

Controlling the Flexibility of Non-Gaussian Processes Through Shrinkage Priors

Rafael Cabral*, David Bolin†, and Håvard Rue‡

Abstract. The normal inverse Gaussian (NIG) and generalized asymmetric Laplace (GAL) distributions can be seen as skewed and semi-heavy-tailed extensions of the Gaussian distribution. Models driven by these more flexible noise distributions are then regarded as flexible extensions of simpler Gaussian models. Inferential procedures tend to overestimate the degree of non-Gaussianity in the data and therefore we propose controlling the flexibility of these non-Gaussian models by adding sensible priors in the inferential framework that contract the model towards Gaussianity. In our venture to derive sensible priors, we also propose a new intuitive parameterization of the non-Gaussian models and discuss how to implement them efficiently in *Stan*. The methods are derived for a generic class of non-Gaussian models that include spatial Matérn fields, autoregressive models for time series, and simultaneous autoregressive models for aerial data. The results are illustrated with a simulation study and geostatistics application, where priors that penalize model complexity were shown to lead to more robust estimation and give preference to the Gaussian model, while at the same time allowing for non-Gaussianity if there is sufficient evidence in the data.

MSC2020 subject classifications: Primary 62F15, 62M20; secondary 62M40.

Keywords: Bayesian, penalized complexity, priors, non-Gaussian, generalized asymmetric Laplace, normal inverse Gaussian, SPDE, Matérn covariance.

1 Introduction

Gaussian processes are the most common class of models to describe spatial and temporal dependence in Bayesian hierarchical models. Due to their well-established theory, flexibility, and practicality, Gaussian processes are also a fundamental building block in spatial and temporal statistics. However, additional flexibility is needed in several applications, and Gaussian processes are replaced by more flexible non-Gaussian Lévy processes (Ken-Iti, 1999). In this paper, we study processes driven by generalized hyperbolic (GH) noise, more specifically the normal inverse Gaussian (NIG) and generalized asymmetric Laplace (GAL) subfamilies, which contain the Gaussian distribution as a particular case and are semi-heavy-tailed.

*Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Kingdom of Saudi Arabia, rafael.medeiroscabral@kaust.edu.sa

†Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Kingdom of Saudi Arabia

‡Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Kingdom of Saudi Arabia

1.1 Literature review

Autoregressive processes with GH innovating terms are discussed in Ghasami et al. (2020) and can model extreme market movements not captured by the Gaussian model (Dhull and Kumar, 2021). Bibby and Sørensen (2003) presented more sophisticated models for financial time series, namely stochastic processes whose marginal distributions or the distributions of increments (or both) are generalized hyperbolic, including classical diffusion processes and stochastic volatility models (Barndorff-Nielsen, 1997; Barndorff-Nielsen and Shephard, 2001, 2002). The empirical distributions of the log-returns of financial time series are too heavy for satisfactory fitting by normal densities, and the GH distribution can better capture the stylized features of financial data.

In the field of spatial statistics, Bolin (2014) provided a class of non-Gaussian random fields with Matérn covariance function, constructed as solutions to stochastic partial differential equations (SPDEs) driven by Generalized Hyperbolic noise. The previous models were applied in the context of geostatistics (Wallin and Bolin, 2015), joint modeling of multivariate random fields (Bolin and Wallin, 2020), and to model continuously repeated measurement data collected longitudinally (Asar et al., 2020), where the non-Gaussian models led to improved predictive power in datasets where there were sharp spikes or jumps in the observed data that the Gaussian models oversmoothed.

The aforementioned papers performed parameter estimation of non-Gaussian models via likelihood maximization techniques or by the method of moments. Bayesian estimation was carried out for the same geostatistical framework of Wallin and Bolin (2015) by Walder and Hanks (2020). A Bayesian framework has also been considered for stochastic volatility models (Xie and Shen, 2021; Nakajima and Omori, 2012), for a vector of autoregressive processes (Karlsson et al., 2021), and GARCH models (Deschamps, 2012), and all previously cited papers utilized the generalized hyperbolic skew Student's t (GHSST) distribution for the driving noise. In the same papers, a gamma prior was chosen for the leptokurtosis parameter and a normal prior for the skewness parameter of the GHSST distribution. There was no principled motivation for the use of these priors other than mathematical or computational convenience, and the gamma prior places 0 mass at the base Gaussian model, thus forcing overfitting of the data (Simpson et al., 2017).

1.2 Motivation and contributions

Here we present the layout of the paper and highlight what are its main contributions to the literature on non-Gaussian processes from a Bayesian perspective. We find in the current literature challenges when it comes to model interpretability, sensible prior selection, and ease of implementation which we address next.

Parameterization and interpretability: Parameterizations of the NIG and GAL distributions involve 4 parameters that are used to regulate the mean, variance, skewness, and kurtosis of the distribution. However, the parameterizations are not themselves property-based leading to a difficult interpretation of the parameters. This also presents a difficulty when constructing priors for the parameters, since, for instance, the mean




Figure 1: The box plots show the posterior means of the leptokurtosis parameter η for i.i.d. Gaussian simulated data with sample sizes of $n = 50, 100, 200$ with penalized complexity (PC) and uniform (Uni) priors, repeated over 1000 experiments. The right plot shows the right tail of the symmetric NIG PDF for several values of η .

and variance will depend not only on the location and scale parameter but also on the other parameters. In Section 2 we provide a property-based parameterization of the NIG and GAL distributions that preserves the mean and covariance structure of the Gaussian model.

Priors that avoid overfitting and lead to a robust estimation: Figure 1 shows posterior means of the leptokurtosis parameter η of the NIG distribution when fitting i.i.d. standard Gaussian data. These were obtained by numerical integration of the unnormalized $\pi(\eta|\mathbf{y}) \propto \prod_i^n \pi_{\text{NIG}}(y_i|\eta)\pi(\eta)$, where π_{NIG} is the probability density function (PDF) of the standardized and symmetric NIG distribution of Subsection 2.1. The Gaussian and Cauchy distributions are limit distributions when $\eta \rightarrow 0$ and $\eta \rightarrow \infty$, respectively. If we use a uniform prior for small samples the results suggest the necessity of a non-Gaussian model when the simpler Gaussian model should be the preferred one. We should prefer the Gaussian model not only for parsimony’s sake but also because for Gaussian simulated data, the expected predictive performance will be the highest for the Gaussian model if a strictly proper scoring rule is used to evaluate the predictive accuracy (Gneiting and Raftery, 2007). The tendency to overestimate the degree of non-Gaussianity in the data is more severe when a non-Gaussian model is added as a latent component in a hierarchical model and can be prevalent even for samples sizes of 500, as our simulation study in Section 5 shows. This issue motivates an inferential framework that contracts non-Gaussian models towards Gaussianity in the absence of sufficient evidence of non-trivial leptokurtosis or asymmetry. In Section 3 we construct prior distributions for the skewness and leptokurtosis parameters of the NIG and GAL distribution based on the penalized complexity (PC) priors principled approach of Simpson et al. (2017).

Many models, one framework: We presented in the previous subsection a wide range of models, defined in discrete or continuous space, for time series or spatial data. We unify these models into a single generic class, which we examine in Section 2. We also offer in Section 4 a set functions for the *Stan* platform (Stan Development Team, 2020) that allow a straightforward implementation of these models.

Bayesian analysis: A simulation study is conducted to compare the PC priors with other priors choices in Section 5, and the developed methods are applied to two spatial

datasets in Section 6. The analysis showed that the PC priors achieve the sought contraction towards the Gaussian model when there is not enough convincing evidence in the data of non-Gaussianity and lead to more robust estimation.

Finally, Section 7 contains a summary and discussion of future work and possible extensions.

2 A flexible extension of Gaussian models

The GH distribution (Barndorff-Nielsen, 1978) can be conveniently represented as a variance mixture of normal distributions, where the mixing distribution is a generalized inverse Gaussian (GIG) random variable. If a random variable Λ follows the GH distribution, then it has the following hierarchical representation

$$\Lambda|V \sim N(\mu + \beta V, V), \quad V \sim \text{GIG}(\lambda, \delta, \sqrt{\alpha^2 - \beta^2}), \quad (1)$$

where δ , μ , and β are scale, location, and skewness parameters, respectively, while α and λ are two shape parameters. The constraints on these parameters are found in Barndorff-Nielsen et al. (2012). The GH distribution includes many widely used distributions as special cases, including the Gaussian, t-Student, Cauchy, NIG, and GAL distributions. The NIG subclass is obtained by setting $\lambda = -1/2$, leading to an inverse Gaussian (IG) mixing distribution. The PDF of a NIG distribution is

$$f(x; \alpha, \beta, \delta, \mu) = \frac{\alpha \delta}{\pi} \frac{K_1\left(\alpha \sqrt{\delta^2 + (x - \mu)^2}\right)}{\sqrt{\delta^2 + (x - \mu)^2}} e^{\delta \sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)},$$

where $K_\lambda(x)$ is the modified Bessel function of the second kind of order λ . On the other hand, the GAL subclass is obtained from the limit $\delta \rightarrow 0$, leading to a gamma mixing variable and the PDF

$$f(x; \lambda, \alpha, \beta, \mu) = \frac{(\alpha^2 - \beta^2)^\lambda}{\sqrt{\pi}(2\alpha)^{\lambda - \frac{1}{2}} \Gamma(\lambda)} |x - \mu|^{\lambda - \frac{1}{2}} K_{\lambda - \frac{1}{2}}(\alpha|x - \mu|) e^{\beta(x - \mu)}.$$

As a necessary step in building explicable priors, in the following Subsections 2.1 and 2.2 we present a standardized parameterization of the NIG and GAL distributions which also aims to achieve an orthogonal interpretation of the parameters. Then in Subsection 2.3, we define the distribution of multivariate models that are driven by NIG or GAL noise.

