and hemispheric averages of combined sea and land temperature

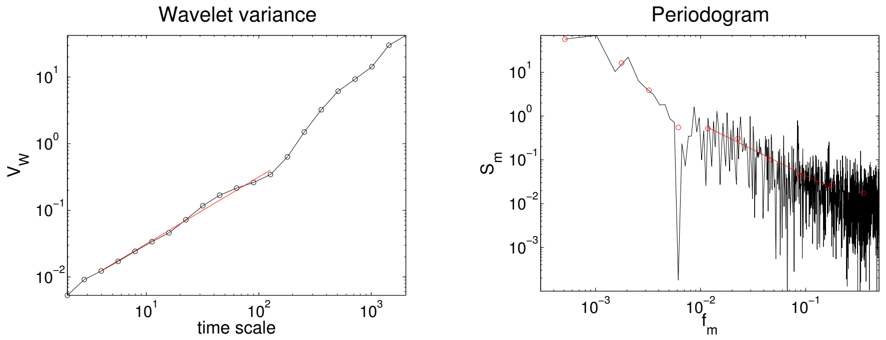

(a) Gives hurst exponent 1.00 (b) Gives hurst exponent 1.09

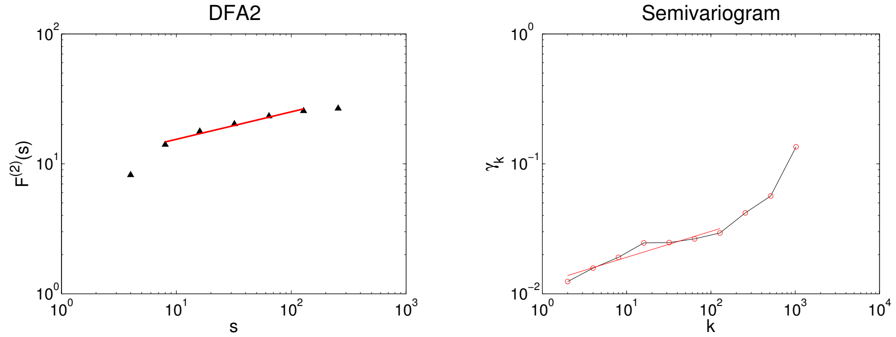

(c) Gives hurst exponent 0.97 (d) Gives hurst exponent 0.97

Figure 8.1: Global hurst exponents. In the estimation of Hurst exponent using semivariogram the time series is assumed to be stationary.

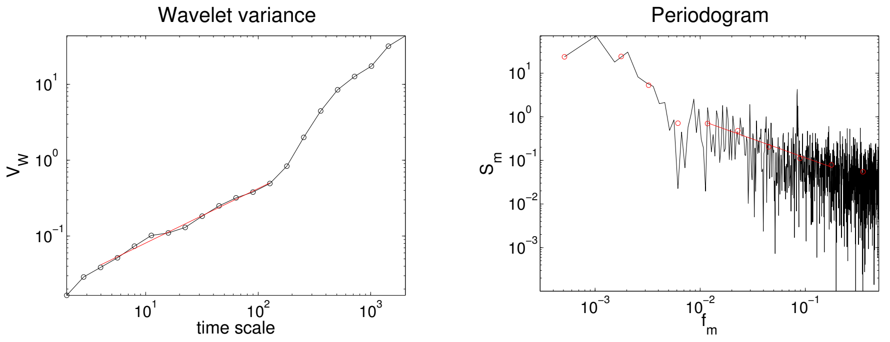

(a) Gives hurst exponent 0.86

(b) Gives hurst exponent 0.93

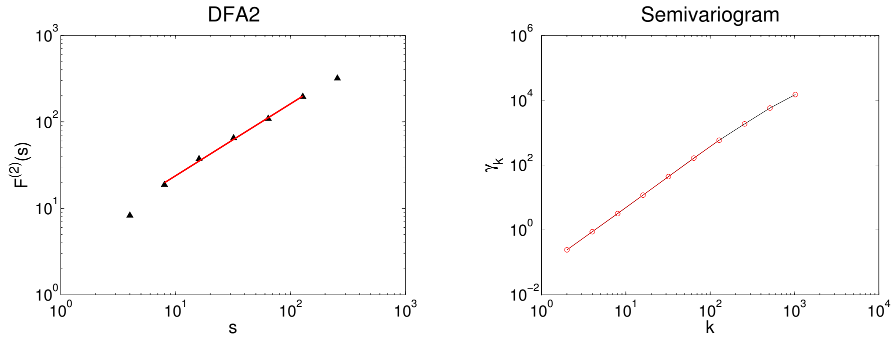

(c) Gives hurst exponent 0.83

(d) Gives hurst exponent 0.94

Figure 8.2: Hurst exponents for northern hemisphere. In the estimation of Hurst exponent using semivariogram the time series is assumed to be stationary.

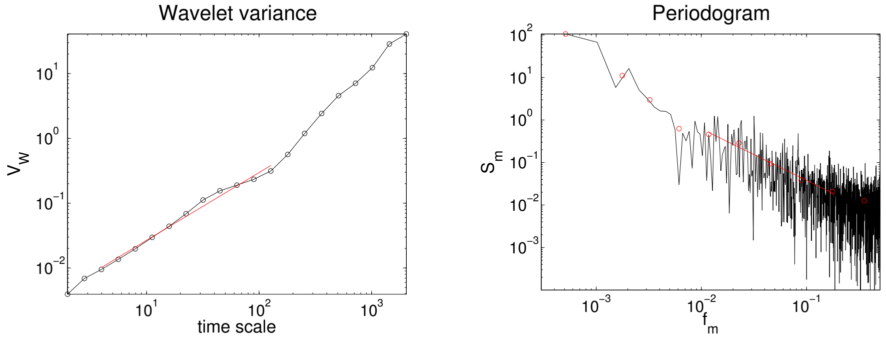

(a) Gives hurst exponent 1.02 (b) Gives hurst exponent 1.11

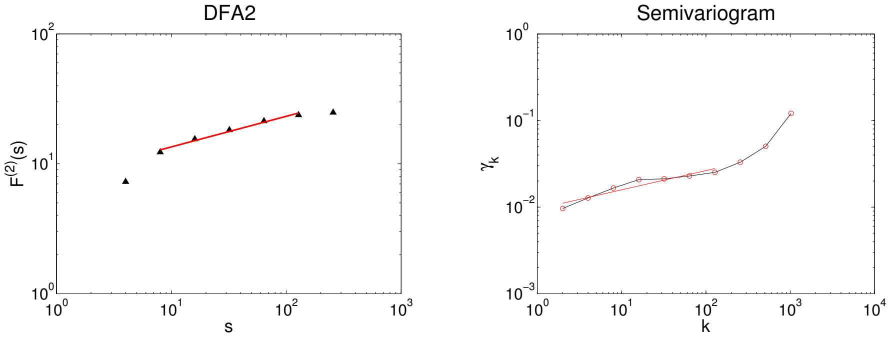

(c) Gives hurst exponent 1.02 (d) Gives hurst exponent 1.11

Figure 8.3: Hurst exponents for southern hemisphere. In the estimation of Hurst exponent using semivariogram the time series is assumed to be nonstationary.

8.2 Global and hemispheric averages of sea surface temperature

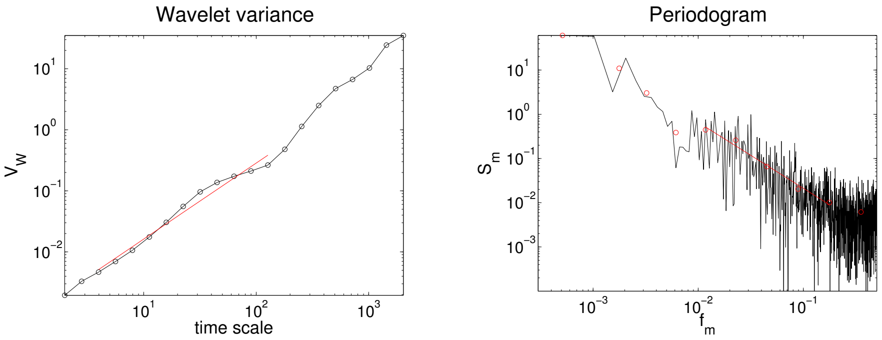

(a) Gives hurst exponent 1.12 (b) Gives hurst exponent 1.25

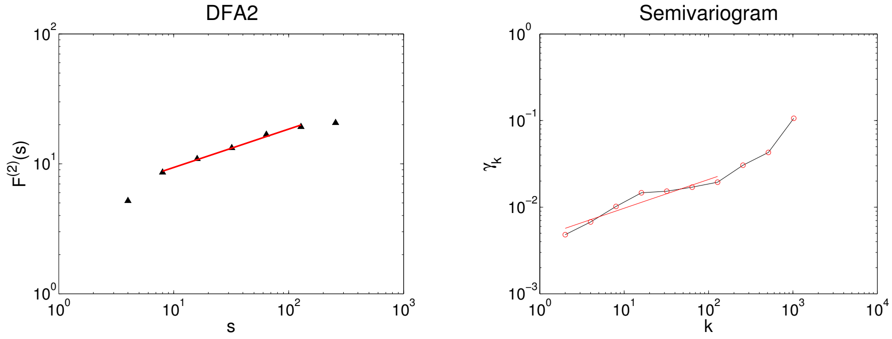

(c) Gives hurst exponent 1.14 (d) Gives hurst exponent 1.17

Figure 8.4: Global hurst exponents for the sea. In the estimation of Hurst exponent using semivariogram the time series is assumed to be nonstationary.

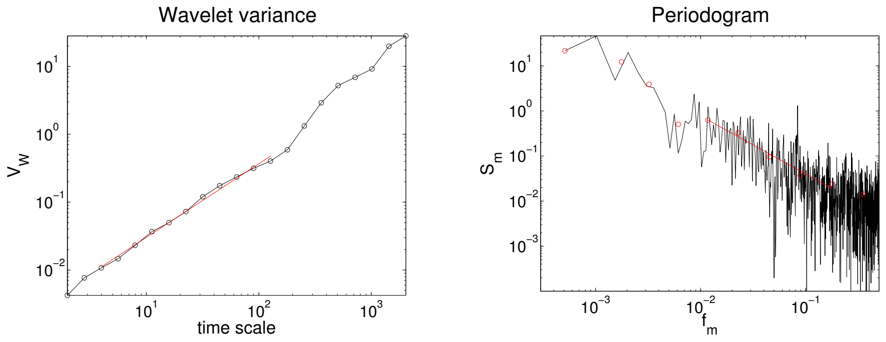

(a) Gives hurst exponent 1.04

(b) Gives hurst exponent 1.14

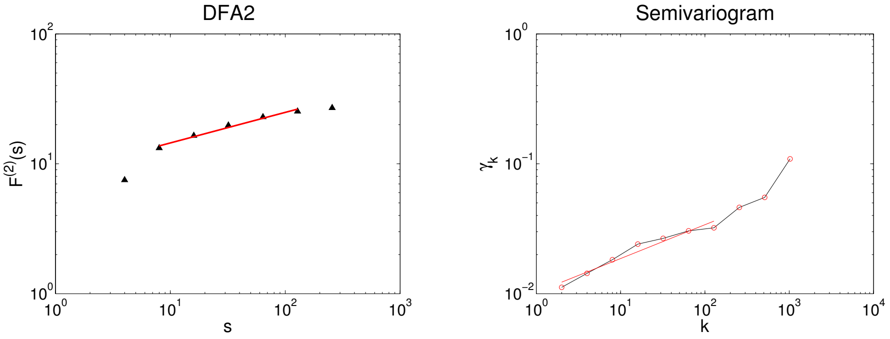

(c) Gives hurst exponent 1.03

(d) Gives hurst exponent 1.13

Figure 8.5: Hurst exponents for northern hemisphere sea. In the estimation of Hurst exponent using semivariogram the time series is assumed to be nonstationary.

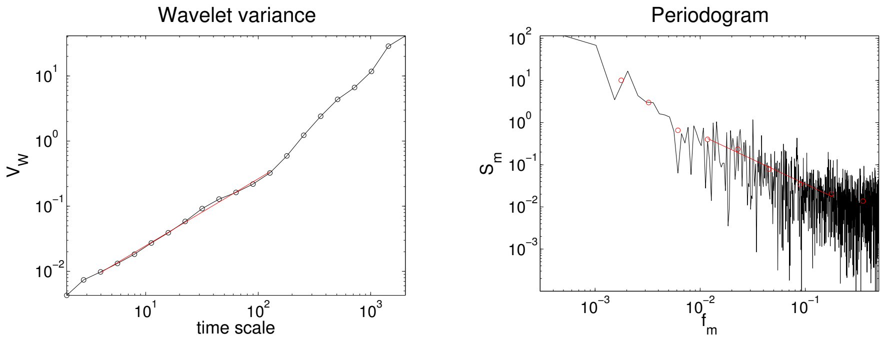

(a) Gives hurst exponent 1.01

(b) Gives hurst exponent 1.08

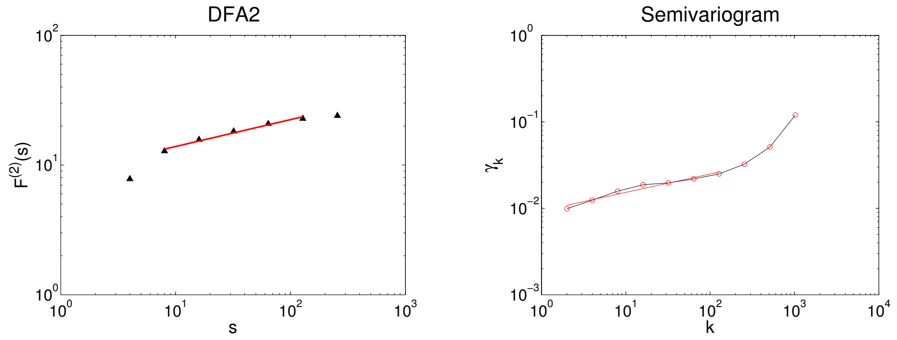

(c) Gives hurst exponent 0.98

(d) Gives hurst exponent 0.98

Figure 8.6: Hurst exponents for southern hemisphere sea. In the estimation of Hurst exponent using semivariogram the time series is assumed to be nonstationary.

8.3 Global and hemispheric averages of land surface temperature

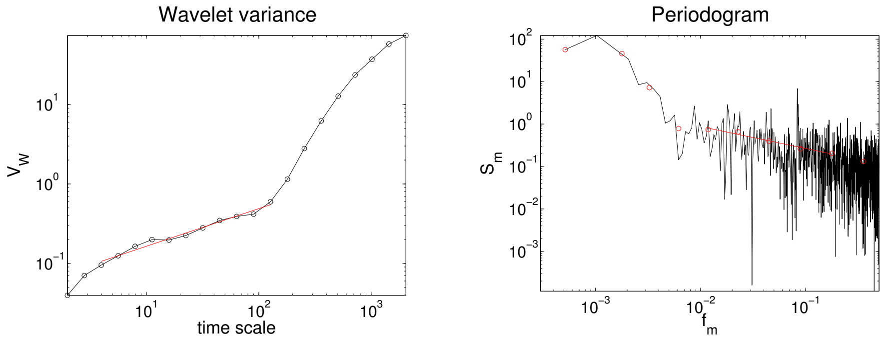

(a) Gives hurst exponent 0.74 (b) Gives hurst exponent 0.76

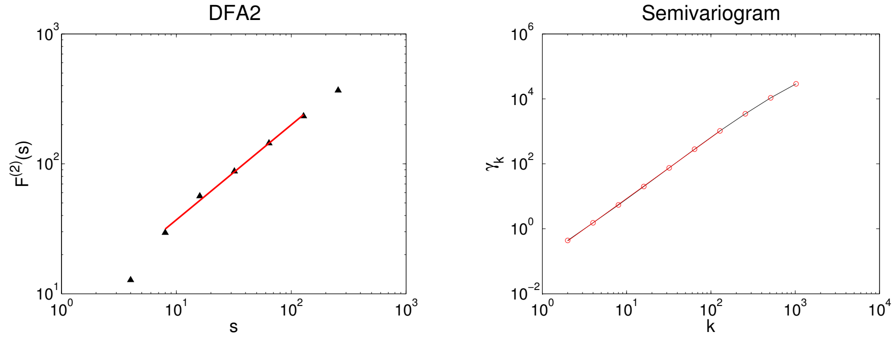

(c) Gives hurst exponent 0.73 (d) Gives hurst exponent 0.94

Figure 8.7: Global hurst exponents for land. In the estimation of Hurst exponent using semi-variogram the time series is assumed to be stationary.

