

Fractional Gaussian noise: Prior specification and model comparison

Sigrunn Holbek Sørbye¹ and Håvard Rue²

¹Department of Mathematics and Statistics, UiT The Arctic University of Norway, 9037 Tromsø, Norway. E-mail: sigrunn.sorbye@uit.no

²CEMSE Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia. E-mail: haavard.rue@kaust.edu.sa

Abstract

Fractional Gaussian noise (fGn) is a stationary stochastic process used to model anti-persistent or persistent dependency structures in observed time series. Properties of the autocovariance function of fGn are characterised by the Hurst exponent (H), which in Bayesian contexts typically has been assigned a uniform prior on the unit interval. This paper argues why a uniform prior is unreasonable and introduces the use of a penalised complexity (PC) prior for H . The PC prior is computed to penalise divergence from the special case of white noise, and is invariant to reparameterisations. An immediate advantage is that the exact same prior can be used for the autocorrelation coefficient ϕ of a first-order autoregressive process AR(1), as this model also reflects a flexible version of white noise. Within the general setting of latent Gaussian models, this allows us to compare an fGn model component with AR(1) using Bayes factors, avoiding confounding effects of prior choices for the two hyperparameters H and ϕ . Among others, this is useful in climate regression models where inference for underlying linear or smooth trends depends heavily on the assumed noise model.

Key words: Autoregressive process; Bayes factor; long-range dependence; PC prior; R-INLA

1 Introduction

Many real time series exhibit statistical properties which can be modelled by self-similar stochastic processes. Such processes have probability distributions which are invariant to changes of scale in time (or space), and are particularly useful to model long-range dependency structures. Long-memory is observed within a wide range of fields, among others economics, climatology, geophysics and network engineering (Beran et al., 2013). Of specific use is fractional Brownian motion (fBm) (Kolmogorov, 1940; Mandelbrot and Ness, 1968), which is the only self-similar continuous-time Gaussian process with stationary increments. Its discrete-time increment process, referred to as fractional Gaussian noise (fGn), has autocovariance function characterised by the self-similarity parameter $H \in (0, 1)$. This parameter is often referred to as the Hurst exponent, known to quantify the Hurst phenomenon (Hurst, 1951).

Several methods to estimate H have been proposed in the literature, among others heuristic approaches like the rescaled range method (Hurst, 1951), detrended fluctuation analysis (Peng et al., 1994) and the rescaled variance method (Giraitis et al., 2003). Alternative approaches include use of wavelets (McCoy and Walden, 1996; Abry and Veitch, 1998), and maximum likelihood and Whittle estimation, see Beran et al. (2013) for a comprehensive overview. Using a Bayesian framework,

H has typically been assigned a uniform prior (Benhmehdi et al., 2011; Makarava et al., 2011; Makarava and Holschneider, 2012). This is computationally simple and in lack of prior knowledge, a non-informative prior might seem like a good choice. However, one major drawback is that the uniform prior is not invariant to reparameterisations, and as seen in Section 2 this can give surprising results when the prior is transformed.

This paper introduces the use of penalised complexity (PC) priors (Simpson et al., 2017), in estimating the Hurst exponent H . We also use a PC prior for the precision of the process. PC priors represent a new principle-based approach to compute priors for a wide range of hyperparameters in hierarchical models, where a given model component can be seen as a flexible version of a simple base model. The simplest base model in the case of fGn is to have no random effect, corresponding to infinite precision. For a fixed precision, the fGn process represents a flexible version of white noise ($H = 0.5$). The PC prior is computed to penalise model component complexity. This is achieved by assigning a prior to a measure of distance from the flexible model to the base model, which is then transformed to give a prior for the hyperparameter of interest. The informativeness of the prior is adjusted by an intuitive and interpretable user-defined scaling criterion.

The framework of PC priors has several beneficial properties, including robustness, invariance to reparameterisations, and it also provides meaningful priors with a clear interpretation (Simpson et al., 2017; Riebler et al., 2016). This paper utilises the invariance property to compare fGn with a first-order autoregressive process, AR(1), using Bayes factor. Similarly to fGn, the AR(1) process represents a flexible version of uncorrelated white noise, where the first-lag autocorrelation coefficient $\phi = 0$. This implies that the PC priors for the hyperparameters H and ϕ can be chosen as transformations of the exact same prior assigned to the distance measure. This eliminates confounding effects of prior choices in the calculation of Bayes factor which is very important as the Bayes factor is known to be sensitive to prior choices (Kass and Raftery, 1995). A comparison of the AR and fGn models is for example relevant in analysing climatic time series (Løvstetten and Rypdal, 2016), identifying short versus long-range dependency structures in temperature series at different temporal and/or spatial scales. However, the given ideas could also be used in comparing other models as long as these represent flexible versions of the same base model.