2.1 Standardized parameterization (η, ζ)

In the previous parameterization $(\lambda, \alpha, \beta, \delta, \mu)$, the location (μ) and scale (δ) parameters do not correspond to the mean and standard deviation of the distribution. For the NIG distribution setting $\mu = 0$ and $\delta = 1$ leads to $E[\Lambda] = \beta/(-\alpha^2 + \beta^2)$ and $V[\Lambda] = \alpha^2/(\alpha^2 - \beta^2)^{3/2}$, and so the mean and variance depend on the degree of asymmetry and

Dist.	λ	α	β	δ	μ	S	EK
NIG	$-1/2$	$\frac{1}{\tilde{\sigma}}\sqrt{\frac{1}{\eta} + \zeta^2}$	$\frac{\zeta}{\tilde{\sigma}}$	$\frac{\tilde{\sigma}}{\sqrt{\eta}}h$	$-\tilde{\sigma}\zeta h$	$\frac{3\zeta\eta}{\sqrt{h+h\zeta^2\eta}}$	$\frac{3\eta(1+5\zeta^2\eta)}{h(1+\zeta^2\eta)}$
GAL	$h\eta^{-1}$	$\frac{1}{\tilde{\sigma}}\sqrt{\frac{2}{\eta} + \zeta^2}$	$\frac{\zeta}{\tilde{\sigma}}$	0	$-\tilde{\sigma}\zeta h$	$\frac{\zeta\eta(3+2\zeta^2\eta)}{\sqrt{h(1+\zeta^2\eta)^3}}$	$\frac{3\eta(1+4\zeta^2\eta+2\zeta^4\eta^2)}{h(1+\zeta^2\eta)^2}$

Table 1: Parameters of the GH distribution, skewness (S), excess kurtosis (EK) for the GAL and NIG special cases, where $\tilde{\sigma} = 1/(\sqrt{1 + \eta\zeta^2})$.

leptokurtosis of the distribution. There are alternative parameterizations of the GH distribution in the literature (Paolella, 2007; Prause et al., 1999), including location-scale invariant ones, but these continue to not be property-based.

As in Niekerk and Rue (2021) we posit that when constructing priors the mean and the standard deviation should be fixed, instead of location and scale parameters. This not only eases interpretation but also allows assigning the priors for the mean and scale of the Gaussian model in non-Gaussian models, which is very convenient when implementing these non-Gaussian models in practice (otherwise reasonable priors for the location and scale would have to depended on α and β). Table 1 shows the conversion between the standardized parameterization and the more conventional one in eq. (1), along with some relevant moments (consider for now $h = 1$). The variance-mean mixture representation belonging to this parameterization is

$$\Lambda|V \sim N\left(\frac{1}{\sqrt{1 + \eta\zeta^2}}\zeta(V - h), \frac{1}{1 + \eta\zeta^2}V\right), V \sim \begin{cases} \text{IG}(1, \eta^{-1}) & \text{(NIG)} \\ \text{Gamma}(\eta^{-1}, \eta^{-1}) & \text{(GAL)}. \end{cases} \quad (2)$$

A location and scale parameter can be added by the usual transformation $m + \sigma\Lambda$, and the mean and variance will be m and σ^2 , regardless of the other parameters' value. The parameter η is related to the degree of non-Gaussianity since Λ converges to a Gaussian random variable when $\eta \rightarrow 0$, and as we increase η , the excess kurtosis increase. The parameter ζ is related to the asymmetry of the random variable since Λ is symmetric when $\zeta = 0$, and when $\zeta > 0$, it is skewed to the right. Here and henceforth, we will make use of the notions of “base model”, “flexible model” and “flexibility parameter” defined in Simpson et al. (2017). We see the non-Gaussian model parameterized by (η, ζ) as a flexible extension of the base Gaussian model since it contains the Gaussian model as a special case (when $\eta = \zeta = 0$) and deviations from the Gaussian model are quantified by the parameters η and ζ to which we refer as flexibility parameters. The densities of the NIG and GAL distributions with the previous parameterization are displayed in Figure 2. We observe that the larger the value of η , the larger the asymmetry induced by the same value of the parameter ζ .

2.2 Standardized and orthogonal parameterization (η^*, ζ^*)

There is still some confounding between η and ζ in the standardized parameterization since the excess kurtosis increases with ζ , so a prior for η derived for a symmetric model may not shrink the models towards Gaussianity as much as we anticipate if the data is

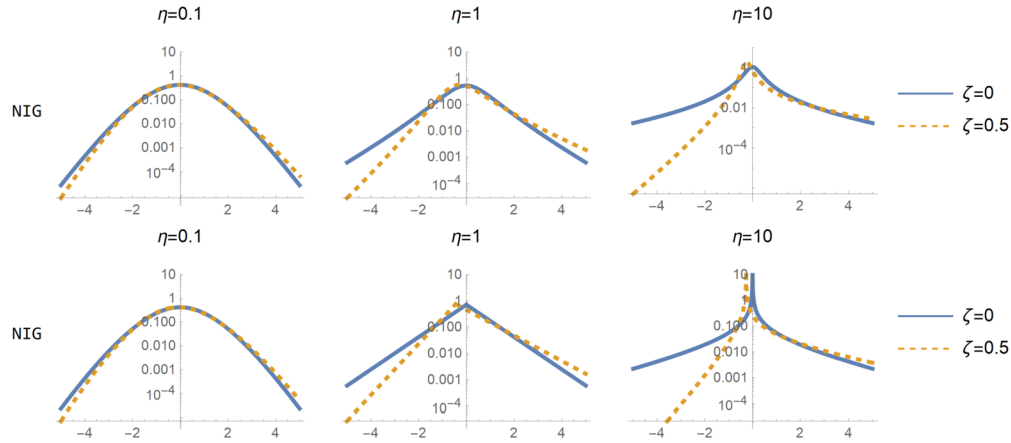


Figure 2: Density of the NIG (top) and GAL (bottom) distributions in log scale for different values of η and ζ with $\sigma = 1$ and $h = 1$.

asymmetric. Thus, we need to associate η with some interpretable model property and then guarantee that the property is invariant with the second flexibility parameter. This will also allow using the PC prior of η , which is derived in Section 3 for the symmetric case, in asymmetric data.

We find the kurtosis to be hard to interpret (it is not clear what an increase in kurtosis of 1 means in practice) and so we prefer to associate η with the likelihood of large events. For the NIG distribution this probability can be approximated by $P(|\Lambda| > x) \propto x^{-3/2} \exp(-\eta^{-1/2} \xi_{NIG} x)$ for large x (Hammerstein, 2016). The dependency of this probability with the skewness parameter ζ comes through the rate ξ_{NIG} :

$$\xi_{NIG} = 1 + \zeta^2 \eta - |\zeta| \sqrt{\eta(1 + \zeta^2 \eta)},$$

which is equal to 1 in the symmetric case ($\zeta = 0$). We require the probability $P(|\Lambda| > x)$ to be invariant with the skewness parameter, at least for large x . This can be achieved by the parameter transformations $\eta^* = \eta \xi_{NIG}^{-2}$ and $\zeta^* = \zeta \sqrt{\eta}$. The same transformation applies for the GAL distribution where one should take

$$\xi_{GAL} = \sqrt{1 + \zeta^2 \eta} (\sqrt{2 + \zeta^2 \eta} - |\zeta| \sqrt{\eta}).$$

We note that these transformations on the flexibility parameters do not affect the mean (0) and the variance (1) of the standardized parameterization. The quantity ξ is also related to the excess kurtosis, namely $\text{EKurt}[\Lambda] \approx 9/\xi^2$, for both the NIG and GAL distributions. Therefore, by guaranteeing that $P(|\Lambda| > x)$ is invariant with ζ^* for large x , we also guarantee that the kurtosis of the noise (and the process' marginal distributions) remains approximately the same for different values of ζ^* , leading to a more orthogonal interpretation of the parameters.