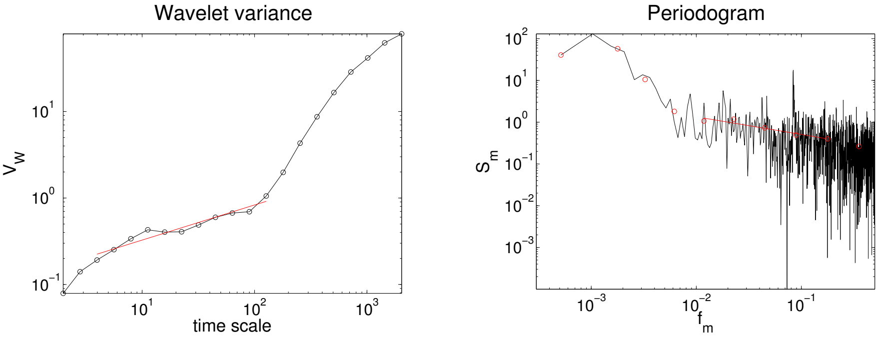

(a) Gives hurst exponent 0.70

(b) Gives hurst exponent 0.71

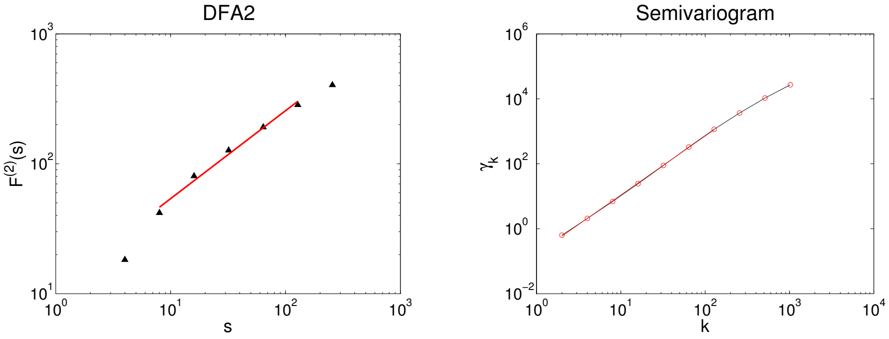

(c) Gives hurst exponent 0.68

(d) Gives hurst exponent 0.91

Figure 8.8: Northern hemisphere hurst exponents for land. In the estimation of Hurst exponent using semivariogram the time series is assumed to be stationary.

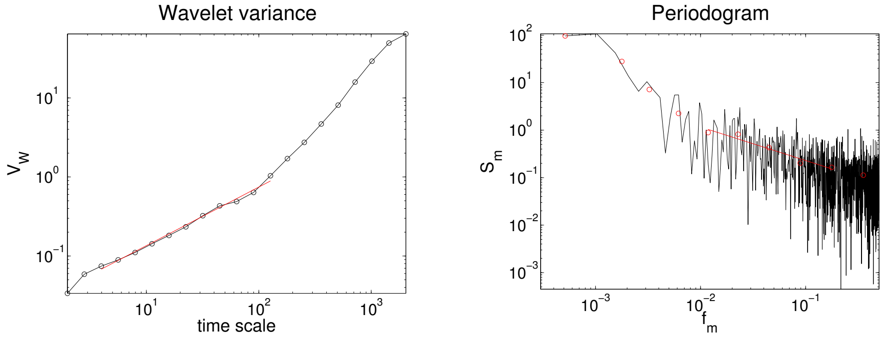

(a) Gives hurst exponent 0.87 (b) Gives hurst exponent 0.85

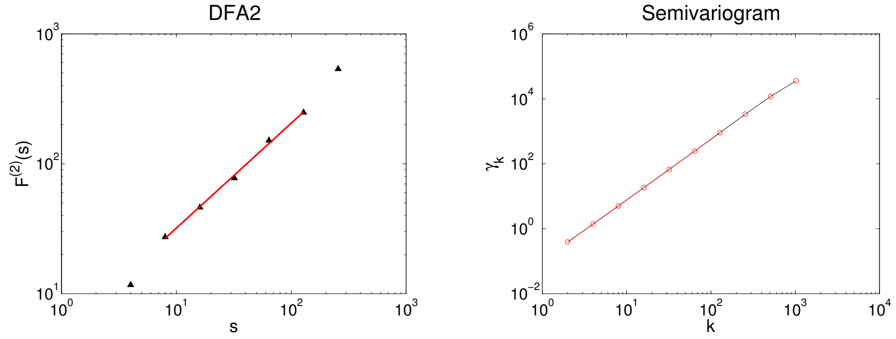

(c) Gives hurst exponent 0.81 (d) Gives hurst exponent 0.93

Figure 8.9: Southern hemisphere hurst exponents for land. In the estimation of Hurst exponent using semivariogram the time series is assumed to be stationary.

		Wavelet	Periodogram	DFA2	Semivariogram
Combined land and sea	Global	1.00	1.09	0.97	1.10 (0.97)
	NH	0.86	0.93	0.83	1.06 (0.94)
	SH	1.02	1.11	1.02	1.11 (0.97)
Only Sea	Global	1.12	1.25	1.14	1.17 (0.97)
	NH	1.04	1.14	1.03	1.13 (0.95)
	SH	1.01	1.08	0.98	1.11 (0.98)
Only Land	Global	0.74	0.76	0.73	1.03 (0.94)
	NH	0.70	0.71	0.68	1.02 (0.91)
	SH	0.87	0.85	0.81	1.06 (0.93)

Table 8.1: This table shows the global and hemispheric Hurst exponents. For semivariogram the numbers in parenthesis are the one we get if we assume stationarity, and the other numbers are the ones we get if we assume nonstationarity. We can see that the Hurst exponents for wavelet and DFA2 are very consistent, and for the periodogram they are slightly larger. That can indicate a larger influence of trends in the periodogram than I initially expected.

Chapter 9

Spatial coarse-graining of HadCRUT4

This chapter and the two following chapters will consist of figures of three types:

- A worldmap where the Hurst exponents are estimated using wavelet variance and denoted by a colour. The criterion that a Hurst exponent will be estimated is that the time series has at least 1200 values and no gaps after passing through the "repair function". For averaged temperatures it is also a criterion that for each time, 80% of the elementary cells have valid measurements. The averaged temperatures are estimated from overlapping grid boxes, and the filled coloured circles are at the center of the areas averaged over. The areas with too few measurements to get a good estimate for the Hurst exponent are left white.
- A plot that shows the shape of the periodograms and the wavelet variances that the Hurst exponents are estimated from. The grey area is where 95% of the plots are found, the red line is the median value at each time scale, and the yellow line is the mean value at each time scale.
- Histograms that shows the distribution of the Hurst exponents estimated using periodogram and wavelet variance.

9.1 $5^\circ \times 5^\circ$ grid boxes

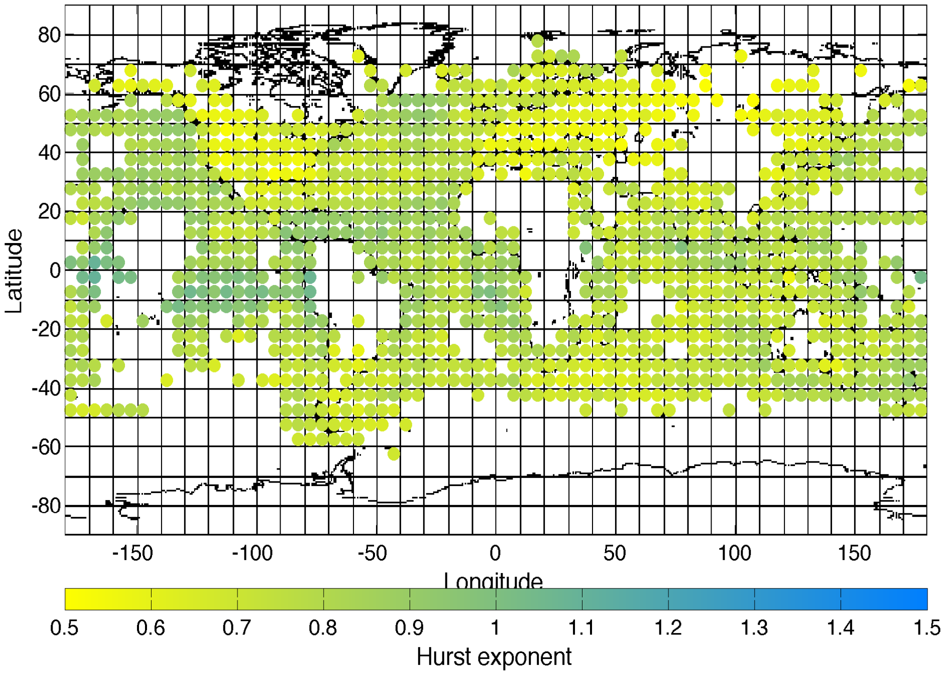

Figure 9.1: The coloured points show the 1206 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

Plots of spatially distributed Hurst exponents, like the one above, have also been made for other data sets. [*Huybers and Curry* (2006)] made one for the NCEP/NCAR reanalysis, and [*Vyushin et al.* (2012)] for a collection of several observation and model data sets. In all of the plots the land-sea difference is very clear. [*Huybers and Curry* (2006)] also find a large latitudinal dependence that they explain by the difference in annual variations between lower and higher latitudes. In their estimation procedures for the spectral exponent the annual variation itself and its harmonics are not used. The reason that the spectrum is still affected by the annual variations, they argue, are that there exist nonlinear interactions that produce phase dependencies between different frequencies. Figure 9.2 shows how my Hurst exponents varies with latitude, which supports the results found by [*Huybers and Curry* (2006)].

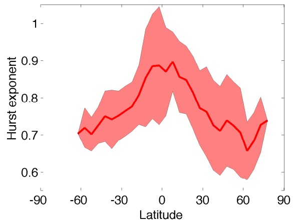

Figure 9.2: The red line shows the mean Hurst exponent, and the shaded red area shows one standard deviation.

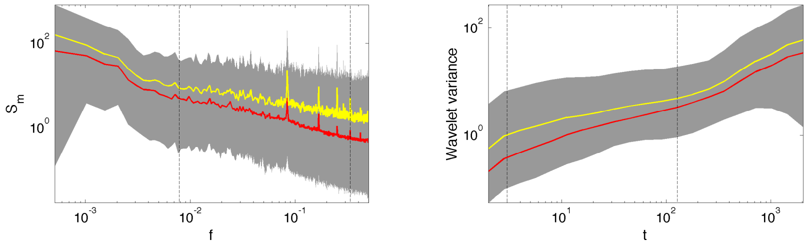

(a) Periodograms

(b) Wavelet variances

Figure 9.3: These plots show the shapes of the periodograms and wavelet variances for temperature time series from elementary grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.81$ in (a) and $H = 0.77$ in (b)) and the yellow line is the mean ($H = 0.72$ in (a) and $H = 0.70$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

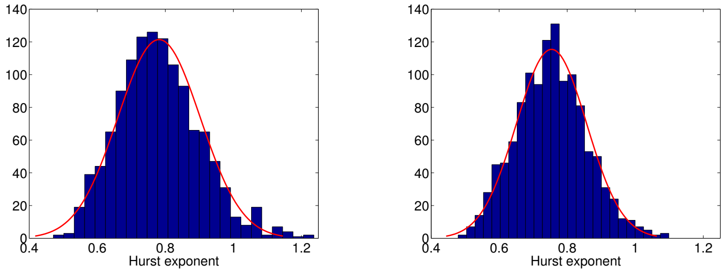

(a) From periodogram

(b) From wavelet analysis

Figure 9.4: The distribution of Hurst exponents estimated from the grid boxes shown in figure 9.1. The periodograms give mean value 0.78, median 0.77 and standard deviation 0.12, and the wavelet variances give mean value 0.75, median 0.75 and standard deviation 0.10.

9.2 $15^\circ \times 15^\circ$ grid boxes

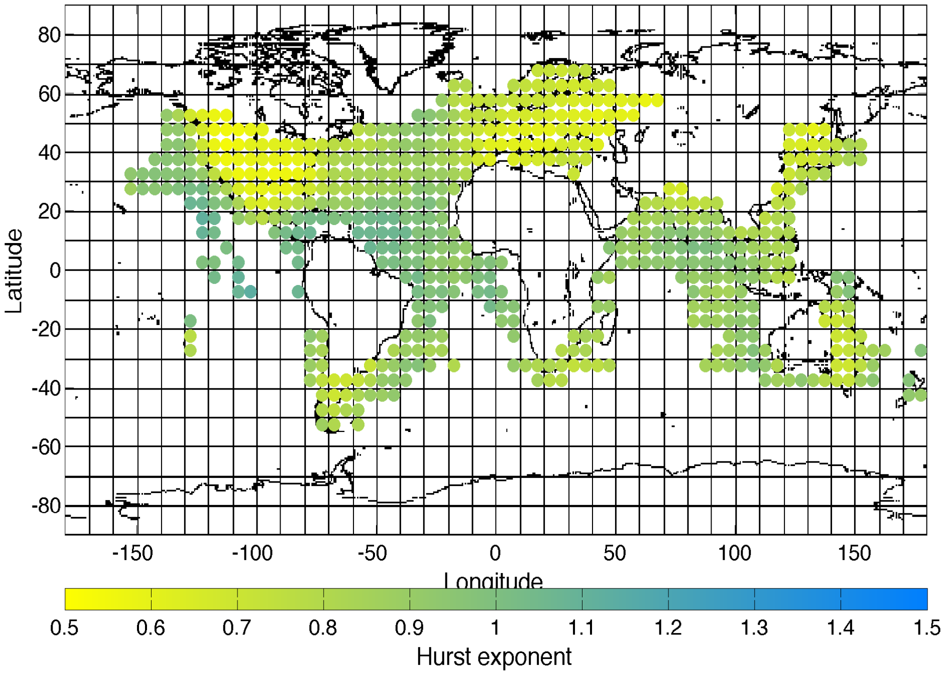

Figure 9.5: The coloured points show the 578 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

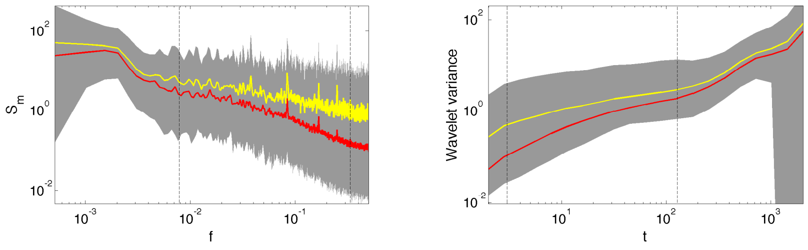

(a) Periodograms (b) Wavelet variances

Figure 9.6: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.91$ in (a) and $H = 0.86$ in (b)) and the yellow line is the mean ($H = 0.75$ in (a) and $H = 0.72$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

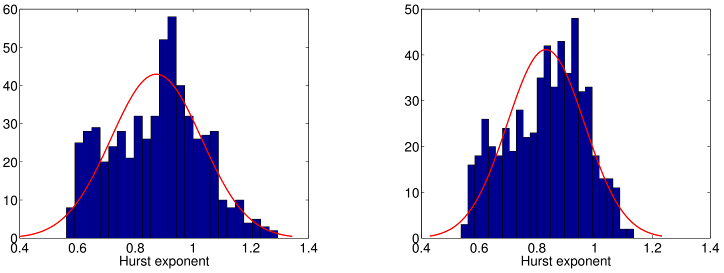

(a) From periodogram (b) From wavelet analysis

Figure 9.7: The distribution of H estimated from the coarse-grained grid boxes in figure 9.5. The periodograms give mean value 0.87, median 0.89 and standard deviation 0.16, and the wavelet variances give mean value 0.83, median 0.85 and standard deviation 0.13.