The structure of this paper is as follows. Section 2 gives a brief review of the principles underlying PC priors and utilises these to derive the PC prior for H . Specifically, the model complexity of fGn is seen to be non-symmetric around $H = 0.5$ which implies that the prior for H should also be non-symmetric. Section 2 also illustrates how Beta priors on H would shift the mode of the prior for distance away from the natural simple base model. In section 3, we focus on model comparison of fGn and the first-order AR process using Bayes factor. We perform a simulation study and also apply the given ideas to check for fGn structure in global land and sea-surface temperatures, respectively. Both the estimation and the comparison of models can be computed within the general framework of latent Gaussian models, using the methodology of integrated nested Laplace approximations (INLA) (Rue et al., 2009). This implies that inference for fGn models is made easily available and that fGn can be combined with other model components in constructing an additive predictor, for example including covariates and non-linear or oscillatory random effects. Concluding remarks are given in section 4, while Appendix 5 describes the implementation of fGn using the R-interface `R-INLA`, combined with generic functions.

2 Penalised complexity priors for the parameters of fractional Gaussian noise

Prior selection for hyperparameters is a difficult issue in Bayesian statistics and often subject to rather ad-hoc choices. A computationally simple choice is to adopt flat, non-informative priors. In

the case of fGn, the Hurst parameter H has typically been assigned a uniform prior, argued for in terms of having no knowledge about the parameter (Benhmehdi et al., 2011; Makarava et al., 2011; Makarava and Holschneider, 2012). Also, it is suggested to use the Jeffrey’s prior for the marginal standard deviation of the model, i.e. $\pi(\tau^{-1/2}) \sim \tau^{1/2}$. In this section, we derive the penalised complexity prior (Simpson et al., 2017) for H and suggest to use a PC prior also for the precision τ . We also describe drawbacks in using a uniform, or more generally a Beta prior, for H .

2.1 Penalising divergence from a base model

Define fractional Gaussian noise as a zero-mean multinormal vector $\mathbf{x}' = (x_1, \dots, x_n) \sim N(\mathbf{0}, \tau^{-1}\Sigma)$. The correlation matrix Σ is Toeplitz with first-row elements

$$\gamma(k) = \frac{1}{2}(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}), \quad k = 0, \dots, n-1,$$

where $H \in (0, 1)$ is referred to as the Hurst exponent, while τ denotes the precision parameter. For large k , the autocorrelation function of fGn has a power-law decay, $\gamma(k) \approx H(2H-1)|k|^{2(H-1)}$. We notice that fGn reduces to uncorrelated white noise when $H = 0.5$. When $H > 0.5$, the process has positive correlation reflecting a persistent autocorrelation structure. Similarly, the autocorrelation is negative when $H < 0.5$, and the resulting process is then referred to as being anti-persistent.

The specific structure of the autocorrelation function of fGn implies that the process can be seen as a flexible version of white noise, where H represents a flexibility parameter. The framework of penalised complexity priors (Simpson et al., 2017) takes advantage of this inherit nested structure, which is seen in many model components where hyperparameters can be used to define flexible versions of a simple base model structure. A key idea is to assign a prior to the divergence from the flexible model to the base model and transform this to give a prior for the hyperparameter of interest. Specifically, the prior is calculated based on four principles, stated in Simpson et al. (2017). These are as follows:

1. The PC prior is calculated to support Occam’s razor, emphasising that model simplicity should be preferred over model complexity. Specifically, a model component is seen as a flexible version of a simple base model specification. In the fGn case, let $f_1(H) = \pi(\mathbf{x} | \Sigma)$ represent the flexible version of a white noise base model $f_0 = \pi(\mathbf{x} | \mathbf{I})$, where \mathbf{I} is the identity matrix. The PC prior for H is then designed to give shrinkage to the fixed value $H = 0.5$.
2. The second principle in deriving a PC prior implies penalisation of model complexity using a measure of distance from the flexible version of a model component to its base model. This is achieved by making use of the Kullback-Leibler divergence, which in the zero-mean Gaussian case simplifies to $\text{KLD}(f_1(H) \| f_0) = -\frac{1}{2} \ln |\Sigma|$. For the fGn process, the determinant will be of order n and we can define a distance measure from f_1 to f_0 as

$$d(H) = \sqrt{\frac{1}{n} 2\text{KLD}(f_1(H) \| f_0)} = \sqrt{-\frac{1}{n} \ln |\Sigma|} = \begin{cases} d_1(H), & 0 \leq H \leq 0.5 \\ d_2(H), & 0.5 \leq H \leq 1. \end{cases} \quad (1)$$

which is invariant to n .