2.3 Multivariate models driven by non-Gaussian noise

Let \mathbf{x}^G be a random vector that follows a multivariate Gaussian distribution with dimension n , mean m , and precision matrix $\mathbf{Q} = \sigma^{-2}\mathbf{D}^T\mathbf{D}$. It can be expressed through

$$\mathbf{D}(\mathbf{x}^G - m) \stackrel{d}{=} \sigma\mathbf{Z}, \tag{3}$$

where $\mathbf{Z} = [Z_1, \dots, Z_n]^T$ is a vector of i.i.d. standard Gaussian variables. When \mathbf{x}^G is derived from a process defined in continuous space, one usually has $Z_i \stackrel{ind.}{\sim} N(0, h_i)$ for some predefined constant h_i (for instance, the distance between locations in an autoregressive process), and so the precision matrix is $\mathbf{Q} = \sigma^{-2}\mathbf{D}^T \text{diag}(\mathbf{h})^{-1}\mathbf{D}$. The non-Gaussian extension for \mathbf{x} consists in replacing the driving noise distribution:

$$\mathbf{D}(\mathbf{x} - m) \stackrel{d}{=} \sigma\mathbf{\Lambda}, \tag{4}$$

where $\mathbf{\Lambda} = [\Lambda_1, \dots, \Lambda_n]^T$ is a vector of independent and standardized NIG or GAL random variables that depend on parameters η and ζ . Using the parameterizations in Subsections 2.1 and 2.2 for $\mathbf{\Lambda}$, the non-Gaussian random vector \mathbf{x} has the same mean and covariance matrix as the Gaussian random vector \mathbf{x}^G but with more flexible sample path properties and marginal distributions whose kurtosis and skewness are regulated by the flexibility parameters η and ζ . Appendix B of the Supplementary Material (Cabral et al., 2022a) contains more properties about non-Gaussian models defined via eq. (4). It is important to mention that these models are not uniquely specified by their covariance or precision matrices but through the matrix \mathbf{D} . In Section 4 we show that many of the non-Gaussian models mentioned in the Introduction belong to the generic class of models defined by eq. (4).

A variance-mean mixture representation of the non-Gaussian random vector \mathbf{x} is obtained by considering the random vector $\mathbf{V} = [V_1, \dots, V_n]^T$ and the predefined vector $\mathbf{h} = [h_1, \dots, h_n]^T$, containing the mixing distributions and some predefined constants, respectively. In the standardized parameterization it is:

$$\begin{aligned} \mathbf{x}|\mathbf{V} &\sim N\left(m + \frac{\sigma}{\sqrt{1 + \zeta^2\eta}}\zeta\mathbf{D}^{-1}(\mathbf{V} - \mathbf{h}), \frac{\sigma^2}{1 + \zeta^2\eta}\mathbf{D}^{-1}\text{diag}(\mathbf{V})\mathbf{D}^{-T}\right), \\ V_i &\stackrel{ind.}{\sim} \begin{cases} \text{IG}(h_i, \eta^{-1}h_i^2) & \text{(NIG Noise)} \\ \text{Gamma}(h_i\eta^{-1}, \eta^{-1}) & \text{(GAL Noise)}. \end{cases} \end{aligned} \tag{5}$$

To broaden the applicability of the multivariate model \mathbf{x} to data measured at irregularly spaced locations or time intervals, we need to consider the predefined vector \mathbf{h} . Take, for instance, a continuous Gaussian random walk of order 1 (RW1) process evaluated at locations x_1, x_2, \dots with distances $h_i = x_{i+1} - x_i$. In this situation, the model is defined by the increments $x_{i+1} - x_i$ following a normal distribution with mean 0 and variance $h_i\sigma^2$ (Rue and Held, 2005). Therefore, a larger distance between observations will lead to a noise Gaussian distribution with increased variance. A similar transformation applies when we discretize a NIG or GAL continuous process, but the

increased distance between locations will not only change the variance to $h_i\sigma^2$ but also will the shape of the distribution. This is a standard result from the theory of Lévy processes, where the new densities are found by raising the NIG and GAL characteristic functions to the power h_i and then taking the inverse-Fourier transform (see Appendix A of the Supplementary Material and Barndorff-Nielsen et al. (2012)). The variables Λ_i will share the σ , η and ζ parameters, but h_i will generally be unique to each noise variable (see Table 1). If the observations are measured at equally spaced locations we set \mathbf{h} to be a vector of ones.

3 Contraction towards Gaussianity

Our goal is to control the flexibility of non-Gaussian models and shrink the model towards Gaussianity if there is insufficient evidence in the data of excess kurtosis and asymmetry. One natural and simple way to achieve this contraction is by adding in the inferential framework priors for the flexibility parameters η and ζ that give preference to the Gaussian model.

The prior distributions derived next are applicable to a large variety of models and can be easily added to pre-existing model implementations. We will make use of the penalized complexity (PC) priors methodology of Simpson et al. (2017). The approach consists of setting an exponential distribution on the unidirectional *distance* $d(\cdot)$, measuring the added complexity of the more flexible model, with respect to the base model. The measure of complexity is based on the Kullbeck-Leibler divergence (KLD), and $d(\eta) = \sqrt{2KLD(\pi(\mathbf{x}|\eta) \parallel \pi(\mathbf{x}|\eta = 0))}$. PC priors give preference to the simpler model and tend to avoid overfitting by default because the mode of the prior distribution is located at the base (Gaussian) model: $\pi(\eta) = \theta|d'(\eta)| \exp(-\theta d(\eta))$. The previous prior is defined up to a rate parameter θ which needs to be calibrated (see Subsection 3.3).

Remark 1. *An important feature of the class of non-Gaussian models that we are studying, which are obtained by replacing Gaussian white noise with more flexible alternatives, is that from an information-theoretic perspective, the deviation from the Gaussian model only depends on the flexible noise and does not depend on the covariance structure encoded by the matrix \mathbf{D} . This may not seem obvious at first, since different choices for the matrix \mathbf{D} lead to different sample path behaviors and marginal properties (see Figure 4). However, if we construe the density of the non-Gaussian random vector as linearly transformed NIG or GAL densities $\mathbf{x} = m + \sigma\mathbf{D}^{-1}\mathbf{\Lambda}$, and similarly for the Gaussian random vector $\mathbf{x}^G = m + \sigma\mathbf{D}^{-1}\mathbf{Z}$, this property follows directly from the well know invariance of the KLD under monotonic transformations:*

$$KLD(\mathbf{x} \parallel \mathbf{x}^G) = KLD(m + \sigma\mathbf{D}^{-1}\mathbf{\Lambda} \parallel m + \sigma\mathbf{D}^{-1}\mathbf{Z}) = KLD(\mathbf{\Lambda} \parallel \mathbf{Z}) = \sum_{i=1}^n KLD(\Lambda_i \parallel Z_i).$$

The previous property implies that the distribution of the PC prior does not depend on \mathbf{D} , σ , or m and is therefore applicable to any non-Gaussian model that can be expressed via eq. (4). Thus, without loss of generality, we consider in this section processes with 0 mean ($m = 0$) and driving noise with unit variance ($\sigma = 1$). In the




Figure 3: KLD computed by numerical integration for $h = 1$ in log-log scale. The plot on the left shows $KLD(\eta, \zeta = 0)$ for NIG and GAL noise, and the plot on the right shows $KLD(\zeta|\eta)$ for three fixed values of η for NIG noise.

following subsection, we derive the PC prior of η for the symmetric case ($\zeta = 0$), and then in Subsection 3.2, we find the PC prior of ζ conditioned on η .

3.1 PC prior distribution for the first flexibility parameter

The non-Gaussian extension presented in Subsection 2.3 preserves the mean and covariance structure of \mathbf{x} , so we can assume that both the flexible and base models have the same scale parameter σ and the same spatial (or temporal) range parameter, such as the parameter κ in the Matérn model of eq. (7). These parameters cancel out when computing $KLD(\mathbf{x} \parallel \mathbf{x}^G)$ as seen in Remark 1, and therefore $KLD(\mathbf{x} \parallel \mathbf{x}^G)$ only depends on η and \mathbf{h} . The non-Gaussian noise Λ_i can follow either a NIG or a GAL distribution, and for both cases, $KLD(\Lambda_i \parallel Z_i)$, which should be seen as a function of η , behaves quadratically near the base model (see Figure 3 (a)). In the symmetric case the parameterizations $(\eta, \zeta = 0)$, and $(\eta^*, \zeta^* = 0)$ are equivalent, therefore the same prior distribution is assigned for η and η^* .

Theorem 3.1. *Let \mathbf{D} be a $n \times n$ non-singular matrix. Also, let the flexible model \mathbf{x} with density $\pi(\mathbf{x}|\eta, \zeta, h)$ be defined by $\mathbf{D}\mathbf{x} = \mathbf{\Lambda}$, where $\mathbf{\Lambda}$ is a vector of independent NIG or GAL noise defined in Subsection 2.1. Then, for small η , the KLD is*

$$KLD(\pi(\mathbf{x}|\eta, \zeta = 0) \parallel \pi^G(\mathbf{x}|\eta = 0, \zeta = 0)) = \frac{3}{16} \left(\sum_{i=1}^n \frac{1}{h_i^2} \right) \eta^2 + \mathcal{O}(\eta^4).$$

where $\pi^G(\mathbf{x}|\eta = 0, \zeta = 0)$ is the density of the base Gaussian model.