9.3 25° × 25° grid boxes

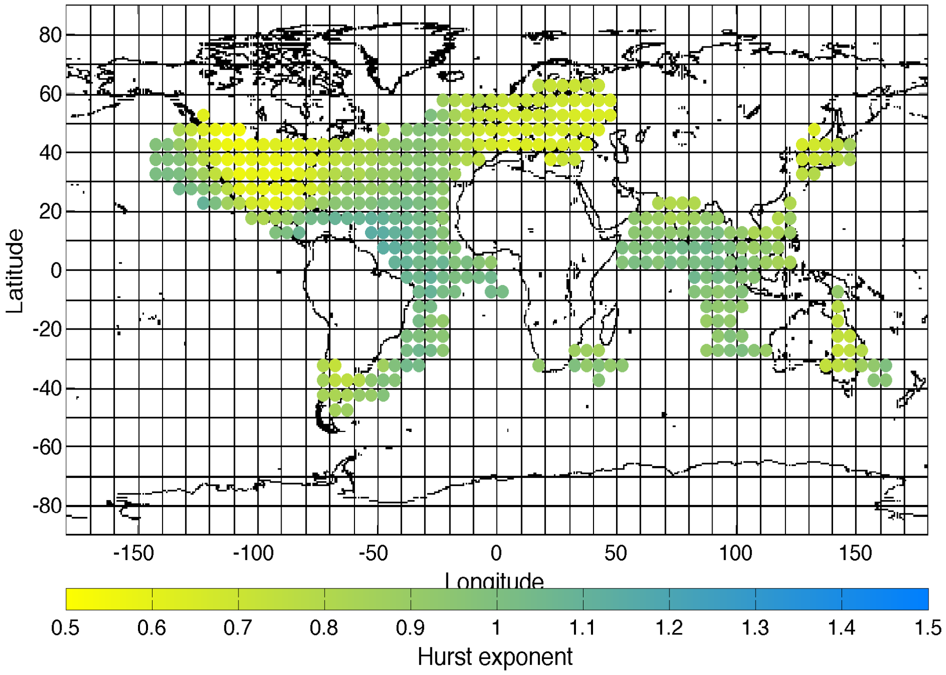

Figure 9.8: The coloured points show the 415 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

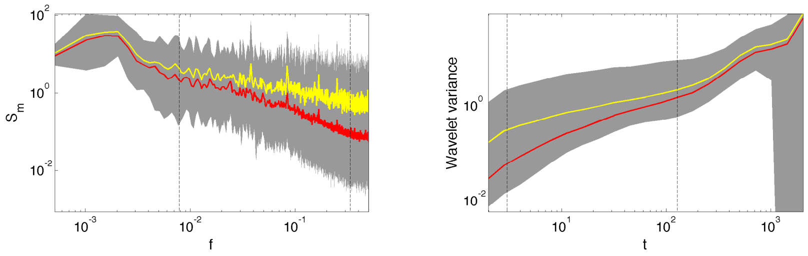

(a) Periodograms (b) Wavelet variances

Figure 9.9: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.94$ in (a) and $H = 0.91$ in (b)) and the yellow line is the mean ($H = 0.76$ in (a) and $H = 0.74$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

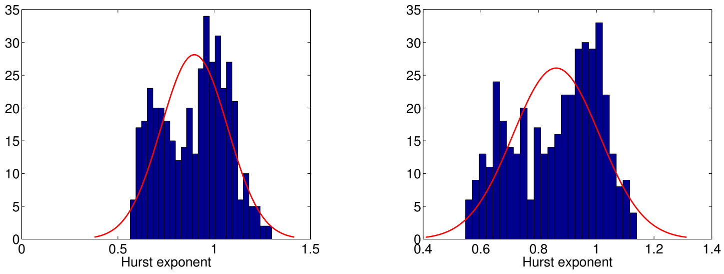

(a) From periodogram (b) From wavelet analysis

Figure 9.10: The distribution of H estimated from the coarse-grained grid boxes in figure 9.8. The periodograms give mean value 0.90, median 0.93 and standard deviation 0.17, and the wavelet variances give mean value 0.86, median 0.89 and standard deviation 0.15.

9.4 $35° \times 35°$ grid boxes

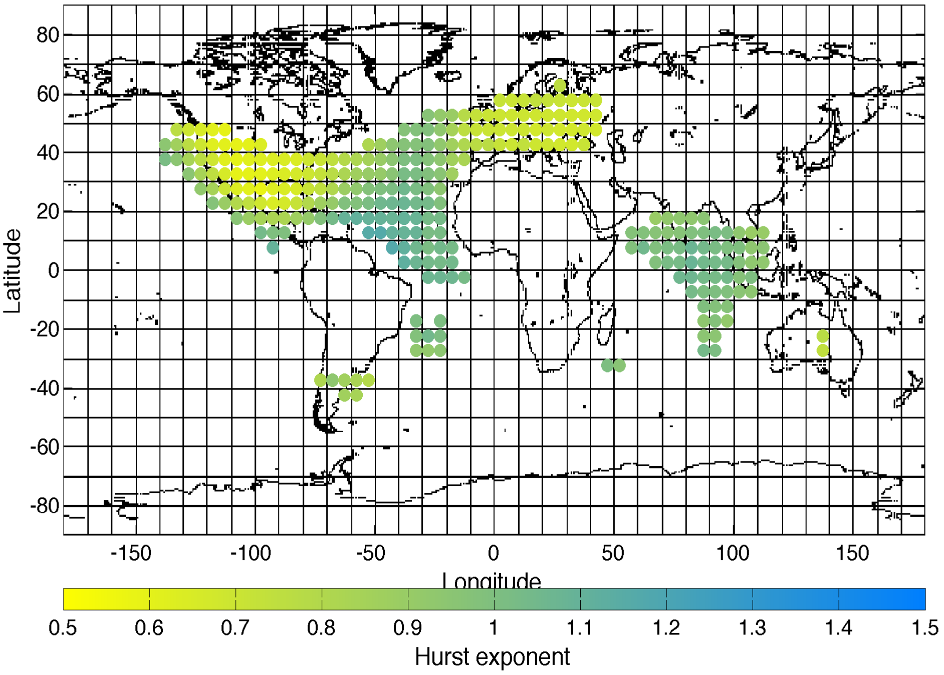

Figure 9.11: The coloured points show the 290 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

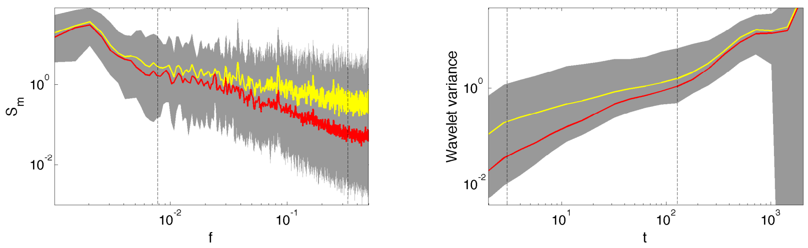

(a) Periodograms

(b) Wavelet variances

Figure 9.12: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.98$ in (a) and $H = 0.95$ in (b)) and the yellow line is the mean ($H = 0.79$ in (a) and $H = 0.76$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

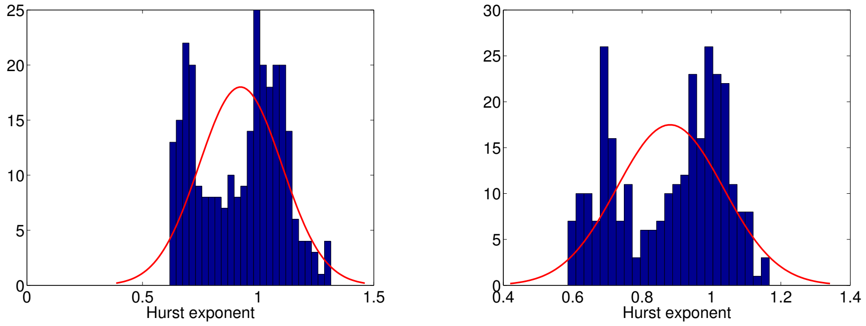

(a) From periodogram

(b) From wavelet analysis

Figure 9.13: The distribution of H estimated from the coarse-grained grid boxes in figure 9.11. The periodograms give mean value 0.92, median 0.97 and standard deviation 0.18, and the wavelet variances give mean value 0.88, median 0.93 and standard deviation 0.15.

9.5 45° × 45° grid boxes

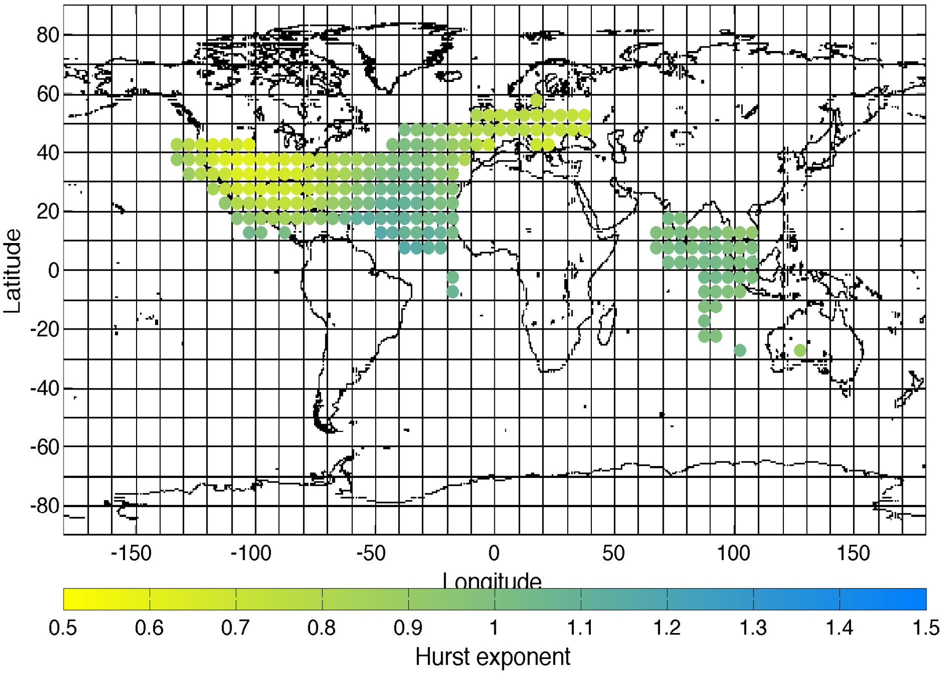

Figure 9.14: The coloured points show the 213 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

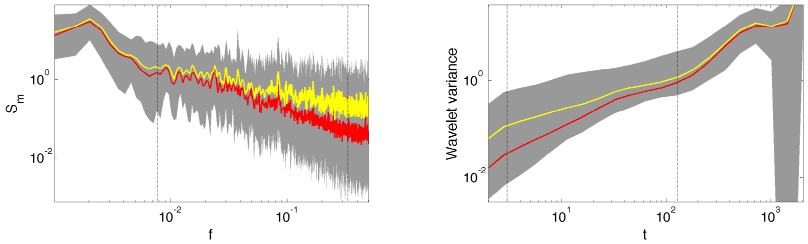

(a) Periodograms (b) Wavelet variances

Figure 9.15: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.00$ in (a) and $H = 0.95$ in (b)) and the yellow line is the mean ($H = 0.84$ in (a) and $H = 0.81$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

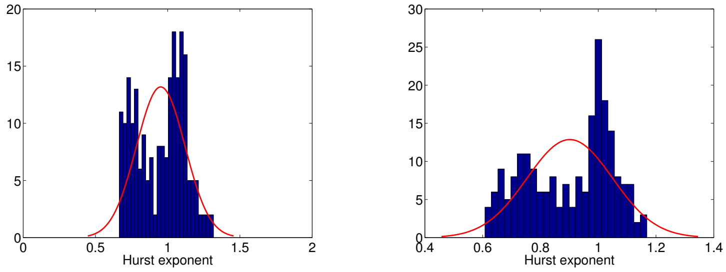

(a) From periodogram (b) From wavelet analysis

Figure 9.16: The distribution of H estimated from the coarse-grained grid boxes in figure 9.14. The periodograms give mean value 0.95, median 0.99 and standard deviation 0.17, and the wavelet variances give mean value 0.90, median 0.95 and standard deviation 0.15.

9.6 55° × 55° grid boxes

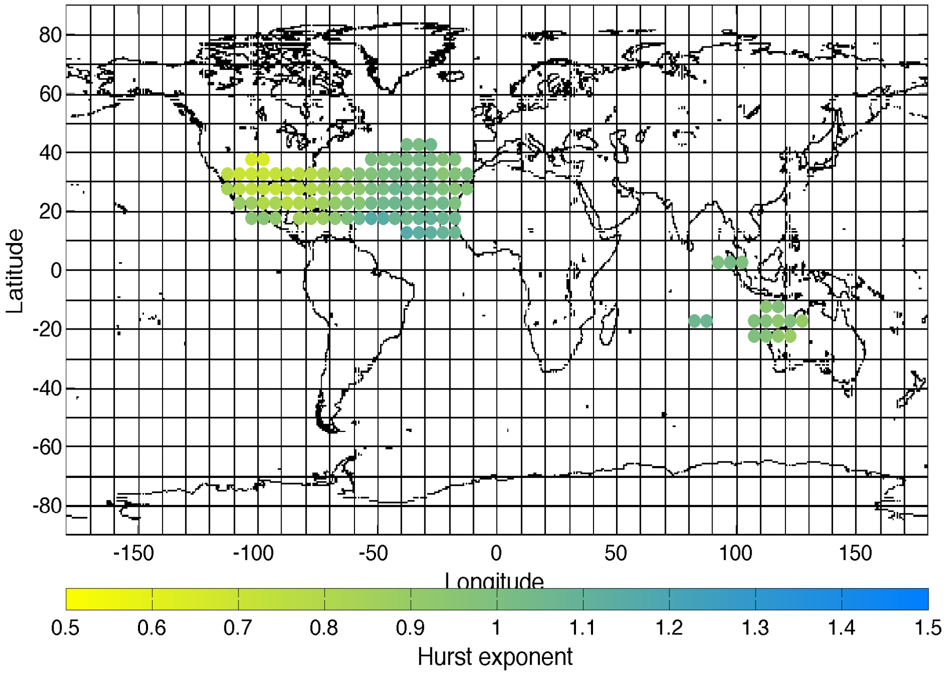

Figure 9.17: The coloured points show the 112 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

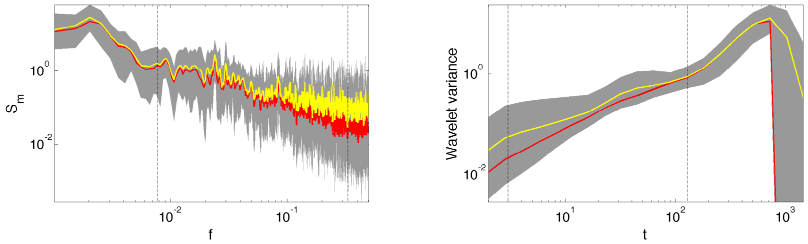

(a) Periodograms (b) Wavelet variances

Figure 9.18: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.04$ in (a) and $H = 0.99$ in (b)) and the yellow line is the mean ($H = 0.93$ in (a) and $H = 0.89$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

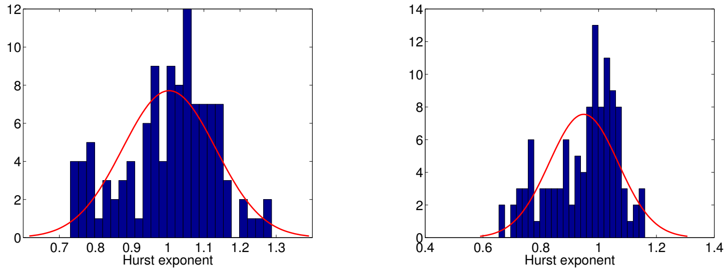

(a) From periodogram (b) From wavelet analysis

Figure 9.19: The distribution of H estimated from the coarse-grained grid boxes in figure 9.17. The periodograms give mean value 1.00, median 1.02 and standard deviation 0.13, and the wavelet variances give mean value 0.95, median 0.98 and standard deviation 0.12.