3. The third principle used to derive PC priors assumes constant rate penalisation. This implies that the relative change in the prior is constant, independent of the actual value of the distance measure. Consequently, the prior assigned to the distance d will be exponential. In the fGn case, the PC prior is calculated separately for the two distances $d_1(H)$ and $d_2(H)$,

$$\pi(d_i(H)) = \frac{1}{2} \lambda \exp(-\lambda d_i(H)), \quad i = 1, 2,$$

where $H \in [0, 1]$ and $\lambda > 0$ denotes the rate parameter. The prior for H is then obtained by an ordinary change of variable transformation,

$$\pi_i(H) = \pi(d_i(H)) \left| \frac{d_i(H)}{dH} \right|, \quad i = 1, 2, \quad (2)$$

where the derivative is found numerically. In practice, truncation due to the bounded interval of H is avoided by calculating the prior for the logit transformation $\rho = \text{logit}(H)$, which also has a variance-stabilising effect.

4. The last principle in deriving PC priors states that the rate λ , controlling shrinkage of the prior, should be governed by a user-defined and interpretable scaling criterion. This needs to be chosen for each specific model component of interest. In the given case, a natural and easily-implemented choice is to define a tail probability, like

$$P(U < H < 1) = \alpha \quad (3)$$

where α is a specified small probability. Alternatively, the given probability statement could for example define the median of the prior. In any of these cases, the corresponding rate parameter is

$$\lambda = \frac{-\ln(2\alpha)}{d(U)}.$$

2.2 Implications using Beta versus PC priors

The calculated distance function in (1) is non-symmetric around $H = 0.5$ (Figure 1), clearly illustrating that the properties of fractional Gaussian noise depend on the value of H . Specifically, the distance measure for volatility or anti-persistent behaviour ($0 < H < 0.5$) is seen to decrease more slowly than in the case of having positive correlation and long memory properties ($0.5 < H < 1$). Also, notice that the given distance measure increases rapidly as $H \rightarrow 0$ and $H \rightarrow 1$, reflecting the major change in the properties of the fGn process as it approaches its limiting non-stationary cases.




Figure 1: The distance measure $d(H)$, from a general fractional Gaussian process, seen as a function of H and to white noise base model where $H = 0.5$.

By transforming the exponential prior on the given distance measure, the resulting PC prior for H will preserve the properties of the distance function. Specifically, the PC prior will always have

its mode at the white noise base model ($H = 0.5$), independent of the parameter specifications for U and α in (3). Also, it preserves the non-symmetry of the distance function and it goes to infinity when $H \rightarrow 0$ and $H \rightarrow 1$. The PC prior for H is illustrated in Figure 2 (a), using three different tail probabilities in (3). Specifically $P(0.9 < H < 1)$ is set equal to 0.10, 0.15 and 0.20, in which the last two cases express an increased prior belief in persistent versus anti-persistent behaviour. Alternatively, if we knew in advance that a given process has long memory properties, we could consider to define the prior only for the interval $H \in (0.5, 1)$.

Using the given tail probabilities, the corresponding values for the rate parameter λ of the exponential distribution are approximately equal to 1.70, 1.27 and 0.97. Figure 2 (b) illustrates the corresponding Beta priors, where the shape parameters are chosen to give the same probability statements as for the PC priors. Naturally, this can be obtained choosing the shape parameters of the Beta distribution in numerous way. Here, we just make a specific choice in which the first shape parameter of the Beta distribution is set equal to 1.0, 1.8 and 2.5, respectively. Specifically, the first case gives the uniform distribution which equals the Beta(1,1) case.