Proof. See Appendix D of the Supplementary Material. □

We are mostly interested in penalizing the added complexity of the more flexible non-Gaussian model in a neighborhood around the base Gaussian model, and as suggested by Simpson et al. (2017), a Taylor expansion around the base model is done, and the second order expansion is used as the measure of added complexity.

Corollary 3.1.1. *The distance measure is $d(\eta) = \sqrt{2KLD(\eta)} \propto \eta$ and so the PC prior for η and η^* will follow an exponential distribution with some rate parameter θ_η . We note that this PC prior induces a LASSO (L-1) style penalty on the parameter η (and η^*), since $\log \pi(\eta) = -\theta_\eta \eta + \text{const}$.*

3.2 PC prior distribution for the second flexibility parameter

We derive in this subsection the PC prior for the second flexibility parameter. As seen in Section 2.1, the impact of the parameter ζ on the NIG and GAL distributions depend on the value of η (if $\eta = 0$, then ζ has no impact). Therefore, we derive the PC prior for ζ conditionally on η , and consider that the base model is a non-Gaussian model driven by symmetric noise Λ^{Sym} (with parameters η and $\zeta = 0$), and the flexible model \mathbf{x} is driven by asymmetric noise Λ (with parameters η and ζ). In Figure 3 (b), we show the KLD between the NIG noises Λ_i and Λ_i^{Sym} for several values of η and ζ .

A Taylor expansion around $\eta = 0$ and $\zeta = 0$ yields a quadratic dependency with η which is only accurate when η is very close to 0:

$$KLD(\Lambda_i \parallel \Lambda_i^{Sym}) = (3\eta^2/(4h_i) + \mathcal{O}(\eta^3))\zeta^2 + \mathcal{O}(\zeta^4).$$

Therefore, we use the following upper bound of the KLD as a measure of added complexity, which provides a more reasonable approximation.

Theorem 3.2. *Under the same conditions of Theorem 3.1, we have for NIG driving noise*

$$KLD(\pi(\mathbf{x}|\eta, \zeta) \parallel \pi^{Sym}(\mathbf{x}|\eta, \zeta = 0)) \leq \frac{n}{2}\eta\zeta^2.$$

This inequality also holds for GAL driving noise when $\eta < \min_{i=1, \dots, n} h_i$.

Proof. See Appendix D of the Supplementary Material. □

If we use the previous upper bound as a measure of complexity, the distance $d(\zeta)$ is $\sqrt{2KLD(\zeta)} \propto \sqrt{\eta}|\zeta|$ and by setting an exponential distribution on $d(\zeta)$ with rate parameter θ_ζ , the density of $\zeta|\eta$ is found to be

$$\pi(\zeta|\eta) = \frac{1}{2}\theta_\zeta\sqrt{\eta}\exp(-\theta_\zeta\sqrt{\eta}|\zeta|). \quad (6)$$

Corollary 3.2.1. *With the standardized and orthogonal parameterization (η^*, ζ^*) , the skewness parameter is $\zeta^* = \sqrt{\eta}\zeta$ and so $KLD \leq n\zeta^{*2}/2$. We then have $d(\zeta) \propto |\zeta^*|$ and the PC prior for ζ^* is a Laplace distribution with rate parameter θ_ζ . This prior also acts as a LASSO (L-1) penalty on the estimation of ζ^* , since $\log \pi(\zeta^*) = -\theta_\zeta|\zeta^*| + \text{const}$.*

3.3 Calibration of the PC priors

In line with the *weakly informative* prior framework of Simpson et al. (2017) and Gelman et al. (2017), the calibration of the PC priors, i.e, the choice of θ_η and θ_ζ , is based on

the user defining the prior probabilities $P(\eta^* > U_\eta) = \alpha_\eta$ and $P(|\zeta^*| > U_\zeta) = \alpha_\zeta$ and in turn the calibration parameters are $\theta_\eta = -\log(\alpha_\eta)/U_\eta$ and $\theta_\zeta = -\log(\alpha_\zeta)/U_\zeta$. Of course, sensible choices for the upper-tail values U and probabilities α must come from an understating of how different values of η^* and ζ^* impact the process. This understanding can be informed by, for instance, plotting sample paths of the process for increasing values of η^* (see Figure 4), stopping when we observe unreasonably large spikes or jumps, at $\eta^* = U_\eta$, and then setting a low probability that $\eta_* > U_\eta$, say equal to 0.01. A similar procedure can be applied when calibrating the PC prior of ζ^* , but one should look for asymmetries in the number of large jumps or spikes, for instance for positive ζ^* , there should be more positive jumps than negative jumps.

More interpretable distribution features can be used in the calibration, such as the probability that large marginal events occur (larger than 3 times the marginal standard deviation, for instance) which increase with η . However, deriving these probabilities for new models can be unhandy, since the marginal distributions do not have a closed-form PDFs or CDFs. Nevertheless, we pursue this path in Appendix E of the Supplementary Material to calibrate the PC prior of η for Matérn and OU processes.

3.4 Comparison with other prior distributions

For the GHSST distribution (subclass of the GH distribution) seen in Subsection 1.1, the flexibility parameters are ν and β , the base model corresponds to $\nu \rightarrow \infty$ and $\beta = 0$ and gamma and normal priors are commonly chosen, respectively. In our parameterization, this would suggest an inverse gamma distribution for η since the base model is at $\eta \rightarrow 0$, and a normal prior for ζ . An inverse gamma prior decays slower to 0 as η increases (see Figure 5), and it has no mass at the base model ($\eta = 0$), so it acts as a repellent from the simpler Gaussian model, having the opposite effect that we necessitate. Also, a Gaussian prior for ζ may not achieve as much contraction as the prior in eq. (6), which follows a Laplace distribution for a fixed value of η , and it does not take into account that the asymmetry induced by a particular value of ζ increases with η .

The Fisher information matrix does not seem to be available in closed form which makes the use of the Jeffreys priors for the flexibility parameters impractical for the class of models defined by eq. (4). Figure 5 shows a numerical approximation of the Jeffreys prior density for η , based on the univariate pdf of the NIG distribution. It has the mode at the base model, but unlike the PC priors, which decay exponentially, it is almost flat.

4 Many models, one framework

Our general model in eq. (4) contains a wide variety of non-Gaussian models as special cases that allow departures from Gaussianity within realizations. Here, we review some non-Gaussian models defined in discrete and continuous space and present a set of functions for *Stan* that allow an easy and efficient implementation of these models.

4.1 Models defined on discrete space

A Gaussian autoregressive process (AR) of order 1 assumes that $\sqrt{1-\rho^2}x_1$ and the differences $\{x_{i+1} - \rho x_i, i > 1\}$ follow independent Gaussian white noise $N(0, \sigma^2)$ and $|\rho| < 1$. If we assume that the differences instead are non-Gaussian white noise, we obtain the linear system $\mathbf{D}_{AR1}\mathbf{x} = \mathbf{\Lambda}$, where $\mathbf{D}_{AR1} = (d_{i,j}) \in \mathbb{R}^{n \times n}$ with $d_{1,1} = \sqrt{1-\rho^2}$, $d_{i,i-1} = -\rho$, $d_{i,i} = 1$ for $i > 1$, and the other matrix elements are 0. Likewise, the matrix \mathbf{D} for higher-order AR processes is a lower triangular Toeplitz matrix containing the autoregression coefficients. A vector of autoregressive processes with no intercept also has a representation $\mathbf{D}_{VAR}\mathbf{x} = \mathbf{\Lambda}$ when extending the model to non-Gaussianity, by stacking the multivariate time series into a single vector \mathbf{x} . To fit areal data Walder and Hanks (2020) proposed a simultaneously autoregressive (SAR) model driven by non-Gaussian noise that also can be represented by eq. (4) for an appropriately specified matrix \mathbf{D} .

4.2 Models defined on continuous space

A famous class of processes in spatial statistics are stationary Gaussian processes with Matérn covariance functions (Matérn, 1960). Gaussian processes with this covariance function can be obtained as solutions of the SPDE:

$$(\kappa^2 - \Delta)^{\alpha/2} X(\mathbf{s}) = \sigma \mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathbb{R}^d, \quad (7)$$

where κ is a spatial range parameter, $\Delta = \sum_i \partial^2 / \partial x_i^2$ is the Laplace operator, $\alpha > d/2$ is a smoothness parameter, and $\mathcal{W}(\mathbf{s})$ is a Gaussian white noise process (Whittle, 1963). The approximation to discrete space in (Lindgren et al., 2011) uses the finite element method (FEM) to the stochastic weak formulation of the previous SPDE. It begins by expressing the process $X(\mathbf{s})$ as a weighted sum of basis functions, $X(\mathbf{s}) = \sum_{i=1}^n w_i \psi_i(\mathbf{s})$, and it was shown that the weights $\mathbf{w} = [w_1, \dots, w_n]^T$ follow the system $\mathbf{D}\mathbf{w} = \mathbf{Z}$, where the predefined vector \mathbf{h} of Section 2.3 has elements $h_i = \int_{\mathbb{R}^d} \psi(\mathbf{f}\mathbf{l}\mathbf{s}) d\mathbf{s}$. The random field $X(\mathbf{s})$ evaluated at locations $\mathbf{s}_1, \mathbf{s}_2, \dots$ composes the vector $\mathbf{x} = [X(\mathbf{s}_1), X(\mathbf{s}_2), \dots]^T$ and it is given by the linear combination $\mathbf{x} = \mathbf{A}\mathbf{w}$, where \mathbf{A} is the projector matrix with elements $A_{ij} = \psi_i(\mathbf{s}_j)$. Building the matrices \mathbf{D} and \mathbf{A} may seem hard at first since it involves the finite element method, but these can be easily built with the functions `inla.mesh.2D`, `inla.mesh.fem` and `inla.spde.make.A` in the R package *INLA*, and the user only needs to supply the location of the observations and some tuning parameters for the discretization mesh (Bakka et al., 2018).