9.7 $65^\circ \times 65^\circ$ grid boxes

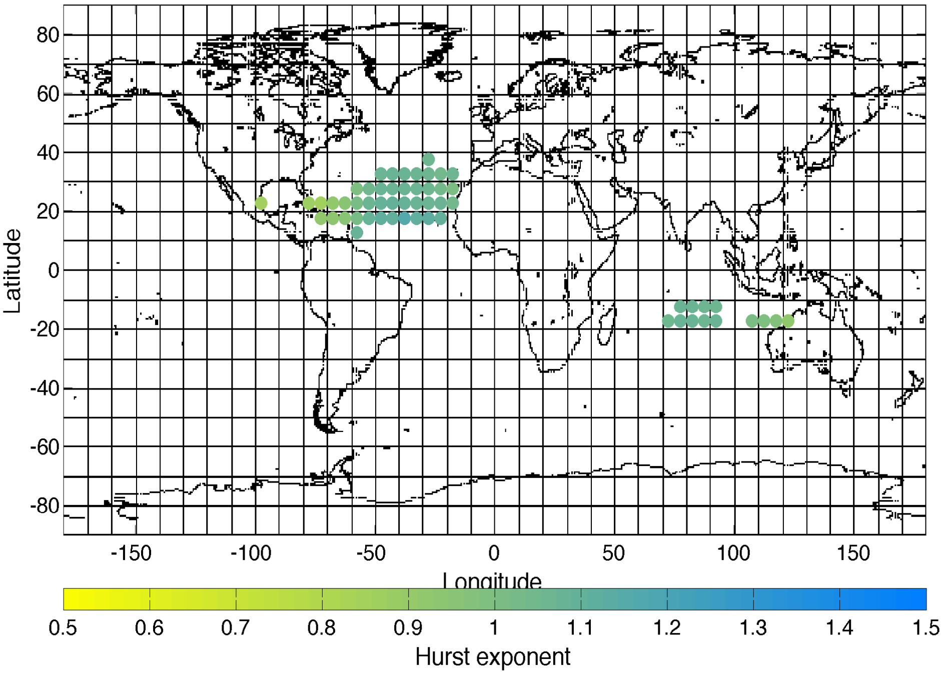

Figure 9.20: The coloured points show the 56 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

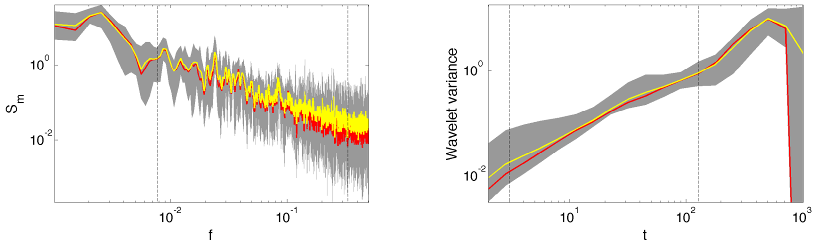

(a) Periodograms (b) Wavelet variances

Figure 9.21: These plots show the shapes of the periodograms and wavelet variances for temperature time series from coarse-grained grid boxes all over the globe. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.14$ in (a) and $H = 1.06$ in (b)) and the yellow line is the mean ($H = 1.09$ in (a) and $H = 1.03$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

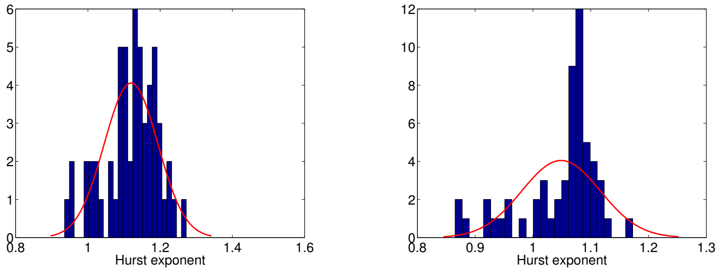

(a) From periodogram (b) From wavelet analysis

Figure 9.22: The distribution of H estimated from the coarse-grained grid boxes in figure 9.20. The periodograms give mean value 1.12 median 1.13 and standard deviation 0.07, and the wavelet variances give mean value 1.05, median 1.07 and standard deviation 0.07.

Chapter 10

Spatial coarse-graining of HadSST3

10.1 5° × 5° grid boxes

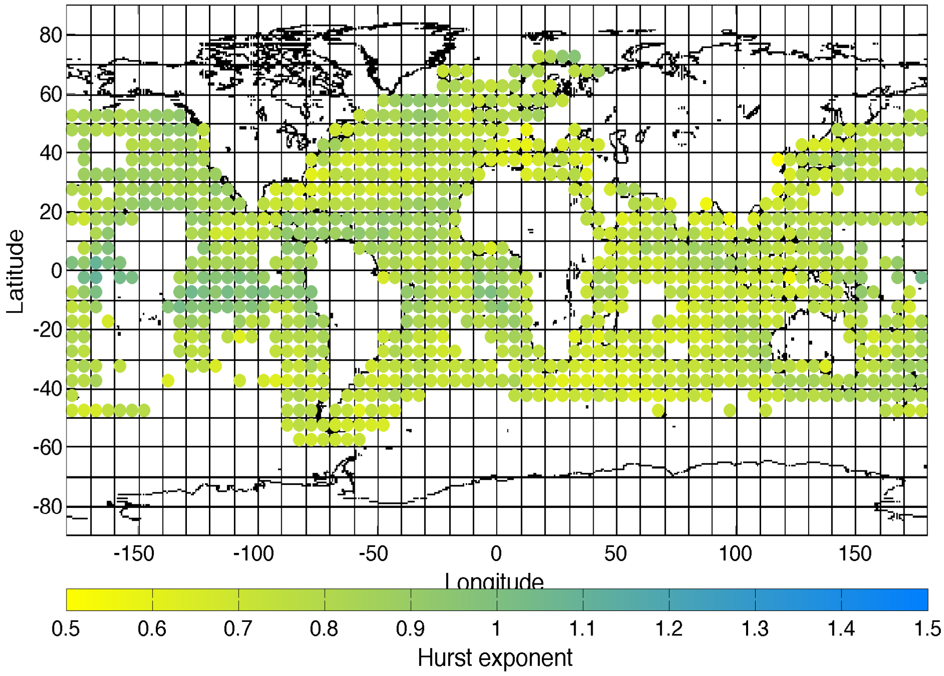

Figure 10.1: The coloured points show the 982 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

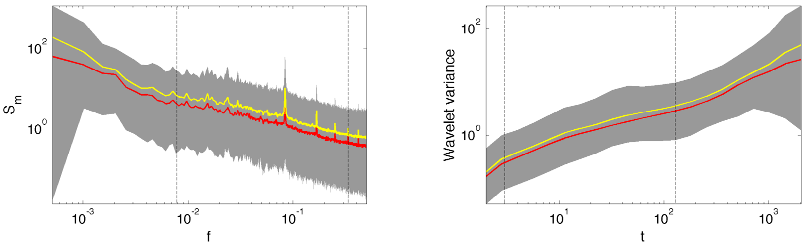

(a) Periodograms (b) Wavelet variances

Figure 10.2: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.82$ in (a) and $H = 0.78$ in (b)) and the yellow line is the mean ($H = 0.82$ in (a) and $H = 0.79$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

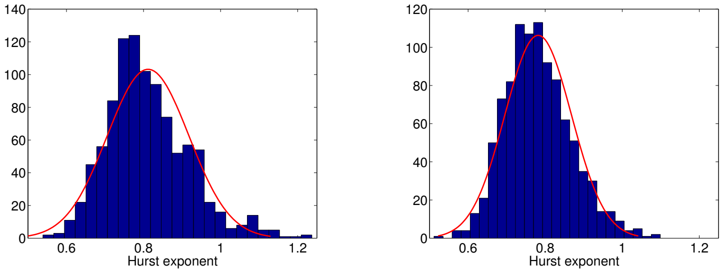

(a) From periodogram (b) From wavelet

Figure 10.3: The distribution of Hurst exponents estimated from the grid boxes in figure 10.1. The periodograms give a mean of 0.81, median 0.80 and standard deviation 0.11. The wavelet variances give a mean of 0.78, median 0.77 and standard deviation 0.09. The red curves are normal fits to the histogram.

10.2 15° × 15° grid boxes

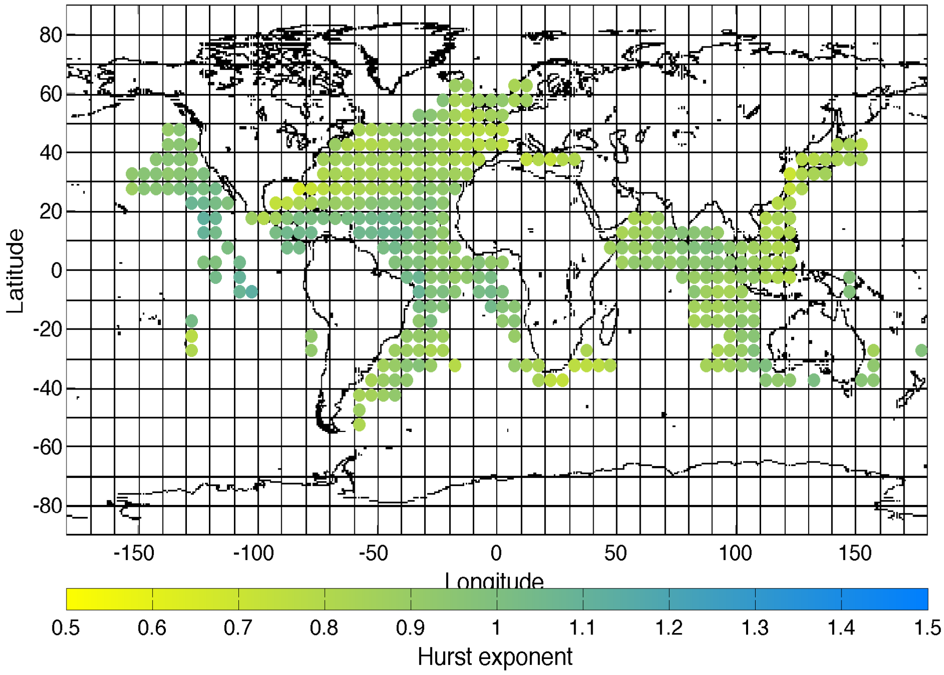

Figure 10.4: The coloured points show the 381 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

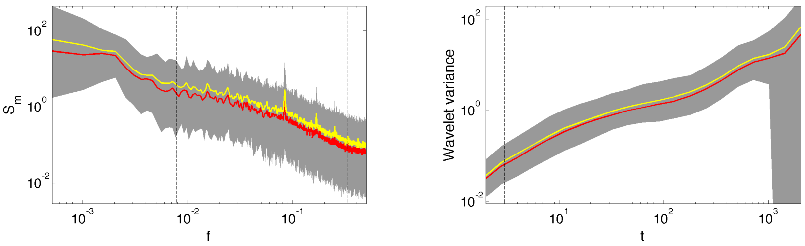

(a) Periodograms (b) Wavelet variances

Figure 10.5: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.96$ in (a) and $H = 0.90$ in (b)) and the yellow line is the mean ($H = 0.96$ in (a) and $H = 0.91$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

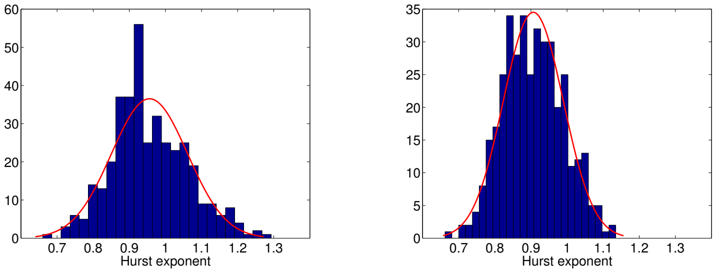

(a) From periodogram (b) From wavelet analysis

Figure 10.6: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 10.4. The periodograms give a mean of 0.96, median 0.94 and standard deviation 0.11. The wavelet variances give a mean of 0.91, median 0.90 and standard deviation 0.08.

10.3 25° × 25° grid boxes

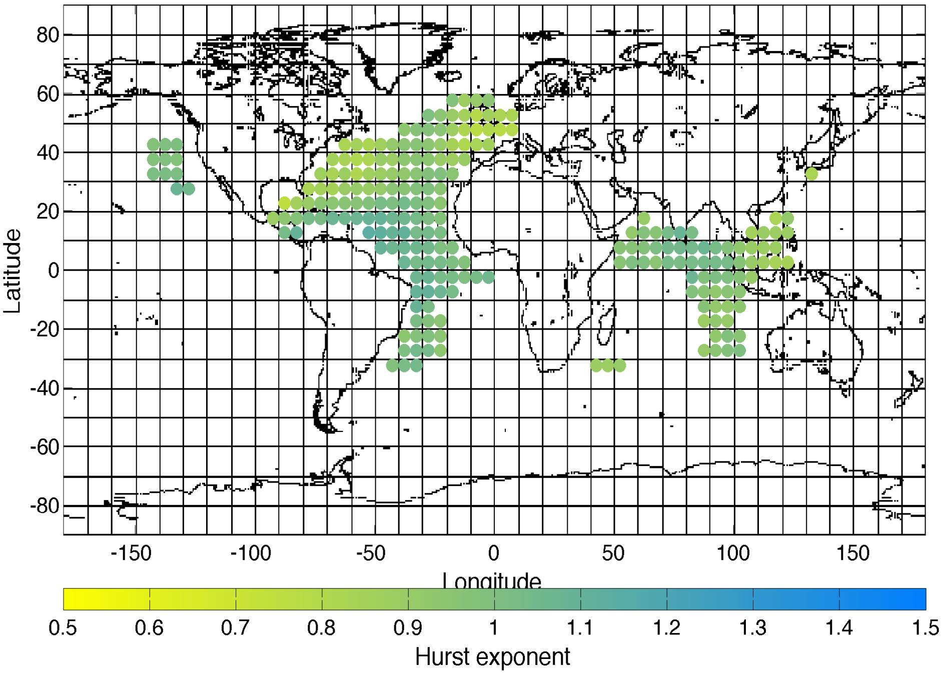

Figure 10.7: The coloured points show the 231 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

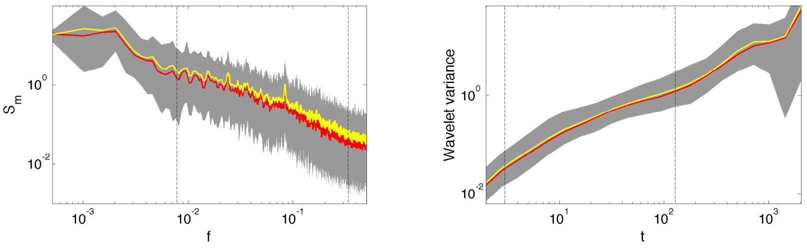

(a) Periodograms (b) Wavelet variances

Figure 10.8: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.03$ in (a) and $H = 0.96$ in (b)) and the yellow line is the mean ($H = 1.00$ in (a) and $H = 0.96$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

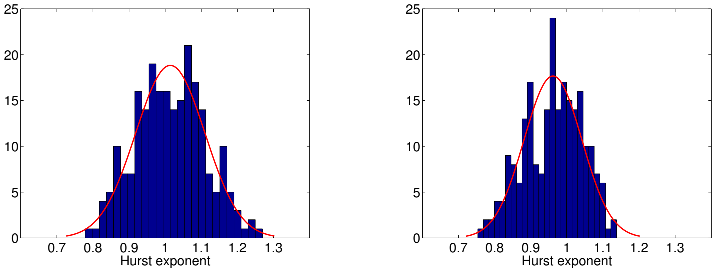

(a) From periodogram (b) From wavelet analysis

Figure 10.9: The distribution ofHurst exponents estimated from the coarse-grained grid boxes in figure 10.7. The periodograms give a mean of 1.01, median 1.01 and standard deviation 0.10. The wavelet variances give a mean of 0.96, median 0.96 and standard deviation 0.08.