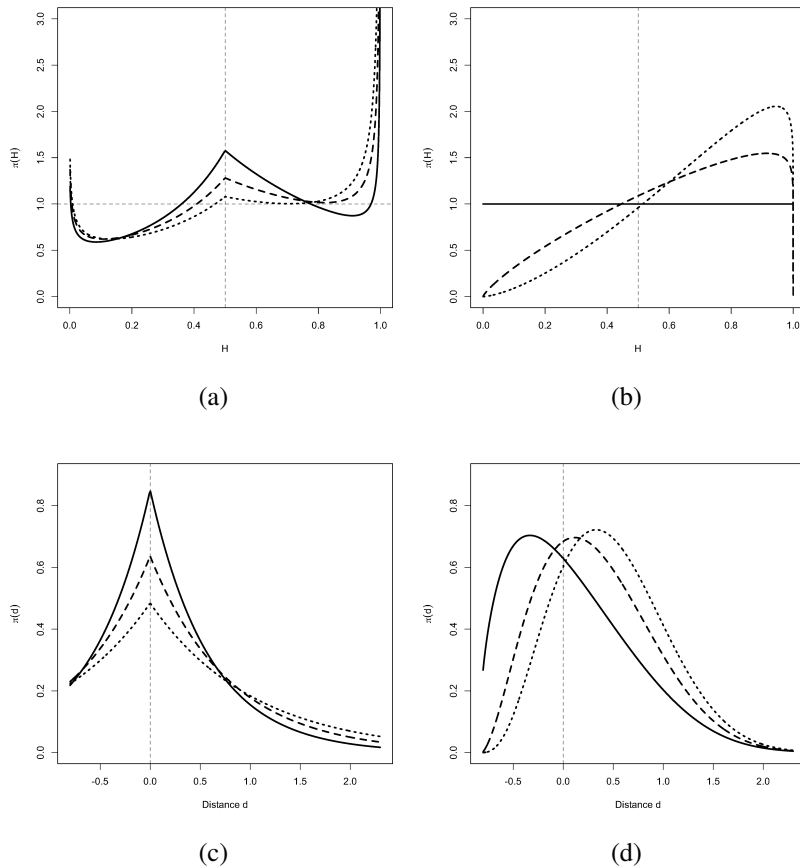


Figure 2: Panel a) and b) illustrate PC priors and Beta priors for H using three different scalings where $P(0.9 < H < 1)$ is set equal to 0.1 (solid), 0.15 (dashed) and 0.2 (dotted). Panel c) and d) give the corresponding priors for the distance measure d .

In order to understand the implications of different priors on H , we consider the corresponding priors for the distance measure, see Figure 2 (c) and (d). Here, we use negative distance to illustrate the priors when $H < 0.5$. Notice that when the parameters of the PC prior is changed, the mode of

the prior for distance is still at the base model, implying that we would not impose a more complicated model if the true data are in fact white noise. However, we do change the rate of shrinkage to the base model and increase the probability of large positive distances. In changing the parameters of the Beta prior, the mode of the prior for distance is shifted and pulled away from the base model. Especially, by assigning a uniform prior to H , the corresponding prior on distance is shifted to the left of the base model. The given negative distance corresponds to a mode at around $H \approx 0.28$. This seems like an unreasonable assumption, expect for cases where prior knowledge actually supports this. For example, theoretical results suggest that the Hurst exponent is around $1/3$ for various turbulence processes (Kang and Vidakovic, 2016).

2.3 The PC prior for the precision

The PC prior for H is derived assuming a fixed precision parameter τ . An alternative base model for fGn is to have no random effect ($1/\tau = 0$). Using the same principles as above, the resulting PC prior for τ is the type-2 Gumbel distribution (Simpson et al., 2017)

$$\pi(\tau) = \frac{\lambda}{2} \tau^{-3/2} \exp(-\lambda \tau^{-1/2}), \quad \lambda > 0, \quad (4)$$

which prescribes the size of the marginal precision of the fGn process. This density corresponds to using an exponential prior on the standard deviation.

Adopting the strategy of Simpson et al. (2017), the rate parameter λ is inferred using the probability statement $P(1/\sqrt{\tau} > U) = \alpha$, where α is a small probability while U specifies an upper limit for the marginal standard deviation $1/\sqrt{\tau}$. For example, if $\alpha = 0.01$, the marginal standard deviation is $0.31U$, after the precision τ is integrated out. In practice, the PC prior for the precision is assigned to the log-precision $\kappa = \log(\tau)$, which is represented in closed form as

$$\pi(\kappa) = \frac{\lambda}{2} \exp\left(-\lambda \exp\left(-\frac{\kappa}{2}\right) - \frac{\kappa}{2}\right), \quad \lambda = -\ln(\alpha)/U. \quad (5)$$

3 Identifying long- versus short-range dependency using Bayes factor

In analysing real time series, it is of major importance to model the true underlying dependency structure correctly as this will influence the inference made. Consider a stochastic regression model,

$$y_t = \nu_t + \epsilon_t, \quad (6)$$

where ν_t represents an underlying smooth trend while the noise term ϵ_t models random fluctuations. In analysing temperature data for different geographical regions, significance of a linear trend $\nu_t = \beta_0 + \beta_1 t$ is seen to depend heavily on whether the noise term is modelled by fGn or a first-order autoregressive process (Løvstletten and Rypdal, 2016).

3.1 Imposing identical priors for the hyperparameters of fGn and a first-order AR model

Autoregressive (AR) processes represent a commonly-applied class of stochastic processes, used to model short-range dependencies in time series. Specifically, assume a first-order AR(1) process defined by

$$x_t = \phi x_{t-1} + w_t, \quad t = 2, \dots, n,$$

where the first-lag autocorrelation coefficient $|\phi| < 1$. The innovations are assumed Gaussian, $w_t \sim N(0, \kappa^{-1})$, and x_1 has zero-mean and precision $\tau = \kappa(1 - \phi^2)$. The resulting covariance

matrix of $\mathbf{x} = (x_1, \dots, x_n)$ is Toeplitz and defined by the autocorrelation function $\gamma(k) = \phi^k$ for lags $k = 0, \dots, n - 1$.

Similarly to fGn, the first-order AR process represents a flexible version of uncorrelated white noise ($\phi = 0$). The Kullback-Leibler divergence is of order $n - 1$ and an invariant distance measure from AR(1) to white noise is

$$d(\phi) = \sqrt{-\ln(1 - \phi^2)}.$$

The resulting PC prior for ϕ (Simpson et al., 2017; Sørbye and Rue, 2017) is then expressed as

$$\pi(\phi) = \frac{\lambda}{2} \exp\left(-\lambda\sqrt{-\ln(1 - \phi^2)}\right) \frac{|\phi|}{(1 - \phi^2)\sqrt{-\ln(1 - \phi^2)}}, \quad |\phi| < 1 \quad (7)$$

where $\lambda > 0$.