Bolin (2014) extended the previous results to Type-G Matérn random fields by replacing the Gaussian noise process $\mathcal{W}(\mathbf{s})$ with a non-Gaussian noise process $\mathring{\mathbf{A}}(\mathbf{s})$. As discussed in (Wallin and Bolin, 2015), for SPDE models, the increments need to be closed under convolution, and so we can only consider the NIG and GAL subclasses of the GH distribution. It was shown that the stochastic weights now follow the system $\mathbf{D}\mathbf{w} = \mathbf{\Lambda}$, where the matrix \mathbf{D} is the same as in the Gaussian case seen before.

Table 2 lists differential operators associated with several stochastic processes and Figure 4 contains their sample paths, obtained by the SPDE approach, considering a

Model	CRW1	OU	CRW2	Matérn $\alpha = 2$
\mathcal{D}	∂_t	$\kappa + \partial_t$	∂_t^2	$\kappa^2 - \partial_t^2$

Table 2: Differential operators associated with different models.




Figure 4: Noise and sample paths of several models for $\eta = 10^{-6}$ (left) and $\eta = 1$ (right), for $\sigma = 1$ and $\zeta = 0$.

Gaussian noise process (left panel) and a NIG noise process (right panel). The characterization of the marginal distributions is in Appendix C of the Supplementary Material. Matérn and Ornstein–Uhlenbeck (OU) processes are widely used in applications. The continuous random walk model of order 2 (CRW2) is also known as an integrated Brownian motion process, and it has been used by Diggle et al. (2014), and Zhu and Dunson (2017). Whenever the noise takes an extreme value (for instance, near location 0.8), the CRW1 and OU processes will exhibit a distinct jump, and the RW2 and Matérn processes will exhibit a kink (discontinuity in the first derivative).

4.3 Implementation in Stan

Estimating non-Gaussian models defined by $\mathbf{D}\mathbf{x} = \mathbf{\Lambda}$ in the *Stan* platform (Stan Development Team, 2020) can be done by declaring the random vectors \mathbf{x} and \mathbf{V} with the hierarchical representation of eq. (5). However, the dimension of \mathbf{V} is the same as the dimension of \mathbf{x} , which can be very large in some applications, and since \mathbf{V} needs to be estimated, one can expect long sampling times. We can integrate out the mixing variables V_i in eq. (5) to reduce the dimension of the parameter space being explored in *Stan*, which can lead to a significant speedup. Note that $\mathbf{x} = \mathbf{D}^{-1}\mathbf{\Lambda}$ and if \mathbf{D} is

non-singular, then the multivariate transformation method yields:

$$\pi(\mathbf{x}) = |\mathbf{D}| \prod_{i=1}^n \pi_{\Lambda_i}([\mathbf{D}\mathbf{x}]_i), \quad (8)$$

where π_{Λ_i} is the PDF of a NIG or GAL distribution.

The *Stan* function `nig_model` returns the log-likelihood of a NIG model based on eq. (8) using the standardized and orthogonal parameterization and it is implemented in github.com/rafaelcabral96/nigstan. The declaration of \mathbf{x} takes the following form:

```
x      ~ nig_model(D, etas, zetas, h, 1)
etas   ~ exp(theta_eta)
zetas  ~ double_exponential(0, 1.0/theta_zeta)
```

where the last argument of `nig_model` is an integer with value 1 if the log-determinant of \mathbf{D} should be computed (if \mathbf{D} depends on parameters), or 0 otherwise. The last two lines are the declaration of the PC priors for η^* and ζ^* . A location and scale parameter can be added by the usual transformation $m + \sigma\mathbf{x}$. We can fit all models in Subsections 4.1 and 4.2 in *Stan* using the `nig_model` function by specifying the appropriate matrix \mathbf{D} . Contrary to the hierarchical representation in eq. (5), if we work with eq. (8) there is no need to estimate the auxiliary random vector \mathbf{V} and invert the matrix \mathbf{D} , and thus sampling times of hours can be reduced to minutes. The results in the following sections were obtained via this implementation. We demonstrate in more detail how *Stan* can be used to fit non-Gaussian models for several applications in rafaelcabral96.github.io/nigstan/. Unfortunately, it is currently not possible to implement models driven by GAL noise in *Stan* based on eq. (8), since modified Bessel functions of the second kind with fractional order are currently not available in *Stan*.

5 Simulation studies

To study how the priors on the flexibility parameters perform under different conditions, we consider two simulation sets. In the first, we investigate if the posterior distributions of η^* and ζ^* are close to the true values of these parameters used to simulate the sample paths. We verify the contraction towards Gaussianity induced by the PC priors and the ability to allow for non-Gaussianity when the latent field is significantly non-Gaussian. For the second simulation set, we check how sensitive the posterior distributions of η^* and ζ^* are to irregularities in the data for different choices of priors. The *R* scripts can be found in the Supplementary Material (Cabral et al., 2022b).

5.1 Implementation details

The following simulation settings were considered. The response is $y_i \sim N(x_i, 0.7)$, where \mathbf{x} is non-Gaussian latent field \mathbf{x} , defined by eq. (4) with mean $m = 0$. The model parameters are $\{\sigma, \eta^*, \zeta^*\}$ and the matrix \mathbf{D} was chosen so that \mathbf{x} corresponds to a Matérn model (with $\alpha = 2$ and $\kappa = 0.2$) as in Subsection 4.2. We fit the model to simulated data with parameters as shown in Table 3. We considered sample sizes n




Figure 5: Prior distributions for the two flexibility parameters in log-log scale (we only plot for positive values of ζ^*).

chosen from $\{50, 100, 500, 1000\}$ for the first simulation set and $n = 500$ for the second simulation set.

The prior distribution for σ is $\text{IGamma}(1, 1)$ and we consider several prior configurations for the flexibility parameters η^* and ζ^* : **PC1**: $\eta^* \sim \exp(\theta_\eta = 30)$, and $\zeta^* \sim \text{Laplace}(\theta_\zeta = 13)$; **PC2**: $\eta^* \sim \exp(\theta_\eta = 2.3)$, and $\zeta^* \sim \text{Laplace}(\theta_\zeta = 1)$; **IG1/N1**: $\eta^* \sim \text{IGamma}(2, 0.1)$, and $\zeta^* \sim \text{N}(0, 0.3)$; **IG2/N2**: $\eta^* \sim \text{IGamma}(2, 0.43)$, and $\zeta^* \sim \text{N}(0, 1)$; **Jeffrey**: η^* follows the Jeffrey prior computed numerically, and $\zeta^* \sim \text{Uni}(-50, 50)$ **Uniform**: $\eta^* \sim \text{Uni}(0, 50)$, and $\zeta^* \sim \text{Uni}(-50, 50)$;

The models were fitted in *Stan* with the *cmdstanr* interface (Gabry and Češnovar, 2021), and $N = 200$ replications were run for each scenario, prior distribution, and sample size. Each replication consisted of 1000 warmup iterations and 1000 sampling iterations. The IG1 and IG2 priors for η^* were chosen so that they would have the same mean as the PC1 and PC2 priors, and therefore they could potentially achieve the same level of contraction to the Gaussian model. Likewise, the variances of the N1 and N2 priors assigned to ζ^* are the same as the variances of the PC1 and PC2 priors for ζ^* . These prior distributions are plotted in Figure 5.

5.2 Results

Figure 6 shows the posterior means and widths of the credible intervals for the parameters η^* and ζ^* . For brevity, we only show the results of scenarios 1 and 3 and for the sample sizes of 50, 100, and 1000. The remaining figures can be found in Appendix F of the Supplementary Material. For sample sizes up to 500 we see a positive bias in the estimation of η^* when using the uniform or Jeffreys priors, namely we have posterior means larger than 1, indicating a clear departure from Gaussianity in

Scenarios	Simulation set 1	Simulation set 2
Scenario 1	$\sigma = 1, \eta^* = \zeta^* = 0$ (Gaussian)	Gaussian with no jumps
Scenario 2	$\sigma = 1, \eta^* = 2, \zeta^* = 0$	Gaussian with two jumps of size 25
Scenario 3	$\sigma = 1, \eta^* = 5, \zeta^* = 1$	Gaussian with two jumps of size 50

Table 3: Parameters of the simulated data for each simulation set and scenario.