10.4 35° × 35° grid boxes

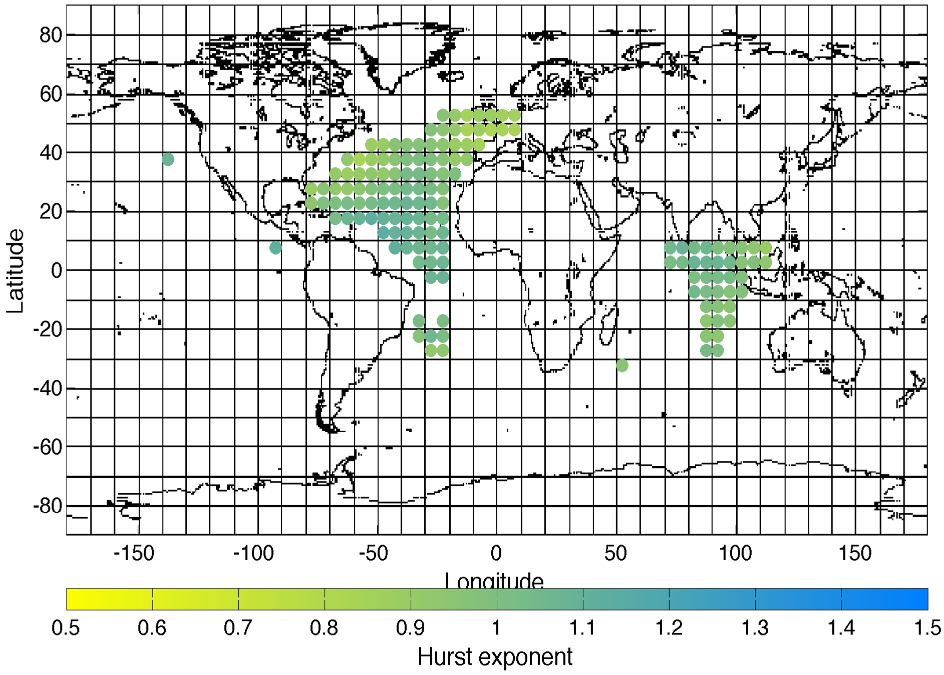

Figure 10.10: The coloured points show the 145 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

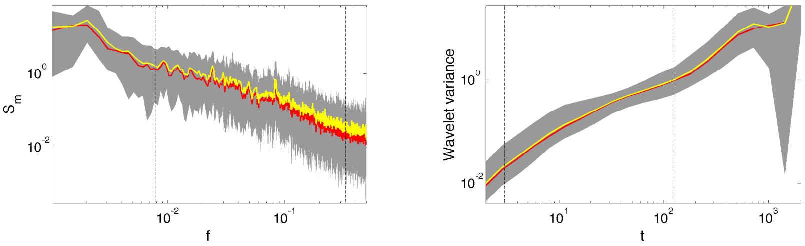

(a) Periodograms (b) Wavelet variances

Figure 10.11: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.07$ in (a) and $H = 1.00$ in (b)) and the yellow line is the mean ($H = 1.04$ in (a) and $H = 0.99$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

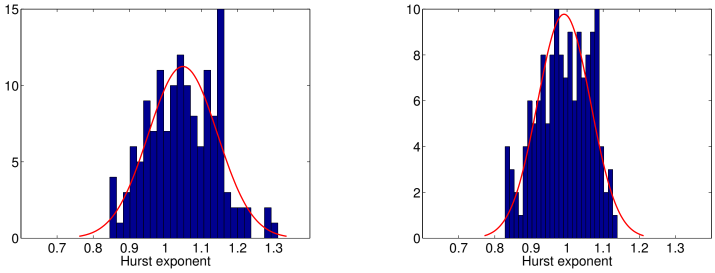

(a) From periodogram (b) From wavelet analysis

Figure 10.12: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 10.10. The periodograms give a mean of 1.05, median 1.05 and standard deviation 0.10. The wavelet variances give a mean of 0.99, median 1.00 and standard deviation 0.07.

10.5 45° × 45° grid boxes

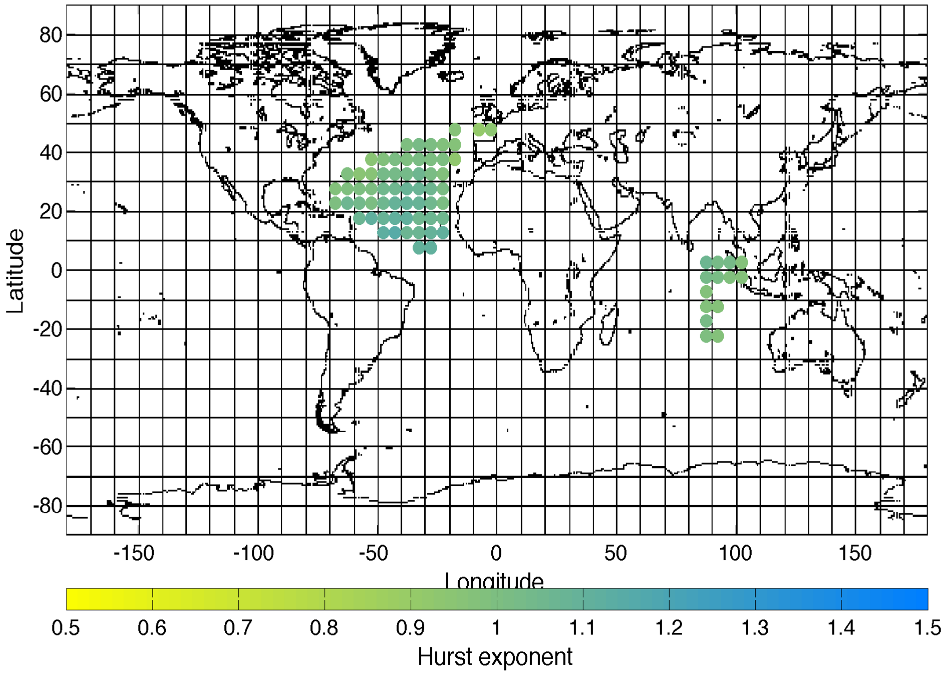

Figure 10.13: The coloured points show the 75 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

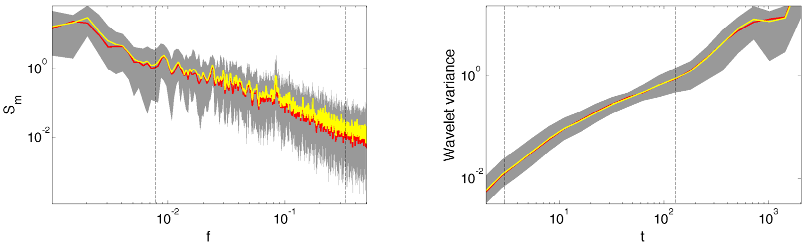

(a) Periodograms

(b) Wavelet variances

Figure 10.14: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.12$ in (a) and $H = 1.04$ in (b)) and the yellow line is the mean ($H = 1.09$ in (a) and $H = 1.03$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

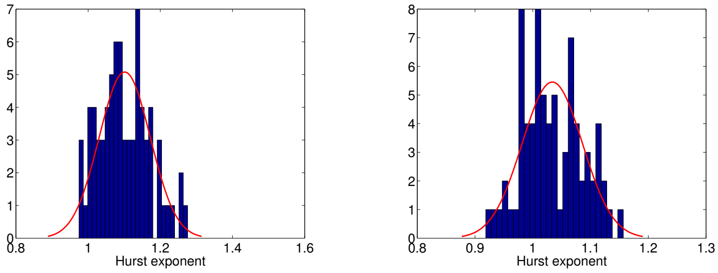

(a) From periodogram

(b) From wavelet analysis

Figure 10.15: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 10.13. The periodograms give a mean of 1.10, median 1.09 and standard deviation 0.07. The wavelet variances give a mean of 1.03, median 1.03 and standard deviation 0.05.

10.6 55° × 55° grid boxes

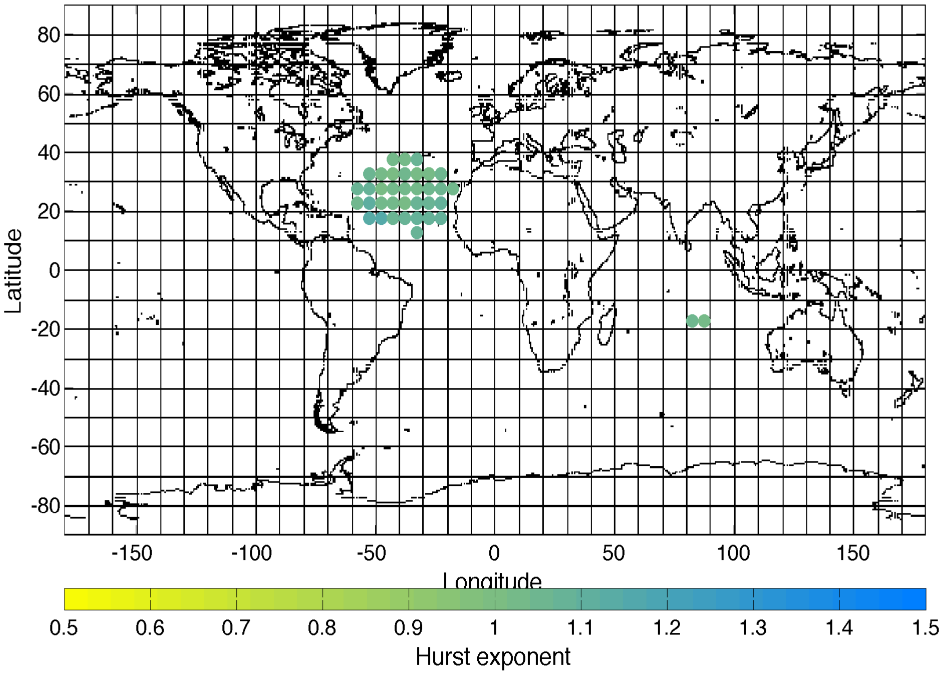

Figure 10.16: The coloured points show the 37 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

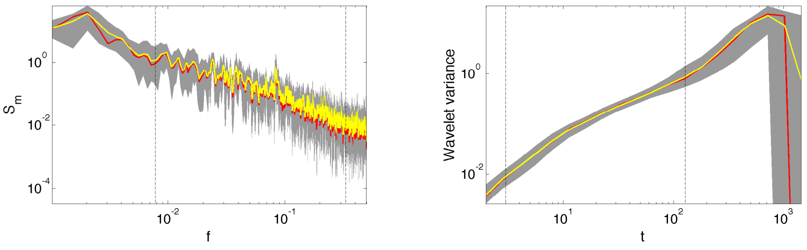

(a) Periodograms (b) Wavelet variances

Figure 10.17: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.15$ in (a) and $H = 1.07$ in (b)) and the yellow line is the mean ($H = 1.13$ in (a) and $H = 1.07$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

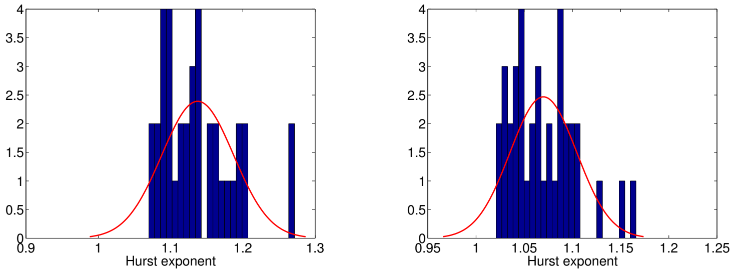

(a) From periodogram (b) From wavelet analysis

Figure 10.18: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 10.16. The periodograms give a mean of 1.14, median 1.13 and standard deviation 0.05. The wavelet variances give a mean of 1.07, median 1.06 and standard deviation 0.03.

10.7 65° × 65° grid boxes

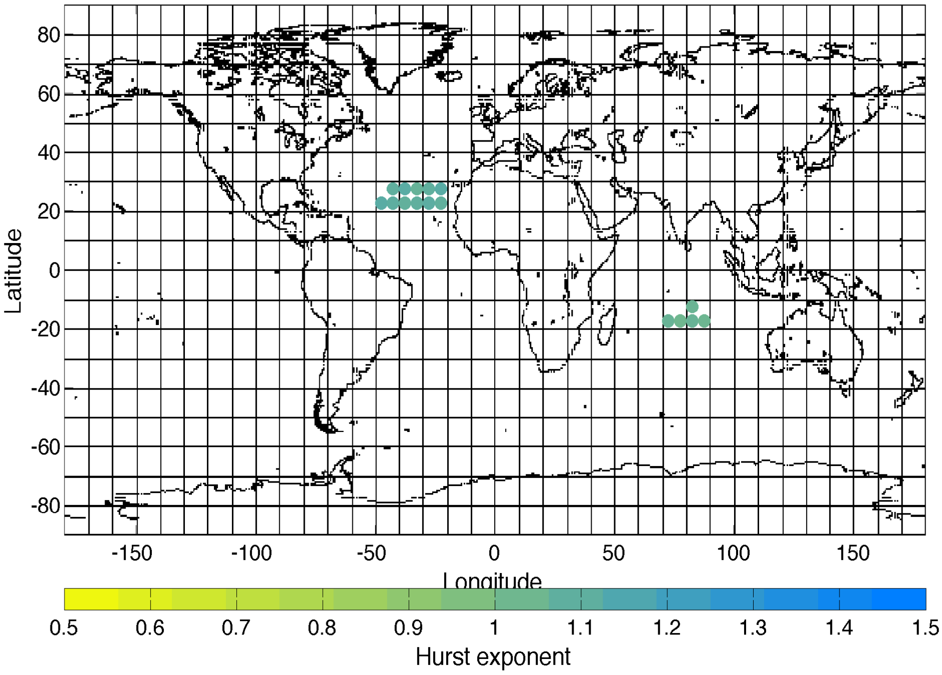

Figure 10.19: The coloured points show the 16 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

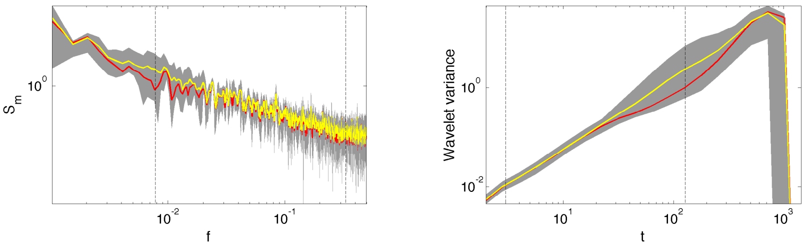

(a) Periodograms (b) Wavelet variances

Figure 10.20: These plots show the shapes of the periodograms and wavelet variances for sea surface temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 1.19$ in (a) and $H = 1.09$ in (b)) and the yellow line is the mean ($H = 1.26$ in (a) and $H = 1.22$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

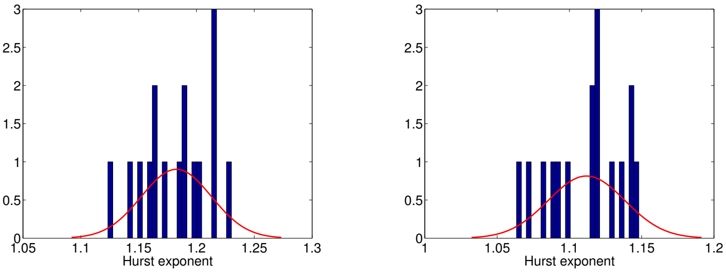

(a) From periodogram (b) From wavelet analysis

Figure 10.21: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 10.19. The periodograms give a mean of 1.18, median 1.19 and standard deviation 0.03. The wavelet variances give a mean of 1.11, median 1.12 and standard deviation 0.03.