In general, let M_0 and M_1 denote two different model components for the noise term ϵ in (6). Within a Bayesian framework, the two models can be compared using Bayes factor which quantifies the evidence in favour of one statistical model. The Bayes factor is defined as the ratio of the marginal likelihoods

$$\frac{\pi(\epsilon | M_0)}{\pi(\epsilon | M_1)} = \frac{\pi(M_0 | \epsilon) \pi(M_1)}{\pi(M_1 | \epsilon) \pi(M_0)}, \quad (8)$$

where $\pi(M_0)$ and $\pi(M_1)$ represent priors for the two model components. In general, Bayes factors are sensitive to the choices of $\pi(M_0)$ and $\pi(M_1)$ (Kass and Raftery, 1995). When the aim is to compare a model with an fGn component to a model with an AR(1) component, the Bayes factor will depend on the prior distributions for the precision parameters of the models and the priors for the parameters ϕ and H . The precision parameters have the same interpretation for both fGn and AR(1), hence we will use the same prior distribution for the precision of both models. Since we are using PC priors for both ϕ and H , these parameters are just a reparameterisation of the distance from the (unit variance) AR(1) and fGn model, to the *same* base model. By choosing the same rate parameter λ for the PC prior of both ϕ and H , the prior distributions will be the same even if the models and hyperparameters are different. This is a very convenient feature of PC priors which we make use of in the following examples.

3.2 Simulation results

To illustrate the capability of Bayes factor to identify fGn versus AR(1) structure, we perform a simulation study and generate fGn processes of different lengths n , for specified values of the Hurst parameter, scaled to have variance 1. The selected values for the Hurst parameter are $H = 0.7$, $H = 0.8$ and $H = 0.9$. These all give persistent fGn processes and are selected to mimic realistic values of the Hurst parameter in analysing long-memory processes. The PC prior for H is implemented using the scaling criterion $P(0.9 < H < 1) = 0.10$. This implies that the rate parameter of the exponential distribution of the distance measure is $\lambda \approx 1.70$. This rate parameter is used to find the corresponding prior for the first-lag correlation coefficient in fitting an AR(1) model to the data. The marginal precisions of both the fGn and AR(1) models are assigned the PC prior using $U = 1$ and $\alpha = 0.01$ in (5).

In evaluating the evidence given by Bayes factors, we have chosen to apply the categories given in Kass and Raftery (1995), in which Bayes factor should be larger than 3 to give a positive evidence in favour of one model. Note that these categories could also be specified in terms of twice the natural logarithm of the Bayes factor, which would have the same scale as the likelihood ratio test statistic. Table 1 illustrates that rather long time series are needed in order for the Bayes factor to

n	H	Bayes factor: Strength of evidence					Total BF > 3
		False BF < 1/3	No conclusion 1/3 < BF < 3	Positive 3 < BF < 20	Strong 20 < BF < 150	Very Strong BF > 150	
100	0.7	0.226	0.685	0.077	0.009	0.003	0.089
200	0.7	0.232	0.493	0.192	0.063	0.020	0.275
500	0.7	0.097	0.224	0.220	0.186	0.273	0.679
100	0.8	0.366	0.513	0.105	0.015	0.001	0.121
200	0.8	0.260	0.353	0.219	0.123	0.045	0.387
500	0.8	0.072	0.099	0.121	0.153	0.555	0.829
100	0.9	0.521	0.352	0.096	0.029	0.002	0.127
200	0.9	0.316	0.252	0.189	0.146	0.097	0.432
500	0.9	0.060	0.073	0.091	0.104	0.672	0.867

Table 1: Estimated proportion of Bayes factors in each of five groups, showing strength of evidence of the underlying process being fGn. Results are based on 1000 simulations for each value of the Hurst parameter H and for each time series length n .

give correct identification of the underlying fGn process. This seems natural as the long-memory properties might not be clearly apparent if the time series is short. We also notice that the results improve with higher values of the Hurst parameter for all lengths n , in which the long-memory structure of fGn would get more apparent. Especially, if $H = 0.9$ and the length of the time series is $n = 500$, we have very strong evidence of fGn in two-thirds of the cases. Naturally, the given results could be improved upon using a prior for H that puts more probability mass to the upper tail of the distribution.

3.3 Temperature data example

In this section, we analyse a real temperature data set to investigate whether we find evidence of fGn compared with AR(1) structure, using Bayes factor. The dataset "NOAA-CIRES 20th Century Reanalysis V2c", downloaded from http://www.esrl.noaa.gov/psd/data/20thC_Rean, contains reanalysed data from the period 1851 - 2014. The data are available for a 2 degree latitude times a 2 degree longitude global grid covering the earth and combine temperature measurements with model estimated data.