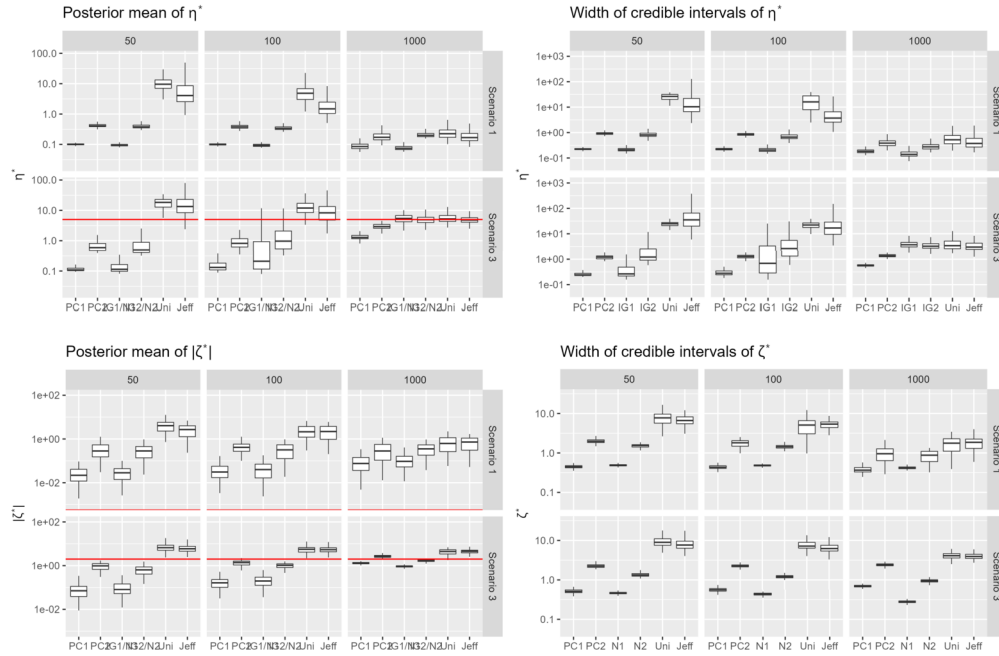


Figure 6: Histograms of the posterior means (top) and widths of the posterior credible intervals (bottom) for η^* (left) and ζ^* (right) and for different sample sizes, prior configurations, and scenarios of the simulation set 1.

the latent field, where there is none. The posterior means of η^* are smaller when utilizing the PC priors compared to the Jeffrey priors. The reason for this is that near the base model the PC prior can be seen as a tilted Jeffrey priors (Simpson et al., 2017), $\pi(\eta^*) = I(\eta^*)^{1/2} \exp(-\theta_\eta m(\eta^*))$, where $I(\eta^*)$ is the Fisher information, and $m(\eta^*) = \int_0^{\eta^*} \sqrt{I(s)} ds$.

The differences in the posterior means of η^* between the PC priors and the Inverse Gamma priors are not substantial because the IG priors were scaled so they would have the same mean as the PC priors. Although the posterior credible intervals of η^* (based on the 5% and 95% quantiles) often have smaller widths when utilizing the PC priors, compared to the IG priors, which is partly because the IG priors are more “spread out” (see Figure 5). The PC and IG priors for η^* lead to posterior means that are closer to 0, which are more consistent between replications and with smaller widths for the credible intervals). Similar observations apply for the posterior inferences of ζ^* .

5.3 Additional simulation studies

The results of simulation set 2 can be found in Appendix F of the Supplementary Material. The PC priors led to an estimation that is less sensitive to jumps in the

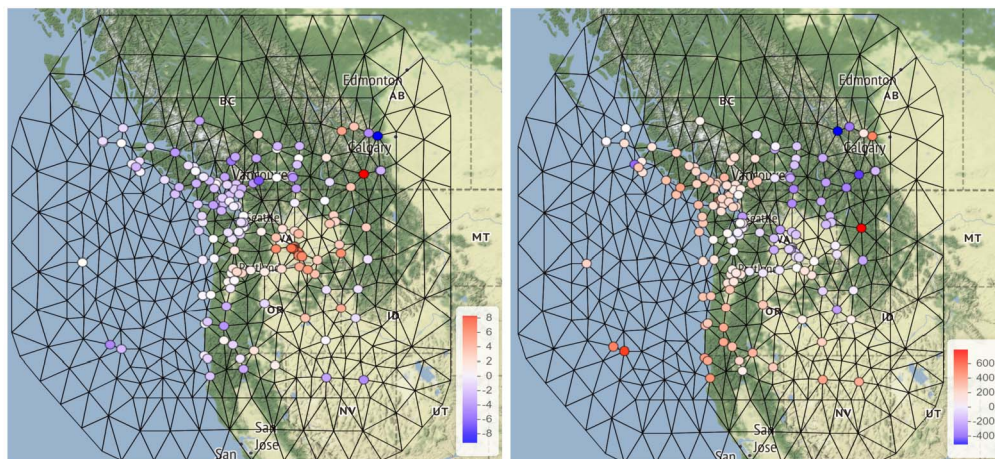


Figure 7: Measurements of temperature (left) and pressure (right) data subtracted by their sample means with the FEM mesh on the background.

processes than the other priors. Overall, the two simulation studies suggest that the PC priors perform well in a variety of scenarios, leading to a more robust estimation with regards to irregularities in the data and giving preference to the Gaussian model, while at the same time allowing for non-Gaussianity if there is enough support in the data.

6 Application

In this section, we illustrate the impact of the PC priors on a geostatistics application and how they achieve the sought contraction towards the Gaussian model when there is not enough convincing evidence in the data of non-Gaussianity. The *R* scripts can be found in the Supplementary Material (Cabral et al., 2022b).

6.1 Dataset and model implementation

Figure 7 shows temperature and pressure measurements at 157 different locations in the North American Pacific Northwest, where the sample mean was subtracted from the data in both cases. Bolin and Wallin (2020) considered a Gaussian model for the temperature data, while the pressure data appeared to have some localized spikes and short-range variations, which were better captured with a non-Gaussian NIG model.

We consider the geostatistical modelling paradigm of Wallin and Bolin (2015) where the field of interest $\mathbf{X}(\mathbf{s})$ is observed at N locations $\mathbf{s}_1, \dots, \mathbf{s}_N$, generating observations y_1, \dots, y_N , that follow $\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}$, where \mathbf{A} is the projector matrix. The vector $\boldsymbol{\epsilon}$ is i.i.d. Gaussian noise with variance $\sigma_{\boldsymbol{\epsilon}}^2$ and \mathbf{x} is a non-Gaussian random field ($\mathbf{D}\mathbf{x} = \sigma_{\mathbf{x}}\boldsymbol{\Lambda}$), where \mathbf{D} is obtained by the FEM approximation on the Matérn SPDE with $\alpha = 2$ (see Subsection 4.2). We only considered non-Gaussian models with NIG driving noise. The

	Prior choice 1 (conservative)		Prior choice 2 (near-uniform)	
	Temperature	Pressure	Temperature	Pressure
σ_ϵ	0.79 (0.67, 0.93)	63.00 (54.59, 72.46)	0.79 (0.68, 0.94)	62.52 (54.74, 71.07)
σ_x	8.25 (6.25, 10.64)	336.77 (257.84, 427.72)	10.19 (2.09, 7.20)	469.34 (325.10, 649.12)
κ	0.93 (0.60, 1.32)	0.37 (0.25, 0.50)	1.11 (0.78, 1.51)	0.50 (0.35, 0.66)
η^*	0.21 (0.02, 0.55)	2.04 (0.70, 4.05)	4.35 (1.02, 10.70)	22.01 (5.91, 48.22)
ζ^*	0.02 (-0.11, 0.16)	-0.01 (-0.11, 0.10)	0.09 (-0.27, 0.46)	-0.25 (-0.69, 0.19)

Table 4: Posterior means and 95% credible intervals of the Matérn SPDE model parameters driven by NIG noise for the two datasets and prior choices.

mesh used for the FEM approximation is shown in Figure 7, and it was built with the *inla.mesh.2d* function of the *R-INLA* package (Rue and Martino, 2007) and consisted of 394 nodes. The models were again implemented in *Stan* with the *cmdstanr* interface, and 4 parallel chains were run with 500 warmup iterations and 1000 sampling iterations.

The PC priors of η^* and ζ^* are exponential and Laplace, respectively, and to study the impact of different calibration choices for the priors on the posterior inferences, we assigned two sets of PC priors. The first set was calibrated by choosing $\alpha_\eta = 0.01$ and $P(|\zeta^*| > 0.3) = 0.01$, a conservative choice which should lead to a significant contraction, where α_η is the likelihood of having twice as much large jumps in the process compared to the Gaussian process (see Appendix E of the Supplementary Material). For the second set we chose $\alpha_\eta = 0.95$ and $P(|\zeta^*| > 4) = 0.01$, leading to near-uniform priors. We refer to each model as \mathcal{M}_{PC} and \mathcal{M}_{Unif} , respectively. There were no warning messages in the *Stan* program output, and the diagnostics indicated a good mixing, namely a split- \hat{R} smaller than 1.05 for the model parameters and large effective sample sizes. We also tried uniform priors for η^* and ζ^* , but the chains did not converge.