Chapter 11

Spatial coarse-graining of Crutem4

11.1 5° × 5° grid boxes

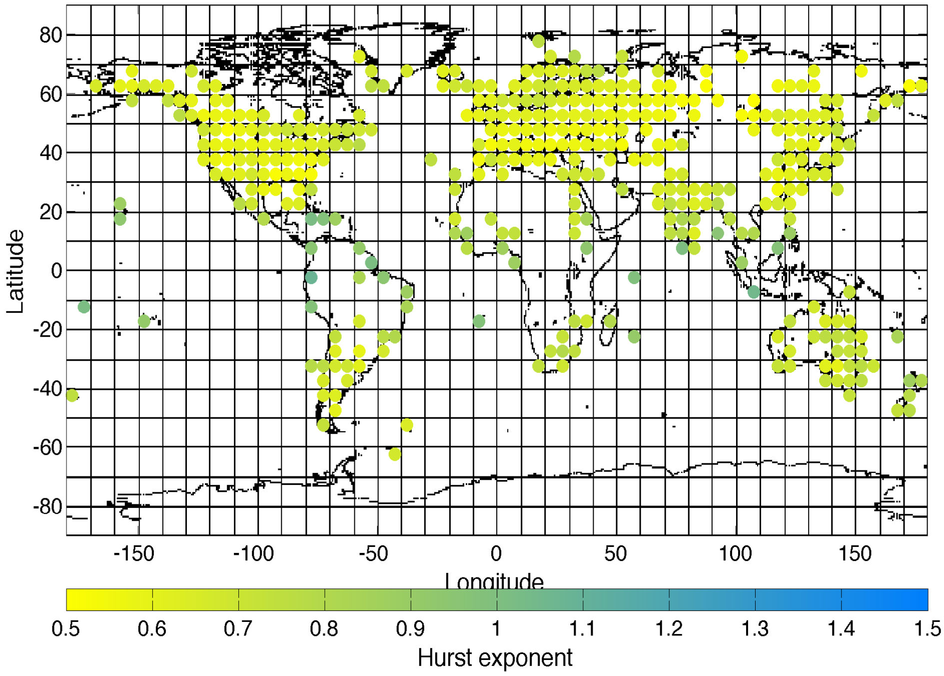

Figure 11.1: The coloured points show the 402 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

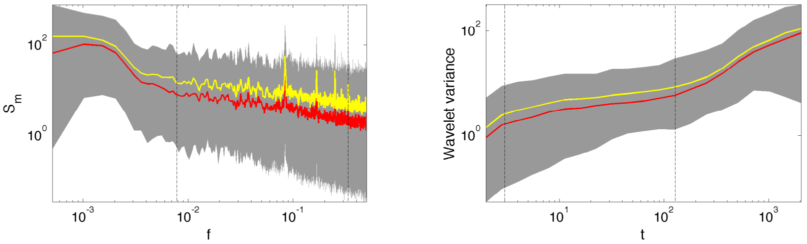

(a) Periodograms

(b) Wavelet variances

Figure 11.2: These plots show the shapes of the periodograms and wavelet variances for land temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.69$ in (a) and $H = 0.65$ in (b)) and the yellow line is the mean ($H = 0.65$ in (a) and $H = 0.63$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

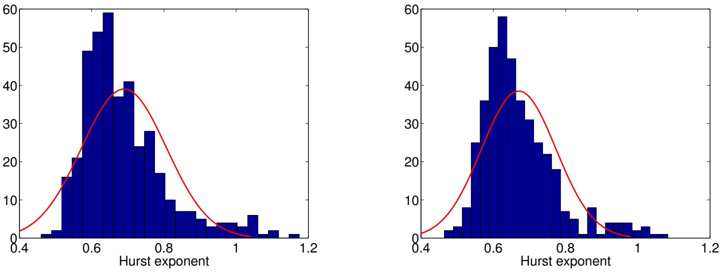

(a) From periodogram

(b) From wavelet analysis

Figure 11.3: The distribution of Hurst exponents estimated from the grid boxes in figure 11.1. We see that this distribution is skewed, and the reason seems to be that there are larger Hurst exponents near the coast. The periodograms give a mean of 0.69, median 0.66 and standard deviation 0.12. The wavelet variances give a mean of 0.67, median 0.65 and standard deviation 0.10.

11.2 15° × 15° grid boxes

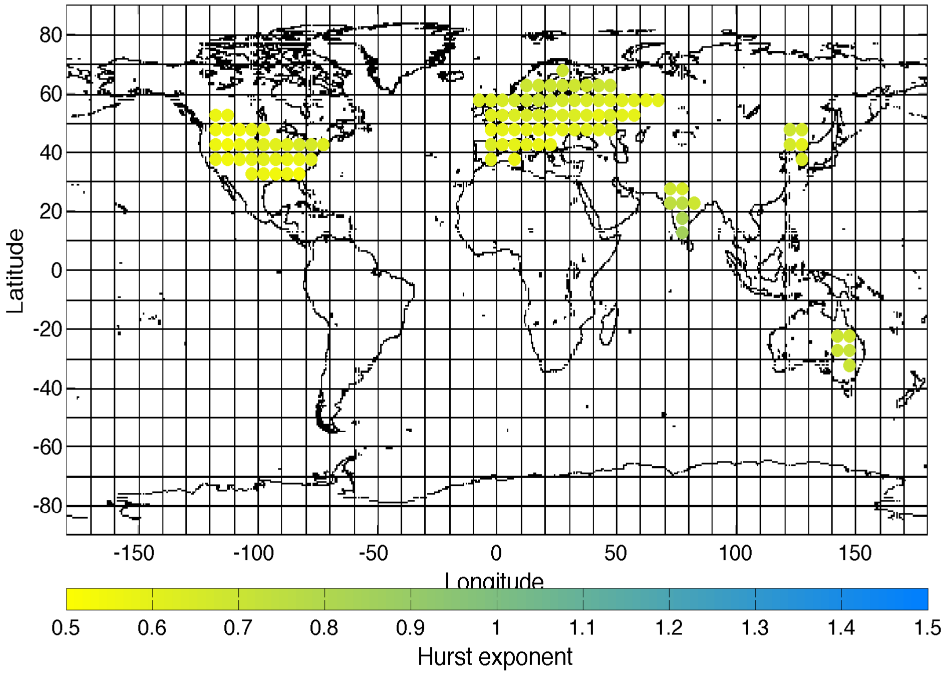

Figure 11.4: The coloured points show the 105 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

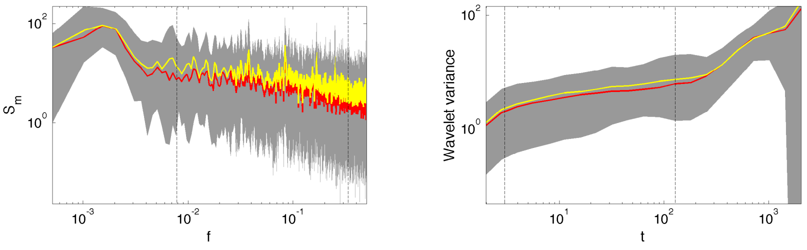

(a) Periodograms (b) Wavelet variances

Figure 11.5: These plots show the shapes of the periodograms and wavelet variances for land temperatures. The grey area is where 95% of the plots are found, the red line is the median (H=0.64 in (a) and H=0.62 in (b)) and the yellow line is the mean (H=0.65 in (a) and H=0.63 in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

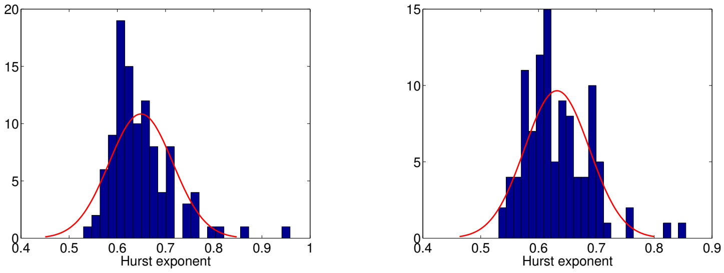

(a) From periodogram (b) From wavelet analysis

Figure 11.6: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 11.4. The periodograms give a mean of 0.65, median 0.63 and standard deviation 0.07. The wavelet variances give a mean of 0.63, median 0.62 and standard deviation 0.06.

11.3 25° × 25° grid boxes

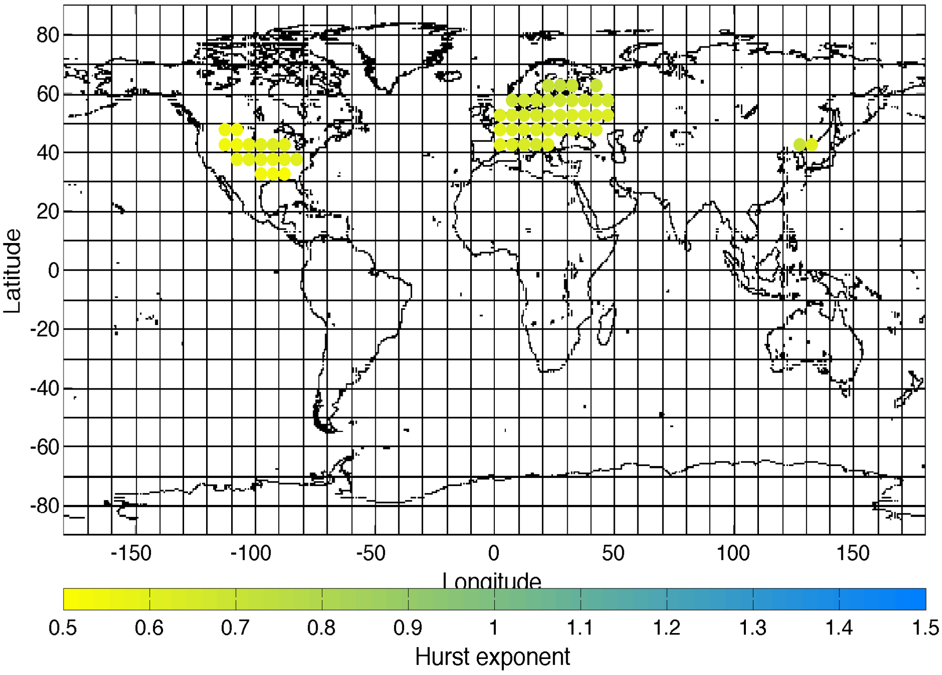

Figure 11.7: The coloured points show the 56 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

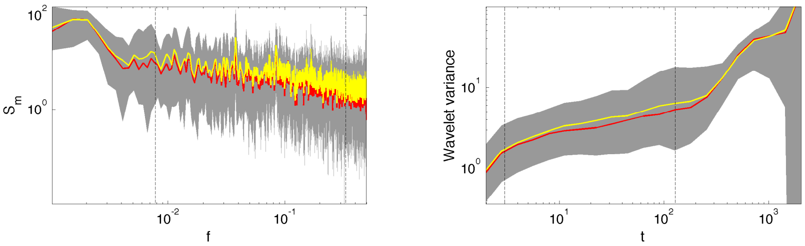

(a) Periodograms (b) Wavelet variances

Figure 11.8: These plots show the shapes of the periodograms and wavelet variances for land temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.66$ in (a) and $H = 0.63$ in (b)) and the yellow line is the mean ($H = 0.67$ in (a) and $H = 0.65$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

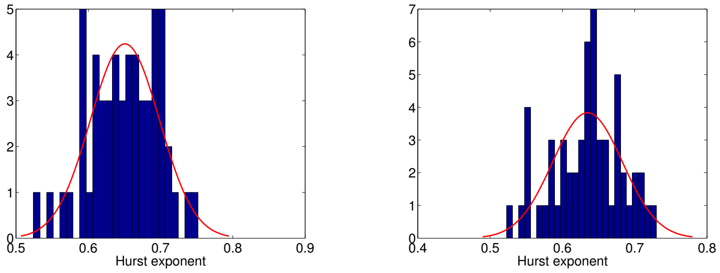

(a) From periodogram (b) From wavelet analysis

Figure 11.9: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 11.7. The periodograms give a mean of 0.65, median 0.65 and standard deviation 0.05. The wavelet variances give a mean of 0.63, median 0.64 and standard deviation 0.05.

11.4 35° × 35° grid boxes

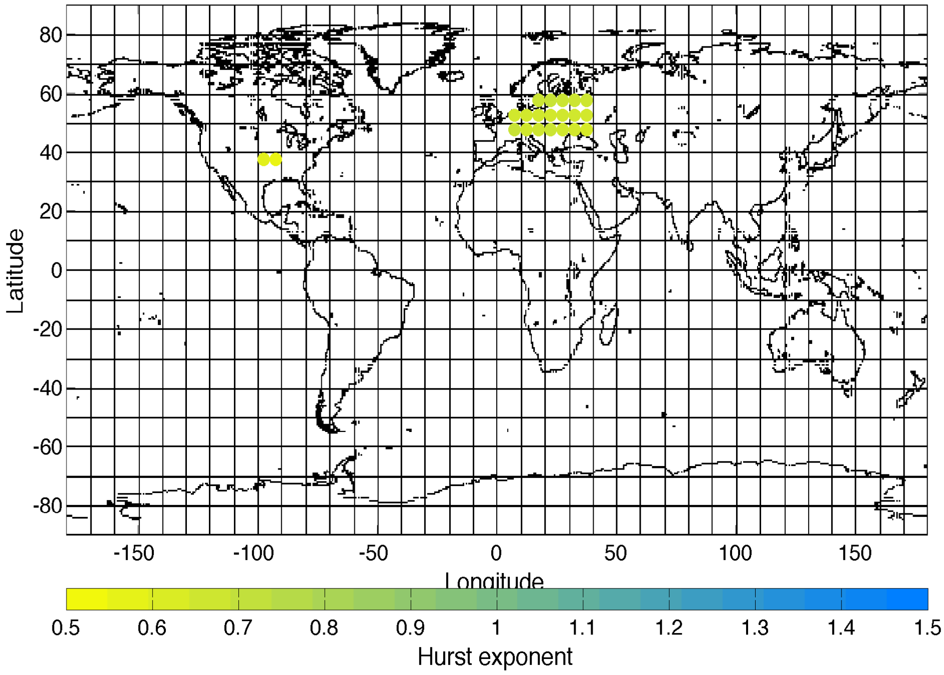

Figure 11.10: The coloured points show the 21 grid boxes with enough measurements to get a good estimate of the Hurst exponent.

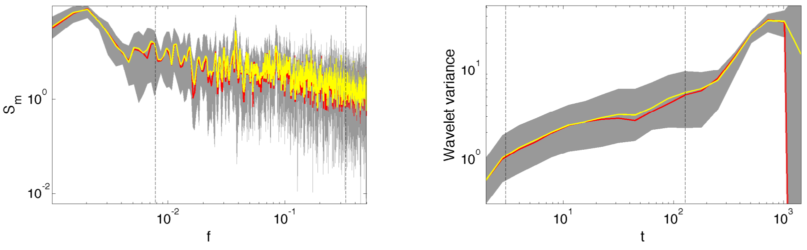

(a) Periodograms (b) Wavelet variances

Figure 11.11: These plots show the shapes of the periodograms and wavelet variances for land temperatures. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.70$ in (a) and $H = 0.67$ in (b)) and the yellow line is the mean ($H = 0.69$ in (a) and $H = 0.68$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these calculations.