We assume a regression model for temperature as given in (6), where the trend is assumed to be either linear or non-linear. The non-linear trend is modelled using an intrinsic CAR model (Besag et al., 1991), also referred to as a second-order intrinsic Gaussian Markov random field defined on a line, see Rue and Held (2005, Ch. 3). The precision parameter of the model is assigned the penalised complexity prior in (4), using scaling parameters $(U, \alpha) = (0.10, 0.01)$ in (5). By choosing a small value of U , the prior imposes only small deviation from a straight line. Note that the given IGMRF model needs to be scaled to have a generalised variance equal to 1, such that we can use the same prior for τ independently of the time scale used for the analysis (Sørbye and Rue, 2014). Further, the hyperprior specifications for the AR(1) and fGn components used to model climate noise, are chosen to be the same as in the simulation study of Section 3.2.

Using the given data, it would be interesting to study evidence of fGn for different geographical regions, but here we only report results for aggregated data over land areas and sea-surfaces, respectively. We calculate the Bayes factor (8) in each case, using both an annual and quarterly time scale. The results are displayed in Table ??, which clearly illustrates that the evidence of fGn is

Time series	Trend parameter estimates		fGn parameter estimates				Evidence fGn Category
	$\hat{\beta}_1$	95% CI	\hat{H}	95% CI	$\hat{\sigma}$	95% CI	
Land (a)	0.006	(0.004, 0.008)	0.76	(0.67, 0.84)	0.22	(0.19, 0.26)	Strong
Land (q)	0.006	(0.004, 0.008)	0.77	(0.72, 0.82)	0.31	(0.29, 0.34)	Very strong
Sea (a)	0.004	(0.003, 0.006)	0.95	(0.91, 0.99)	0.21	(0.13, 0.37)	Strong
Sea (q)	0.004	(0.003, 0.006)	0.99	(0.97, 1.00)	0.39	(0.21, 0.74)	Very strong

Table 2: The mean estimates, 95% credible intervals and the evidence of fGn versus AR(1) given by Bayes factor, fitting a regression model with linear trend to global land and sea temperatures. We use both annual (a) and quarterly (q) time scales for the years 1851 - 2014.

Time series	Trend parameter estimates		fGn parameter estimates				Evidence fGn Category
	$\hat{\sigma}$	95% CI	\hat{H}	95% CI	$\hat{\sigma}$	95% CI	
Land (a)	0.047	(0.015, 0.102)	0.69	(0.58, 0.79)	0.20	(0.18, 0.23)	No conclusion
Land (q)	0.038	(0.008, 0.093)	0.75	(0.69, 0.80)	0.30	(0.28, 0.33)	Positive
Sea (a)	0.041	(0.012, 0.095)	0.92	(0.83, 0.98)	0.15	(0.10, 0.23)	No conclusion
Sea (q)	0.039	(0.012, 0.091)	0.98	(0.95, 1.00)	0.34	(0.16, 0.66)	No conclusion

Table 3: The mean estimates, 95% credible intervals and the evidence of fGn versus AR(1) given by Bayes factor, fitting a regression model with a smooth non-linear trend to global land and sea temperatures. We use both annual (a) and quarterly (q) time scales for the years 1851 - 2014.

very sensitive to the trend model. Using a linear trend, the evidence of fGn compared with AR(1) is strong ($\text{BF} > 20$) or very strong ($\text{BF} > 150$). Using a non-linear trend, the test using Bayes factor is inconclusive, except for the quarterly land data where there is weak evidence of fGn structure. It is well-known that non-linear trends can easily be interpreted incorrectly as long memory (Beran et al., 2013). The given results illustrates that a non-linear trend, even a very weak one (Figure 3), typically captures some of the long-range dependency structure of the time series, influencing the Bayes factor. However, it does not change the estimate of H , radically. We also notice that the quarterly data give higher values for the Hurst parameter and stronger evidence of fGn, than the annual data, which among others could be explained by the length of these series being four times longer than for the annual data.

The estimated Hurst parameter is seen to be higher for the sea-surface temperatures than the land temperatures, both using different trend models and different time scales. In fact, in the cases where the Hurst parameter is very close to 1 this could indicate that more components should be included in the model, for example to explain oscillatory behaviour. It is well-known that estimation of the Hurst parameter is not robust against departures from stationarity and trends and maybe also other models for the climate noise should be considered. Alternatively, the variability of the trend model could be adjusted to explain more of the fluctuations of the temperatures. A big advantage of using PC priors is that such adjustments are easily made, simply choosing a larger upper value of U in (5).