6.2 Estimation results

Table 4 shows that the posterior means of η^* and ζ^* are closer to 0 for the first set of conservative PC priors, compared to the second set of near-uniform priors. The posterior credible intervals of ζ^* include the value 0, so a symmetric model seems adequate for both datasets. The posterior means and standard deviations of the field $X(\mathbf{s})$ are plotted in Figure 8 in a prediction grid consisting of 100000 nodes for the \mathcal{M}_{PC} model. We can observe a smoother field for the temperature data and several localized spikes for the pressure data.

To compare the performance of the models we performed a leave-one-out cross-validation (loocv) study. We also compared the two non-Gaussian models we fitted before (\mathcal{M}_{PC} and \mathcal{M}_{Unif}) with a Gaussian model for the latent vector \mathbf{x} (\mathcal{M}_{Gauss}). Performing a loocv study for non-Gaussian models by refitting a *Stan* model at every held-out observation can be very expensive. The computation of the loocv estimates is made efficient in *Stan* through the *loo* package (Vehtari et al., 2017), however, the *loo* function released several warnings that the approximated loo estimates were not reliable. Instead, we performed a pseudo loocv study similar to the ones performed in Bolin and Wallin (2020) and Bolin and Kirchner (2020), which assumes that the

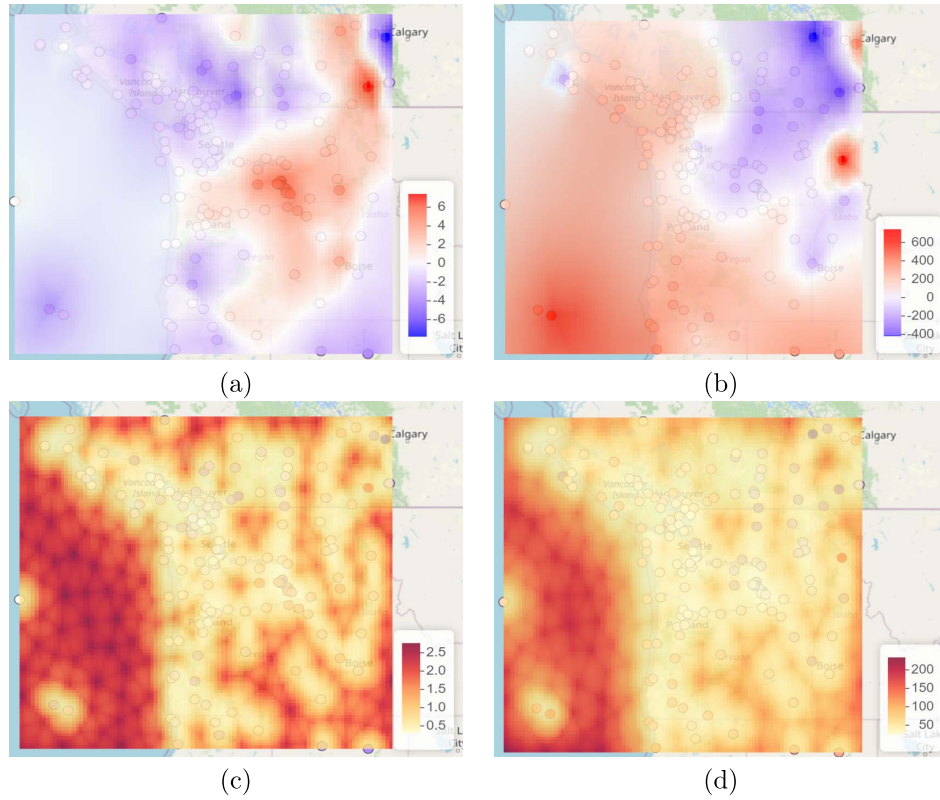


Figure 8: Plots related to temperature and pressure data are on the left and right, respectively. Plots (a)-(b) show the posterior means and plots (c)-(d) show the standard deviations of the latent field $X(\mathbf{s})$ for the first set of conservative PC priors.

posterior distributions of the parameters when removing observation i , $\pi(\boldsymbol{\theta}|\mathbf{y}_{-i})$, is equal in distribution to $\pi(\boldsymbol{\theta}|\mathbf{y})$, where $\boldsymbol{\theta} = (\sigma_x, \sigma_y, \kappa, \eta^*, \zeta^*)$. Still, the posterior distribution of the latent field \mathbf{x} and of the mixing variables \mathbf{V} is affected when removing an observation from the dataset.

We computed the mean squared error (MSE), mean absolute error (MAE), and continuously ranked probability score (CRPS) of Gneiting and Raftery (2007) between the observed data y_i and the held-out predictions for each observation $i = 1, \dots, 157$, and then averaged the results. We show these results in Table 5. The NIG model with near-uniform priors had a higher MSE and CRPS than the Gaussian model for the temperature data, so a very flexible non-Gaussian model for the temperature data did not translate into higher predictive performance. If we look at the MSE and MAE estimates, the NIG model with PC priors had the best predictive power for both datasets, while the CRPS estimates indicate that the NIG model with near-uniform priors performed the best for the pressure data. These results suggest that the NIG model with PC priors did not overfit the data and had the highest predictive power if we use the posterior

Data	Model	MSE	MAE	CRPS
Temperature	Gaussian	4.477	1.506	0.859
	NIG with PC priors	4.331	1.471	0.856
	NIG with Unif. priors	4.881	1.491	0.863
Pressure	Gaussian model	24036.848	98.647	54.893
	NIG model with PC priors	20246.371	87.207	53.339
	NIG model Unif. priors	21155.342	87.743	52.887

Table 5: Mean squared error (MSE), mean absolute error (MAE), and continuously ranked probability score (CRPS) for the leave-one-out cross-validation predictions.

mean of the latent field $\mathbf{X}(\mathbf{s})$ as a spatial predictor.

7 Conclusions and discussion

There is a need for an inferential framework that shrinks models driven by non-Gaussian noise towards Gaussianity in order to avoid overfitting and considering a non-Gaussian model when there is not enough evidence in the data for asymmetry and leptokurtosis. We have proposed achieving this contraction by selecting weakly informative priors that penalize deviations from the simpler Gaussian model, quantified by the KLD. With an appropriate parameterization of the non-Gaussian models, this approach leads to priors that behave essentially like a LASSO penalty on the non-Gaussian flexibility parameters.

The processes presented in this paper are more flexible, while at the same time being tractable, and are an attractive inherently robust alternative to Gaussian processes when the adequacy of the Gaussianity assumption is questionable. Barndorff-Nielsen and Blaesild (1981) suggested that the GH distribution is well-qualified for robustness studies when the deviation from Gaussianity is in the form of asymmetry and leptokurtosis, and work is currently being done to investigate the robustness properties of this more general class of models.

Future work also includes approximating these non-Gaussian models in the *INLA* framework (Rue and Martino, 2007). A direct implementation of these models in *INLA* is not possible because *INLA* is only feasible for latent Gaussian models with up to 20 hyperparameters. If we condition the non-Gaussian vector \mathbf{x} on the mixing variables V_i , we obtain a latent Gaussian model but the number of the mixing variables V_i , which would be hyperparameters in *INLA*, is larger than 20 in almost all applications.

Supplementary Material

Supplementary material for “Controlling the flexibility of non-Gaussian processes through shrinkage priors” (DOI: [10.1214/22-BA1342SUPPA](https://doi.org/10.1214/22-BA1342SUPPA); .pdf). The Supplementary Material contains the proofs of the theorems, more properties of the non-Gaussian vector \mathbf{x} and process $X(t)$, and additional tables and simulation figures.

Supplementary code for “Controlling the flexibility of non-Gaussian processes through

shrinkage priors” (DOI: [10.1214/22-BA1342SUPPB](https://doi.org/10.1214/22-BA1342SUPPB); .zip). This supplementary file contains the *R* code of the motivation example in section 2, simulation study of section 5, and the application in section 6.