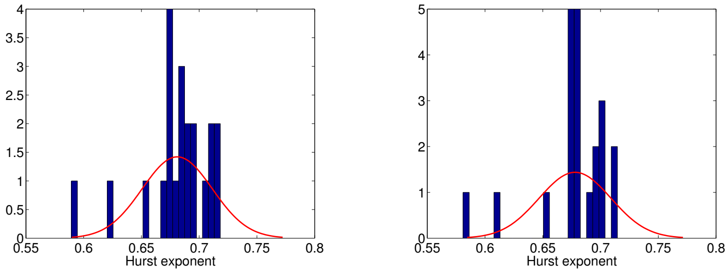

(a) From periodogram (b) From wavelet analysis

Figure 11.12: The distribution of Hurst exponents estimated from the coarse-grained grid boxes in figure 11.10. The periodograms give a mean of 0.68, median 0.68 and standard deviation 0.03. The wavelet variances give a mean of 0.68, median 0.68 and standard deviation 0.03.

Chapter 12

Regional results

In the power spectrum and wavelet variance for global temperature records we observe a small deviation from the power law for time scales of a few years. Since it is well known that El Niño Southern Oscillation (ENSO) is a strong process that varies on these time scales, it is natural to suspect that it can be the reason for the deviations. To test that I have studied the wavelet variances of time series from 5° × 5° grid boxes in different regions, and the results are shown in figure 12.2 and 12.1. From these figures we see that ENSO gives a large bump at these time scales. Not all the other time series have perfect power laws either, but the ENSO region has the most clear deviations. It also appears that part of the Indian Ocean is influenced by ENSO, and that we have some processes stronger than others in the North Atlantic.
The deviations from the power law for global temperatures are largest at time scales around 50 months, so Hurst exponents estimated only up to this scale should be quite large in the regions that gives rise to this bump. The result of this, given in figure 12.3, also supports the result that ENSO is the reason for the bump.

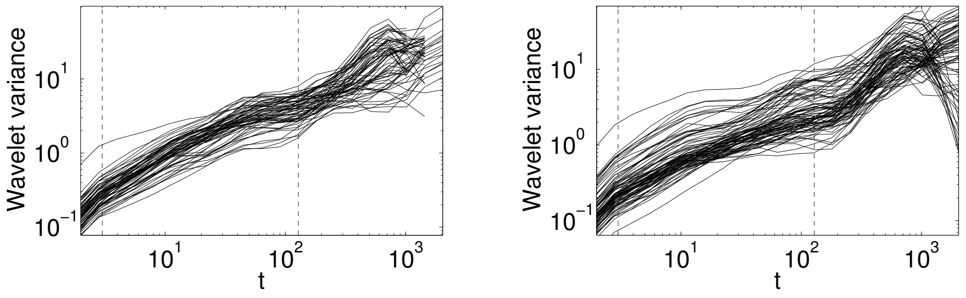

(a) North Pacific (b) North Atlantic

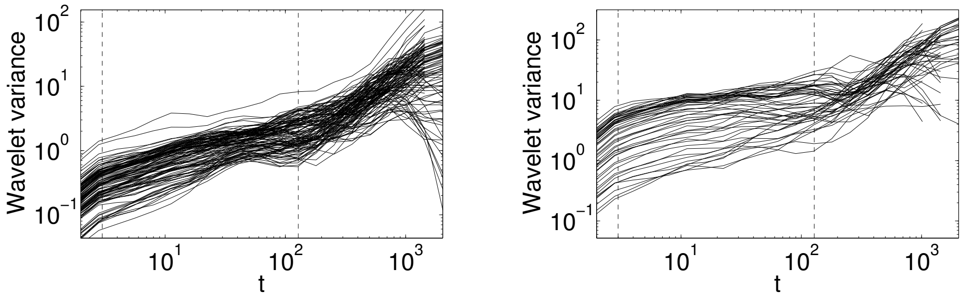

(c) Indian Ocean (d) Some land temperatures

Figure 12.1: Wavelet variances from 5° × 5° grid boxes in the North Pacific ocean on latitudes between 20° and 55° and longitudes between $-180°$ and $-125°$, North Atlantic ocean on latitudes between 10° and 60° and longitudes between $-60°$ and $-15°$, Indian ocean on latitudes between $-50°$ and 10° and longitudes between 50° and 110°, and some land temperatures. The region between the dashed lines is where I fit the straight line when estimating the Hurst exponent, and the time scale t is measured in months.

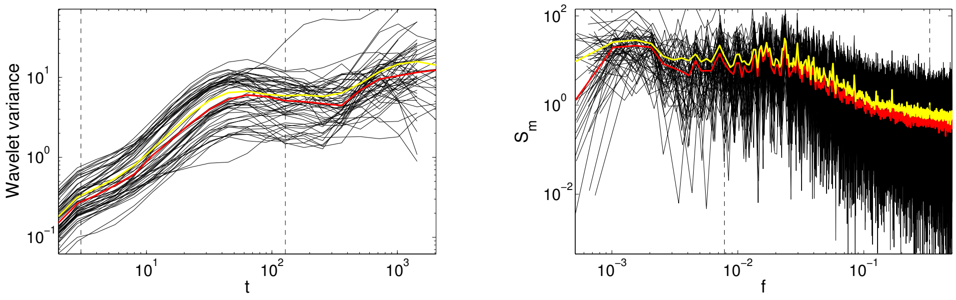

(a) Wavelet variances

(b) Periodograms

Figure 12.2: Wavelet variances and periodograms from 5° × 5° grid boxes in the Pacific ocean on latitudes between $-10°$ and $10°$ and longitudes between $-180°$ and $-75°$.

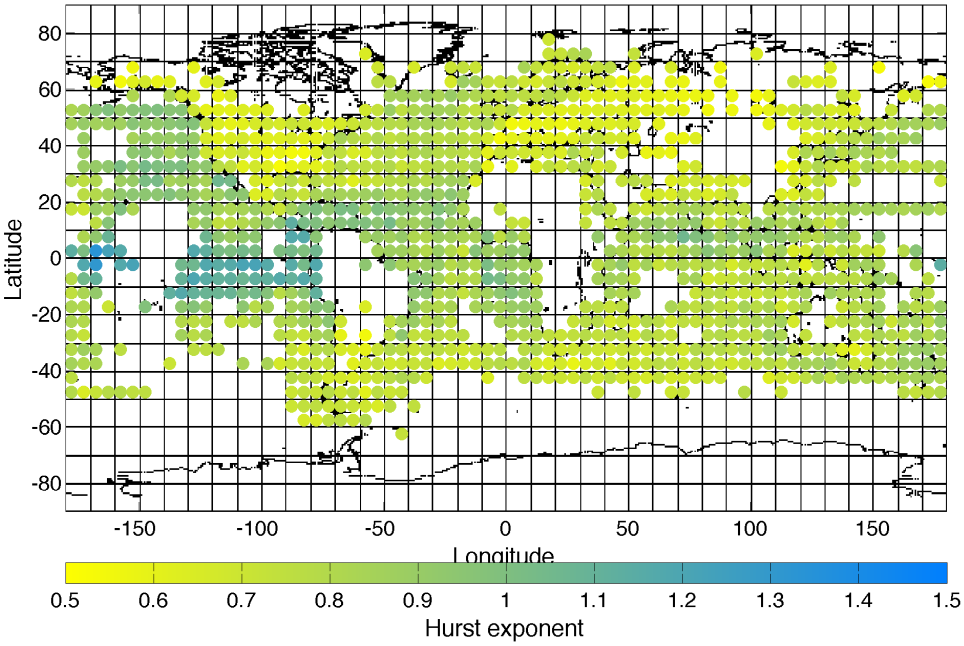

Figure 12.3: Hurst exponents estimated on scales from 3 to 50 months. In figure 12.2 we see that the bump in the wavelet variance has a maximum at around 50 months, so the high Hurst exponents in this figure should indicate where we have these bumps.

12.1 Averaging of E-OBS land temperatures

For the E-OBS temperatures I started out with daily resolved temperatures in a 0.5°×0.5° grid. The Hurst exponents estimated using periodogram are shown in figure 12.4 for all grid boxes that contained at least 4000 contiguous data points after running the time series through my "repair function". In this figure we notice that the highest Hurst exponents can be found at the lower latitudes and near the coast, and the lowest Hurst exponents are found near mountains. One possible explanation of this can be that the mountains are obstacles to horizontal diffusion of heat, such that we can get very cold or warm periods resulting in large temperature variations at the shortest time scales analyzed. Near the coast and equator we know that the annual variations are smaller. At lower latitudes we also find the Hadley cell which gives us trade winds near the surface. That could possibly give us even less variations on the short time scales by smoothing out the spatial temperature variations.

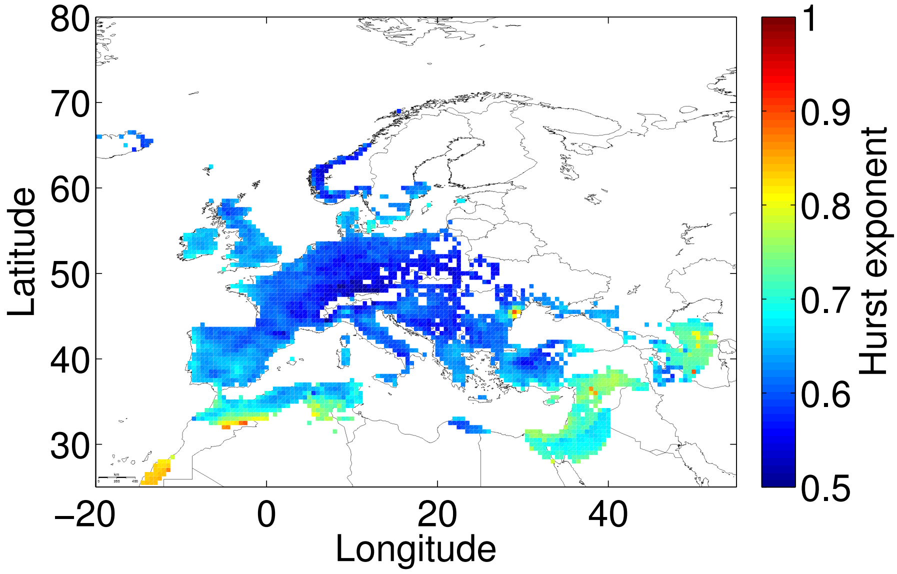

Figure 12.4: The coloured points show the Hurst exponent estimated by periodograms for all time series that had a length of 4000 days or more. The time scales used are in the range 3 weeks to 10 years.

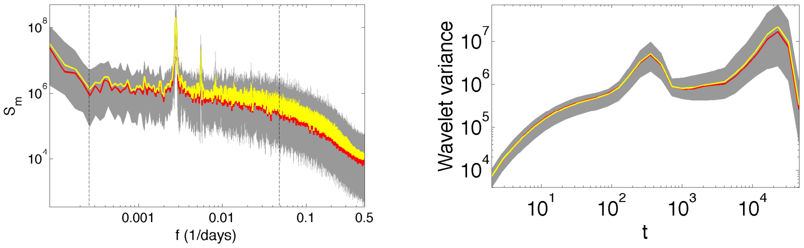

(a) Periodograms (b) Wavelet variances

Figure 12.5: These plots show the shapes of the periodograms and wavelet variances. The grey area is where 95% of the plots are found, the red line is the median ($H = 0.63$ in (a) and $H = 0.62$ in (b)) and the yellow line is the mean ($H = 0.63$ in (a) and $H = 0.63$ in (b)). Since the time series do not have the same lengths, all the periodograms and wavelet variances are interpolated before they are used in these plots.

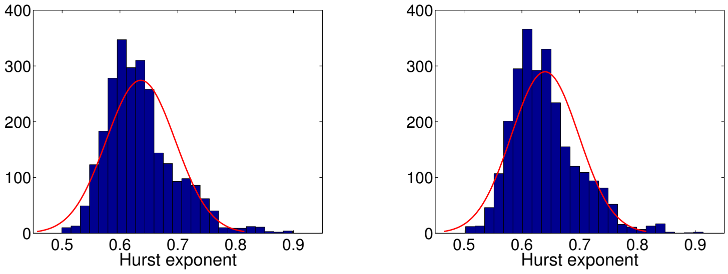

(a) From periodogram (b) From wavelet analysis

Figure 12.6: The distribution of Hurst exponents estimated from the grid boxes in figure 12.4. The periodograms give a mean of 0.64, median 0.63 and standard deviation 0.06. The wavelet variances give the same.

In figure 12.5 we observe that the annual variations are seen very clearly in both the periodogram and the wavelet variance, even though deseasonalization have been attempted by subtraction of the climatology. For the periodogram it is a sharp peak, while in the wavelet variance a smaller bump. In both of the methods these regions are not included when fitting a slope. As seen in figure 12.6 the results for wavelet variance and periodogram are very similar, but I have chosen to put a little more trust

	Wavelet			Periodogram		
	Mean	Median	s	Mean	Median	s
0.5° × 0.5°	0.60	0.60	0.06	0.61	0.61	0.06
1.5° × 1.5°	0.59	0.59	0.05	0.60	0.60	0.06
2.5° × 2.5°	0.60	0.60	0.06	0.61	0.61	0.06
3.5° × 3.5°	0.58	0.59	0.05	0.59	0.60	0.06
4.5° × 4.5°	0.58	0.59	0.05	0.59	0.59	0.06
5.5° × 5.5°	0.57	0.57	0.06	0.58	0.58	0.07
6.5° × 6.5°	0.55	0.55	0.06	0.55	0.55	0.08

Table 12.1: This table shows the means, medians and standard deviations for the Hurst exponents estimated for different degrees of spatial coarse-graining of the E-OBS land temperatures.

the periodogram in these analyzes because of the large bump in the wavelet variance.

When I started to average the E-OBS temperatures in space I found that the daily resolved data set contained more datapoints than the memory of my computer could handle, so I chose to average the temperatures in time to get monthly mean values for the temperature first. The number of measurements I then got in each grid box is shown in figure 12.7. The Hurst exponents obtained for these new time series are shown in figure 12.8 for the three smallest grid box sizes used. For the coarse-grained grid boxes some stripes of missing data or similar Hurst exponents appeared, but nothing seemed to be wrong about my code. So I suspect that the reason for this is either the short length of the time series used, or the way the datapoints are interpolated in constructing the gridded temperatures. Despite of this I tried to find the mean and median Hurst exponents for the averaged time series, which are given in table 8.1. The results did not turn out to be very meaningful since the Hurst exponents actually decreased with increasing coarse graining.

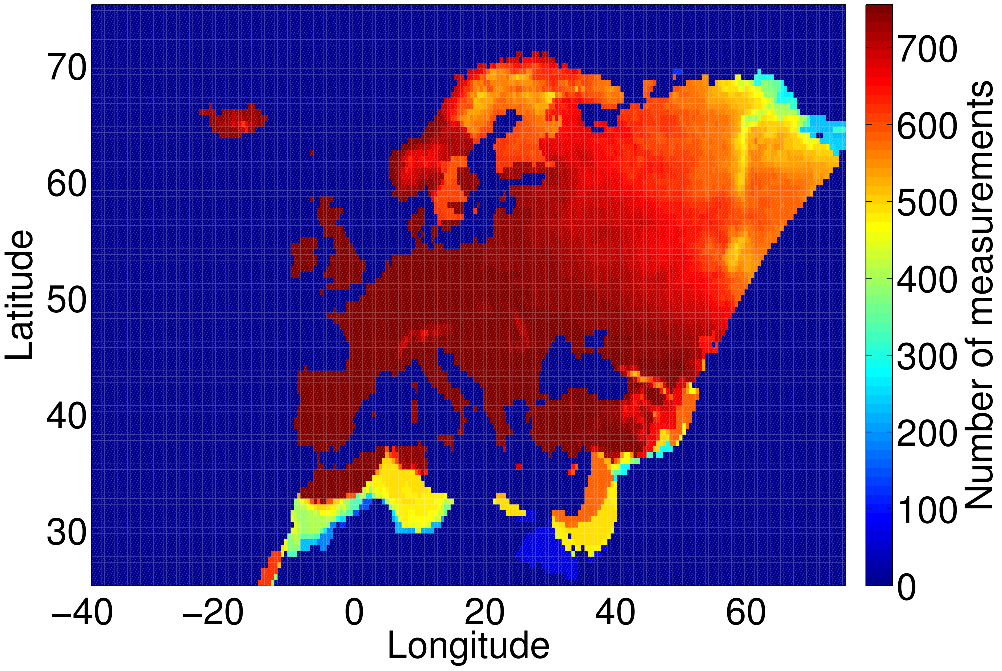

Figure 12.7: The coloured points show the number of monthly data points in each elementary grid box since the recordings started in 1950. A monthly measurement is here defined to exist if there is at least one measurement done in that month.