4 Concluding remarks

Penalised complexity priors (Simpson et al., 2017) is a recently developed framework to calculate priors for hyperparameters in Bayesian hierarchical models. The PC priors are intuitive in the sense that they are designed to shrink towards well-defined simple base model and will always have a




Figure 3: The upper panels show annual global land temperature (a) and annual global sea temperature (b) in the period 1851 - 2014, with estimated linear (red) and non-linear (blue) trends. The lower panels show the corresponding quarterly global land temperature (c) and sea temperature (d).

local mode at this base model where the distance measure is 0. Also, the parameters of the prior have a clear interpretation as they govern the rate of shrinkage to the base model and thereby the informativeness of the prior. In estimating the Hurst parameter of fractional Gaussian noise, the argument of having no knowledge about the parameter has encouraged the use of a non-informative prior. If we accept that a given prior on H should give reasonable results also on the distance scale, the uniform prior needs be ruled out. Similarly, changes in the parameters of Beta priors would shift the mode at distance scale and pull estimates away from the base model, making interpretation of these parameters non-intuitive and non-transparent.

In calculating Bayes factor, we take advantage of the fact that PC priors are invariant to reparameterisations. This implies that priors for two different hyperparameters simply represent two different transformations of the same prior on distance scale, as long as the corresponding model components represent flexible versions of the same base model. This is very useful in calculation of Bayes factors as these are known to be rather sensitive to prior choices. In the current analysis, we have focused on identifying the underlying noise term in a regression model as either AR(1) or fGn but the given ideas could also be applied for other model components as long as these share the same base model. The given analysis demonstrates that it is difficult to separate the trend from the model noise and that potential non-linearity in the trend could easily be interpreted as long memory. This makes it

important to control the smoothness of the underlying trend, in which a linear model is considered as the base model and deviation from the base model is controlled by a PC prior on the precision parameter. In general, the use of PC priors is very helpful to control the effect of different random effects and make the model components identifiable.

The option to combine generic functions with `R-INLA`, is a flexible and easy way to incorporate an fGn model component within the general class of latent Gaussian models. This makes it possible to analyse also more complex additive regression models than the ones presented here, in which an fGn component can be combined with for example smooth effects of covariates, seasonal and oscillatory effects, non-linear trend models and also spatial model components. However, to make full use of the computational power of `R-INLA`, the fGn model should be approximated to have Markov properties as this would give a sparse precision matrix of the underlying latent field. One alternative to achieve this is to approximate fGn with a weighted sum of AR(1) processes. Preliminary results (Myrvoll-Nilsen, 2016) indicate that such an approximation is promising, using only a few AR(1) components. In fact, we have seen that aggregation of just three or four AR(1) processes give an excellent approximation of fGn, when the weights and coefficients of the AR components are chosen to match the autocorrelation structure of fGn (Myrvoll-Nilsen et al., 2017). The resulting decomposition of fGn can potentially be linked to linear multibox energy balance models (Fredriksen and Rypdal, 2017), in analysing temperature data. This requires further study and will be reported elsewhere.

Acknowledgement

The 20th Century Reanalysis V2c data is provided by the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA, from their Web site at <http://www.esrl.noaa.gov/psd/>. The authors wish to thank Hege-Beate Fredriksen for valuable discussions and for processing the data to give aggregated data for land and sea-surface temperatures.

References

- Abry, P. and Veitch, D. (1998). Wavelet analysis of long range dependent traffic. *IEEE Transactions on Information Theory*, 44:2–15.
- Benhmehdi, S., Makarava, N., Menhamidouche, N., and Holschneider, M. (2011). Bayesian estimation of the self-similarity exponent of the Nile River fluctuation. *Nonlinear Processes in Geophysics*, 18:441–446.
- Beran, J., Feng, Y., Ghosh, S., and Kulik, R. (2013). *Long-Memory Processes: Probabilistic Properties and Statistical Methods*. Springer, Heidelberg.
- Besag, J., York, J., and Mollié, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43:1–59.
- Fredriksen, H.-B. and Rypdal, M. (2017). Long-range persistence in global surface temperatures explained by linear multi-box energy balance models. *arXiv:1701.07355*.
- Giraitis, L., Kokoszka, P., Leipus, R., and Teyssiere, T. (2003). Rescaled variance and related tests for long memory in volatility and levels. *Journal of Econometrics*, 112:265–294.
- Hurst, H. E. (1951). Long term storage capacities of reservoirs. *Transactions of the American Society for Civil Engineers*, 116:776–808.