References

- Asar, Ö., Bolin, D., Diggle, P. J., and Wallin, J. (2020). “Linear mixed effects models for non-Gaussian continuous repeated measurement data.” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 69(5): 1015–1065. [MR4166856](#). doi: <https://doi.org/10.1111/rssc.12405>. 1224
- Bakka, H., Rue, H., Fuglstad, G.-A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D., and Lindgren, F. (2018). “Spatial modeling with R-INLA: A review.” *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(6): e1443. [MR3873676](#). doi: <https://doi.org/10.1002/wics.1443>. 1234
- Barndorff-Nielsen, O. (1978). “Hyperbolic distributions and distributions on hyperbolae.” *Scandinavian Journal of statistics*, 151–157. [MR0509451](#). 1226
- Barndorff-Nielsen, O. and Blaesild, P. (1981). “Hyperbolic distributions and ramifications: Contributions to theory and application.” In *Statistical distributions in scientific work*, 19–44. Springer. [MR0656143](#). 1242
- Barndorff-Nielsen, O. E. (1997). “Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling.” *Scandinavian Journal of Statistics*, 24(1): 1–13. URL <http://www.jstor.org/stable/4616433> [MR1436619](#). doi: <https://doi.org/10.1111/1467-9469.00045>. 1224
- Barndorff-Nielsen, O. E., Mikosch, T., and Resnick, S. I. (2012). *Lévy processes: theory and applications*. Springer Science & Business Media. 1226, 1230
- Barndorff-Nielsen, O. E. and Shephard, N. (2001). “Non-Gaussian Ornstein-Uhlenbeck-Based Models and Some of Their Uses in Financial Economics.” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(2): 167–241. URL <http://www.jstor.org/stable/2680596> [MR1841412](#). doi: <https://doi.org/10.1111/1467-9868.00282>. 1224
- Barndorff-Nielsen, O. E. and Shephard, N. (2002). “Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models.” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 64(2): 253–280. URL <http://www.jstor.org/stable/3088799> [MR1904704](#). doi: <https://doi.org/10.1111/1467-9868.00336>. 1224
- Bibby, B. M. and Sørensen, M. (2003). “Chapter 6 - Hyperbolic Processes in Finance.” In Rachev, S. T. (ed.), *Handbook of Heavy Tailed Distributions in Finance*, volume 1 of *Handbooks in Finance*, 211–248. Amsterdam: North-Holland. URL <https://www.sciencedirect.com/science/article/pii/B978044450896650008X> 1224
- Bolin, D. (2014). “Spatial Matérn fields driven by non-Gaussian noise.” *Scandinavian*

- Journal of Statistics*, 41(3): 557–579. MR3249417. doi: <https://doi.org/10.1111/sjos.12046>. 1224, 1234
- Bolin, D. and Kirchner, K. (2020). “The rational SPDE approach for Gaussian random fields with general smoothness.” *Journal of Computational and Graphical Statistics*, 29(2): 274–285. MR4116041. doi: <https://doi.org/10.1080/10618600.2019.1665537>. 1240
- Bolin, D. and Wallin, J. (2020). “Multivariate type G Matérn stochastic partial differential equation random fields.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1): 215–239. MR4060983. 1224, 1239, 1240
- Cabral, R., Bolin, D. and Rue, H. (2022). “Supplementary Material for “Controlling the Flexibility of Non-Gaussian Processes Through Shrinkage Priors”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/22-BA1342SUPPA>. 1229
- Cabral, R., Bolin, D. and Rue, H. (2022). “Supplementary code for “Controlling the Flexibility of Non-Gaussian Processes Through Shrinkage Priors”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/22-BA1342SUPPB>. 1236, 1239
- Deschamps, P. J. (2012). “Bayesian Estimation of Generalized Hyperbolic Skewed Student GARCH Models.” *Comput. Stat. Data Anal.*, 56(11): 3035–3054. URL <https://doi.org/10.1016/j.csda.2011.10.021> MR2943879. doi: <https://doi.org/10.1016/j.csda.2011.10.021>. 1224
- Dhull, M. S. and Kumar, A. (2021). “Normal inverse Gaussian autoregressive model using EM algorithm.” *International Journal of Advances in Engineering Sciences and Applied Mathematics*. MR4330355. doi: <https://doi.org/10.1007/s12572-021-00303-y>. 1224
- Diggle, P. J., Sousa, I., and Asar, Ö. (2014). “Real-time monitoring of progression towards renal failure in primary care patients.” *Biostatistics*, 16(3): 522–536. URL <https://doi.org/10.1093/biostatistics/kxu053> MR3365444. doi: <https://doi.org/10.1093/biostatistics/kxu053>. 1235
- Gabry, J. and Češnovar, R. (2021). *cmdstanr: R Interface to CmdStan*. URL <https://mc-stan.org/cmdstanr> 1237
- Gelman, A., Simpson, D., and Betancourt, M. (2017). “The prior can often only be understood in the context of the likelihood.” *Entropy*, 19(10): 555. 1232
- Ghasami, S., Khodadadi, Z., and Maleki, M. (2020). “Autoregressive processes with generalized hyperbolic innovations.” *Communications in Statistics - Simulation and Computation*, 49(12): 3080–3092. MR4186551. doi: <https://doi.org/10.1080/03610918.2018.1535066>. 1224
- Gneiting, T. and Raftery, A. E. (2007). “Strictly proper scoring rules, prediction, and estimation.” *Journal of the American Statistical Association*, 102(477): 359–378. MR2345548. doi: <https://doi.org/10.1198/016214506000001437>. 1225, 1241
- Hammerstein, E. A. v. (2016). “Tail behaviour and tail dependence of generalized hyper-

- bolic distributions.” In *Advanced modelling in mathematical finance*, 3–40. Springer. MR3630531. doi: https://doi.org/10.1007/978-3-319-45875-5_1. 1228
- Karlsson, S., Mazur, S., and Nguyen, H. (2021). “Vector autoregression models with skewness and heavy tails.” *arXiv preprint arXiv:2105.11182*. 1224
- Ken-Iti, S. (1999). *Lévy processes and infinitely divisible distributions*. Cambridge university press. 1223
- Lindgren, F., Rue, H., and Lindström, J. (2011). “An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4): 423–498. MR2853727. doi: <https://doi.org/10.1111/j.1467-9868.2011.00777.x>. 1234
- Matérn, B. (1960). *Spatial variation: Stochastic models and their application to some problems in forest surveys and other sampling investigations*. Statens Skogsforskningsinstitut, Stockholm. Meddelanden Fran Statens Skogsforskningsinstitut, Band 49, Nr. 5. MR0169346. 1234
- Nakajima, J. and Omori, Y. (2012). “Stochastic volatility model with leverage and asymmetrically heavy-tailed error using GH skew Student’s t-distribution.” *Computational Statistics & Data Analysis*, 56(11): 3690–3704. URL <https://www.sciencedirect.com/science/article/pii/S0167947310002859> MR2943921. doi: <https://doi.org/10.1016/j.csda.2010.07.012>. 1224
- Niekerk, J. and Rue, H. (2021). “Skewed Probit Regression — Identifiability, Contraction, and Reformulation, volume = 19.” *REVSTAT*, (1): 1–22. MR4241745. 1227
- Paolella, M. S. (2007). *Intermediate probability: A computational approach*. John Wiley & Sons. 1227
- Prause, K. et al. (1999). “The generalized hyperbolic model: Estimation, financial derivatives, and risk measures.” Ph.D. thesis, Citeseer. 1227
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC press. MR2130347. doi: <https://doi.org/10.1201/9780203492024>. 1229
- Rue, H. and Martino, S. (2007). “Approximate Bayesian inference for hierarchical Gaussian Markov random field models.” *Journal of statistical planning and inference*, 137(10): 3177–3192. URL <https://www.r-inla.org/> MR2365120. doi: <https://doi.org/10.1016/j.jspi.2006.07.016>. 1240, 1242
- Simpson, D., 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*, 1–28. MR3634300. doi: <https://doi.org/10.1214/16-STS576>. 1224, 1225, 1227, 1230, 1231, 1232, 1238
- Stan Development Team (2020). “Stan Modeling Language Users Guide and Reference Manual, 2.28.” URL <http://mc-stan.org/> 1225, 1235
- Vehtari, A., Gelman, A., and Gabry, J. (2017). “Practical Bayesian model evaluation us-

- ing leave-one-out cross-validation and WAIC.” *Statistics and computing*, 27(5): 1413–1432. MR3647105. doi: <https://doi.org/10.1007/s11222-016-9696-4>. 1240
- Walder, A. and Hanks, E. M. (2020). “Bayesian analysis of spatial generalized linear mixed models with Laplace moving average random fields.” *Computational Statistics & Data Analysis*, 144: 106861. MR4024672. doi: <https://doi.org/10.1016/j.csda.2019.106861>. 1224, 1234
- Wallin, J. and Bolin, D. (2015). “Geostatistical modelling using non-Gaussian Matérn fields.” *Scandinavian Journal of Statistics*, 42(3): 872–890. MR3391697. doi: <https://doi.org/10.1111/sjos.12141>. 1224, 1234, 1239
- Whittle, P. (1963). “Stochastic processes in several dimensions.” *Bulletin of the International Statistical Institute*, 40(2): 974–994. MR0173287. 1234
- Xie, F.-C. and Shen, Y.-Y. (2021). “Bayesian estimation for stochastic volatility model with jumps, leverage effect and generalized hyperbolic skew Student’s t-distribution.” *Communications in Statistics - Simulation and Computation*, 0(0): 1–18. 1224
- Zhu, B. and Dunson, D. B. (2017). “Bayesian Functional Data Modeling for Heterogeneous Volatility.” *Bayesian Analysis*, 12(2): 335–350. URL <https://doi.org/10.1214/16-BA1004> MR3620736. doi: <https://doi.org/10.1214/16-BA1004>. 1235