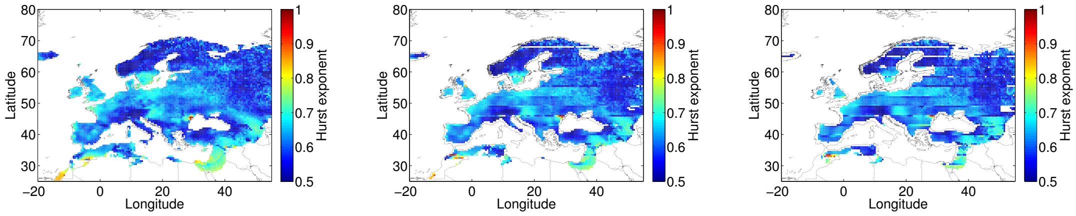

(a) 0.5° × 0.5°

(b) 1.5° × 1.5°

(c) 2.5° × 2.5°

Figure 12.8: Hurst exponents for monthly resolved time series of the E-OBS dataset.

Chapter 13

Conclusions

The mean values and standard deviations from all histograms from chapter 9, 10 and 11 are summed up in figure 13.1. Also included in this figure are all relevant Hurst exponents from table 8.1. For the sea surface temperature we seem to have a significant increase in the Hurst exponent when we coarse-grain temperature time series by averaging over larger areas. For the land temperatures the increase is very small, and the statistics giving us the estimated values are not that good. The land data are influenced by the sea especially for the smallest grid boxes, such that the mean Hurst exponent is larger than it would be for only interior land temperatures. I have tried to eliminate the effect from the coastal areas by plotting the most probable Hurst exponent instead of the mean. Still the increase in the Hurst exponent with increasing area is very small, and perhaps not significant.

We can also see that the periodogram in most cases gives higher Hurst exponents than the wavelet variance. The time series are not detrended before the analysis is done, so this may have something to do with the detrending abilities of the wavelet. The mother wavelet is a second derivative of a gaussian, so it can remove a linear trend, while the periodogram does not remove any trends. Another difference between the methods is that the wavelet variance gives a smoother curve than the periodogram. Not only does it look smoother, but it also tends to smooth out peaks in the spectrum, which can be seen very clearly in figure 12.5.

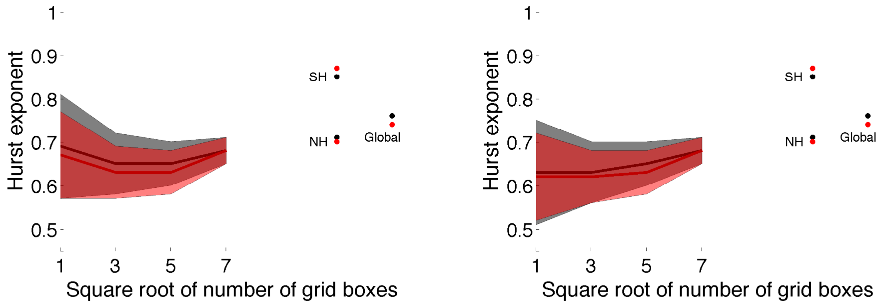

(a) Land (b) Land

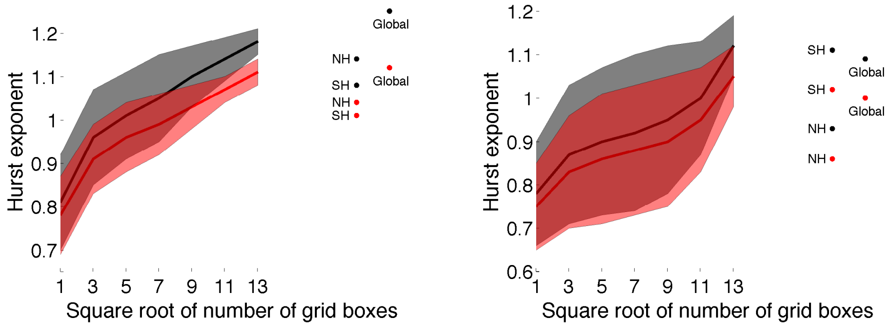

(c) Sea (d) Combined land and sea

Figure 13.1: These figures show the Hurst exponents we get as a function of the size of the coarse grained grid cell, where the black curves and dots show the result when we use the periodogram as an estimator, and the red curves and dots when we use wavelet variance. The shaded areas are one standard deviation on each side of the mean, and by one grid point I mean a grid box of size $5^\circ \times 5^\circ$. The high Hurst exponents for the smallest scale land temperatures in (a) is a result of the influence of the sea on coastal land temperatures. In (b) the black and the red lines are estimates for the most probable Hurst exponent instead of the mean, which will be a better estimate for the true Hurst exponent for interior land areas.

The results in figure 13.1 are consistent with what we could expect from the analyses done in section 7.3, which show that spatial correlations for sea surface temperatures have much more influence on spatially averaged temperatures than land temperatures. Hence we can expect the increase in Hurst exponent with increasing coarse graining to be higher for sea than for land. The reason that we find higher Hurst exponents for some of the coarse-grained grid boxes than for global temperature is probably because the largest coarse-grained boxes are averages over only low-latitude temperatures, while for the global temperature high-latitude temperatures are also used in the average.

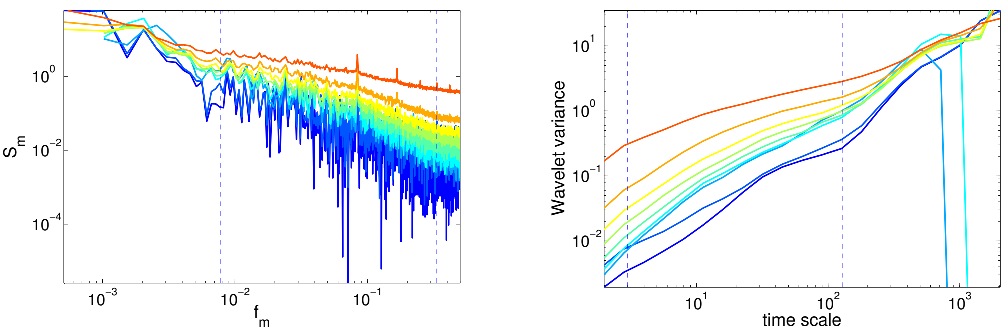

(a) Periodograms (b) Wavelet variances

Figure 13.2: The upper curves in these plots show the median for 5° × 5° grid boxes, the second are for 15° × 15° grid boxes, and so on up to 65° × 65° grid boxes, all for sea-surface temperatures. The second lowest curves are the means of the periodograms/wavelet variances of the northern and southern hemispheres and the lowest curves are the periodogram/wavelet variance for global sea surface temperature.

Figure 13.2 shows the median periodograms and wavelet variances for increasing grid box size of sea surface temperatures. The global curves have less power than the regional and local since the fluctuations become less when we do the averaging. This is just as expected, even if the processes were independent and had the same Hurst exponent as seen in equation (7.3). Some of these curves are not perfect power laws, since we have combined time series resulting from many different processes, where the local time series in some cases have large deviations from power laws.

In section 7.3 I also found that the reason for the nonstationarity observed in global temperature probably is mainly related to cross-covariances on distances longer than 14000 km, which corresponds to around 35% of the earth's circumference. On the time scales from one to ten years the Hurst exponent is expected to be highest if we use only the time scales up to around 6 years according to the results in figure 7.19 (b). This is also what we find in the global sea surface temperature because of the bump we have on time scales of a few years.

Summed up in a few points, my conclusions are that:

- Hurst exponents for interior land temperatures are smaller than Hurst exponents for sea surface temperatures. This result is also consistent with previous analyses.

- Hurst exponents do in general increase with increasing spatial scale. This increase is larger for sea surface temperatures than for land temperatures.

- Both local and global variations in temperatures as a function of time scale follow quite good power laws, although there are some deviations. The largest deviations we find in the regions where ENSO is dominant. This process is so strong that it also gives us a small deviation for global temperature.

- The trend in the temperature is affecting the temperature variations mainly on the largest time scales, and is seen more clearly for land temperatures than for sea surface temperatures in these studies.

13.1 Further work

The spatial averaging should be performed on more datasets, including both observational data and data from global circulation models. In that way the validity of my results can be tested. The cross-covariances should also be studied in more detail, since there exists so many large deviations from the mean that I have studied. More specifically, one should at least study a belt around the equator to find out more about the correlations that ENSO gives rise to, as this seem to be one of the most dominant process in the climate system. The results of this analysis can then be compared to earlier analysis of climate networks. That may give us a lot of information about spatial correlation structures.

Appendix A

Area of grid box

Infinitesimal area at latitude θ:

$$\mathrm{d}A = \mathrm{d}l \cdot \mathrm{d}b = r^2 \mathrm{d}\theta \; \mathrm{d}\varphi \; \cos(\theta)$$

where dl is an infinitesimal length element in longitudinal direction and db is an infinitesimal length element in latitudinal direction. Area of a grid box with $\Delta\theta \cdot \Delta\varphi$ degrees:

$$\Delta A = \int_{\varphi_0}^{\varphi_0+\Delta\varphi} \int_{\theta_0}^{\theta_0+\Delta\theta} r^2 \cos(\theta) \; \mathrm{d}\theta \; \mathrm{d}\varphi = \Delta\varphi \; r^2 [\sin(\theta_0 + \Delta\theta) - \sin(\theta_0)] \tag{A.1}$$

For a $5^\circ \times 5^\circ$ grid box we have that $\Delta\varphi = \Delta\theta = \frac{\pi}{36}$

The radius of the earth varies between 6353 km and 6384 km. The mean radius can be calculated in many different ways. I choose to use the mean value $r = 6367$ km.

Appendix B

Bibliography

[*Beran* (1994)] Jan Beran (1994), Statistics for Long-Memory Processes, *Chapman & Hall.* ISBN 0-412-04901-5.

[*Brohan et al.* (2006)] Brohan, P., J. J. Kennedy, I. Harris, S. F. B. Tett, and P. D. Jones (2006), Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850, *J. Geophys. Res.*, 111, D12106, doi:10.1029/2005JD006548.

[*Burnham et al.* (2010)] Kenneth P. Burnham, David R. Anderson, and Kathryn P. Huyvaert (2010), AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons, *Behav Ecol Sociobiol (2011) 65:23–35*, DOI 10.1007/s00265-010-1029-6

[*Gade* (2010)] Kenneth Gade (2010), A Non-singular Horizontal Position Representation, *THE JOURNAL OF NAVIGATION*, 63, 395–417. f The Royal Institute of Navigation doi:10.1017/S0373463309990415

[*Haylock et al.* (2008)] Haylock, M.R., N. Hofstra, A.M.G. Klein Tank, E.J. Klok, P.D. Jones and M. New. 2008: A European daily high-resolution gridded dataset of surface temperature and precipitation. *J. Geophys. Res (Atmospheres)*, 113, D20119, doi:10.1029/2008JD10201

[*Hu et al.* (2001)] Kun Hu, Plamen Ch. Ivanov, Zhi Chen, Pedro Carpena, and H. Eugene Stanley (2001), *The American Physical Society*, DOI:10.1103/PhysRevE.64.011114

[*Huybers and Curry* (2006)] Peter Huybers and William Curry, Links between annual, Milankovitch and continuum temperature variability (2006), *Nature* 441, 329-332 (18 May 2006), doi:10.1038/nature04745

[*Jokinen et al.* (2000)] Heikki Jokinen, Jaakko Ollila, Olli Aumala, On windowing effects in estimating averaged periodograms of noisy signals, *Measurement*, Volume 28, Issue 3, October 2000, Pages 197-207, ISSN 0263-2241, 10.1016/S0263-2241(00)00013-0.

[*Jones et al.* (2012)] Jones, P. D., D. H. Lister, T. J. Osborn, C. Harpham, M. Salmon, and C. P. Morice (2012), Hemispheric and large-scale land-surface air temperature variations: An extensive revision and an update to 2010, *J. Geophys. Res.*, 117, D05127, doi:10.1029/2011JD017139.

[*Kantelhardt et al.* (2001)] Jan W. Kantelhardt, Eva Koscielny-Bunde, Henio H. A. Rego, Shlomo Havlin and Armin Bunde, Detecting Long-range Correlations with Detrended Fluctuation Analysis (2001), *Physica A, 295, 441-454*, DOI: 10.1016/S0378-4371(01)00144-3

[*Kennedy et al.* (2011a)] Kennedy, J. J., N. A. Rayner, R. O. Smith, M. Saunby, and D. E. Parker (2011a), Reassessing biases and other uncertainties in sea-surface temperature observations measured in situ since 1850: 1. Measurement and sampling errors, *J. Geophys. Res.*, 116, D14103, doi:10.1029/2010JD015218.

[*Kennedy et al.* (2011b)] Kennedy, J. J., N. A. Rayner, R. O. Smith, M. Saunby, and D. E. Parker (2011b), Reassessing biases and other uncertainties in sea-surface temperature observations measured in situ since 1850: 2. Biases and homogenization, *J. Geophys. Res.*, 116, D14104, doi:10.1029/2010JD015220.

[*Lovejoy and Schertzer* (2013)] Shaun Lovejoy and Daniel Schertzer (2013), The Weather and the Climate, emergent Laws and Multifractal Cascades, *Cambridge University Press*

[*Malamud and Turcotte* (1999)] Bruce D. Malamud and Donald L. Turcotte, Self-Affine Time Series: I. Generation and Analyses, *Advances in geophysics* pp. 1-90

[*Mandelbrot and Van Ness* (1968)] Benoit B. Mandelbrot and John W. Van Ness, Fractional Brownian Motions, Fractional Noises and Applications, *SIAM Rev.*, 10(4), 422–437

[*McLeod et al.* (2007)] A. Ian McLeod, Hao Yu, Zinovi L. Krougly, Algorithms for Linear Time Series Analysis: With R Package, *Journal of statistical software*, December 2007, Volume 23, Issue 5.

[*Morice et al.* (2012)] Morice, C. P., J. J. Kennedy, N. A. Rayner, and P. D. Jones (2012), Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set, *J. Geophys. Res.*, 117, D08101, doi:10.1029/2011JD017187.

[*Pelletier and Turcotte* (1999)] Jon D. Pelletier and Donald L. Turcotte, Self-Affine Time Series: II. Generation and Analyses, *Advances in geophysics* pp. 91-166

[*Rypdal et al.* (2013)] Kristoffer Rypdal, Lene Østvand and Martin Rypdal (2013), Long-range memory in Earth's surface temperature on time scales from months to centuries, *Journal of Geophysical Research: Atmospheres*, DOI: 10.1002/jgrd.50399

[*Qian* (2003)] Hong Qian, 2003: Fractional Brownian Motion and Fractional Gaussian Noise, G. Rangarajan, M. Ding (Eds.): LNP 621, pp. 22–33.

[*Torrence and Compo* (1998)] Torrence, Christopher, Gilbert P. Compo, 1998: A Practical Guide to Wavelet Analysis. *Bull. Amer. Meteor. Soc.*, 79, 61–78, doi: http://dx.doi.org/10.1175/1520-0477(1998)079<0061:APGTWA>2.0.CO;2

[*Vyushin et al.* (2012)] Vyushin, D. I., P. J. Kushner, and F. Zwiers (2012), Modeling and understanding persistence of climate variability, *J. Geophys. Res.*, 117, D21106, doi:10.1029/2012JD018240.

UNIVERSITY OF TROMSØ