- Kang, M. and Vidakovic, B. (2016). A note on Bayesian wavelet-based estimation of scaling. *arXiv:1605.01146*.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:773–795.
- Kolmogorov, A. N. (1940). Wienersche spiralen und einige andere interessante Kurven im Hilbertschen Raum. *C. R. (Doklady) Acad. Sci. URSS (N.S.)*, 26:115–118.
- Løvstetten, O. and Rypdal, M. (2016). Statistics of regional surface temperatures after 1900: Long-range versus short-range dependence and significance of warming trends. *Journal of Climate*, 29:4057–4068.
- Makarava, N., Benmehdi, S., and Holschneider, M. (2011). Bayesian estimation of self-similarity exponent. *Physical Review E*, 84:021109.
- Makarava, N. and Holschneider, M. (2012). Estimation of the Hurst exponent from noisy data: a Bayesian approach. *The European Physical Journal B*, 85:1–6.
- Mandelbrot, B. B. and Ness, J. W. V. (1968). Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10:422–437.
- Martins, T. G., Simpson, D., Lindgren, F., and Rue, H. (2013). Bayesian computing with INLA: New features. *Computational Statistics & Data Analysis*, 67:68–83.
- McCoy, E. and Walden, A. (1996). Wavelet analysis and synthesis of stationary long-memory processes. *Water Resources Research*, 5:26–56.
- Myrvoll-Nilsen, E. (2016). *Computationally efficient Bayesian approximation of fractional Gaussian noise using ARI processes*. Master thesis, Norwegian University of Science and Technology.
- Myrvoll-Nilsen, E., Sørbye, S. H., and Rue, H. (2017). An approximate fractional Gaussian noise model with linear computational complexity. *Manuscript in progress*.
- Peng, C.-K., Buldyrev, S., Havlin, S., Simons, M., Stanley, H. E., and Goldberger, A. (1994). Mosaic organization of DNA nucleotides. *Physical Review E*, 49:1685–1689.
- Riebler, A., Sørbye, S. H., Simpson, D., and Rue, H. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, 25:1145–1165.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*, volume 104 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*, 71:319–392.
- Simpson, D. P., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H. (2017). Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, 32:1–28.
- Sørbye, S. H. and Rue, H. (2014). Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spatial Statistics*, 8:39–51.

Sørbye, S. H. and Rue, H. (2017). Penalised complexity priors for stationary autoregressive processes. *Journal of Time Series Analysis*, doi:10.1111/jtsa.12242.

5 Appendix: Incorporating fGn models within latent Gaussian models

We have implemented the fGn model within the R-INLA package (see www.r-inla.org) which uses nested Laplace approximations to do Bayesian inference for the class of latent Gaussian models (Rue et al., 2009; Martins et al., 2013). We used the generic model “`rgeneric`” allowing a (Gaussian) latent model component to be defined using R and it is passed into the `inla-program` which is written in C. This is a convenient feature that allows us to prototype new models easily, despite the fact that the fGn model does not have any Markov properties. In general, Markov properties would improve the computational efficiency significantly, see Rue and Held (2005, Ch. 2).

The `rgeneric` code to represent the fGn model is given below. We make use of the Toeplitz structure of the covariance matrix to speed up the computations. The model is defined as a function returning properties of the Gaussian model, like its graph, the precision matrix, the normalising constant, and so on.

```
rgeneric.fgn =
  function (cmd = c("graph", "Q", "initial", "log.norm.const",
                  "log.prior", "quit", "mu"), theta = NULL)
{
  stopifnot(require(ltsa) && require(FGN) && require(HKprocess))
  theta.fun = function() {
    return(list(prec = function(x) exp(x),
              H = function(x) exp(x)/(1 + exp(x))))
  }
  interpret.theta = function() {
    fun = theta.fun()
    return(list(prec = fun[[1]](theta[1]),
              H = fun[[2]](theta[2])))
  }
  graph = function()
    return(inla.as.sparse(matrix(1, n, n)))
  Q = function() {
    param = interpret.theta()
    S = toeplitz(acvfFGN(param$H, n-1))
    return(inla.as.sparse(param$prec * TrenchInverse(S)))
  }
  log.norm.const = function() {
    param = interpret.theta()
    r = acvfFGN(param$H, n-1)
    S = toeplitz(r) / param$prec
    val = ltzc(r, rep(0, n))
    val2 = -n/2 * log(2*pi) -0.5*val[2] + n/2 * log(param$prec)
    return(val2)
  }
  log.prior = function() {
    param = interpret.theta()
    lprior = inla.pc.dprec(param$prec, u=1, alpha=0.01, log=TRUE) +
      log(param$prec)
    return(lprior + LPRIOR.FUN.H(theta[2]))
  }
}
```

```

mu = function() {
  return (numeric(0))
}
initial = function() {
  return(c(4, 0))
}
quit = function() {
  return(invisible())
}

cmd = match.arg(cmd)
val = do.call(cmd, args = list())
return(val)
}

```

We can now include the given fGn model component together with other latent components. This is an easy example, where the function returning the log-prior for ρ (the internal representation of H) is defined in the file `fgn-prior.R` (not included here).

```

library(INLA); library(FGN)
n=200; H=0.8
lprior.H = function(theta) return (0) ## replace this fun
model = inla.rgeneric.define(model = rgeneric.fgn, n=n,
                             LPRIOR.FUN.H = lprior.H)
y = scale(SimulateFGN(n,H))
r = inla(y ~ -1 + f(idx, model=model),
        data = data.frame(y = y, idx = 1:n, iid = 1:n),
        control.family = list(hyper = list(
                               prec = list(initial = 12, fixed=TRUE))))
summary(r)

```