



# Skew Gaussian processes for classification

Alessio Benavoli<sup>1</sup> · Dario Azzimonti<sup>2</sup> · Dario Piga<sup>2</sup>

Received: 9 January 2020 / Revised: 5 May 2020 / Accepted: 12 August 2020 / Published online: 4 September 2020  
© The Author(s) 2020

## Abstract

Gaussian processes (GPs) are distributions over functions, which provide a Bayesian non-parametric approach to regression and classification. In spite of their success, GPs have limited use in some applications, for example, in some cases a symmetric distribution with respect to its mean is an unreasonable model. This implies, for instance, that the mean and the median coincide, while the mean and median in an asymmetric (skewed) distribution can be different numbers. In this paper, we propose skew-Gaussian processes (SkewGPs) as a non-parametric prior over functions. A SkewGP extends the multivariate *unified skew-normal* distribution over finite dimensional vectors to a stochastic processes. The SkewGP class of distributions includes GPs and, therefore, SkewGPs inherit all good properties of GPs and increase their flexibility by allowing asymmetry in the probabilistic model. By exploiting the fact that SkewGP and probit likelihood are conjugate model, we derive closed form expressions for the marginal likelihood and predictive distribution of this new nonparametric classifier. We verify empirically that the proposed SkewGP classifier provides a better performance than a GP classifier based on either Laplace's method or expectation propagation.

**Keywords** Skew Gaussian Process · Nonparametric · Classifier · Probit · Conjugate · Skew

## 1 Introduction

Gaussian processes (GPs) extend multivariate Gaussian distributions over finite dimensional vectors to infinite dimensionality. Specifically, a GP defines a distribution over functions, that is each draw from a Gaussian process is a function. Therefore, GPs provide a principled, practical, and probabilistic approach to nonparametric regression and classification and they have successfully been applied to different domains (Rasmussen and Williams 2006).

---

Editors: Ira Assent, Carlotta Domeniconi, Aristides Gionis, Eyke Hüllermeier.

---

✉ Alessio Benavoli  
alessio.benavoli@ul.ie

<sup>1</sup> Department of Computer Science and Information Systems, University of Limerick, Limerick, Ireland

<sup>2</sup> Dalle Molle Institute for Artificial Intelligence Research (IDSIA) - USI/SUPSI, Manno, Switzerland

GPs have several desirable mathematical properties. The most appealing one is that, for regression with Gaussian noise, the prior distribution is *conjugate* for the likelihood function. Therefore the Bayesian update step is analytic, as is computing the predictive distribution for the function behavior at unknown locations. In spite of their success, GPs have several known shortcomings.

First, the Gaussian distribution is not a “heavy-tailed” distribution, and so it is not robust to extreme outliers. Alternative to GPs have been proposed of which the most notable example is represented by the class of elliptical processes (Fang 2018), such as Student-t processes (O’Hagan 1991; Zhang et al. 2007), where any collection of function values has a desired elliptical distribution, with a covariance matrix built using a kernel.

Second, the Gaussian distribution is symmetric with respect to its mean. This implies, for instance, that its mean and median coincide, while the mean and median in an asymmetric (skewed) distribution can be different numbers. This constraint limits GPs’ flexibility and affects the coverage of their credible intervals (regions)—especially when considering that symmetry must hold for all components of the (latent) function and that, as for instance discussed by Kuss and Rasmussen (2005); Nickisch and Rasmussen (2008), the exact posterior of a GP classifier is skewed.

To overcome this second limitation, in this paper, we propose *skew-Gaussian processes* (SkewGPs) as a non-parametric prior over functions. A SkewGP extends the multivariate *unified skew-normal* distribution defined over finite dimensional vectors to a stochastic process, i.e. a distribution over infinite dimensional objects. A SkewGP is completely defined by a location function, a scale function and three additional parameters that depend on a latent dimension: a skewness function, a truncation vector and a covariance matrix. It is worth noting that a SkewGP reduces to a GP when the latent variables have dimension zero. Therefore, SkewGPs inherit all good properties of GPs and increase their flexibility by allowing asymmetry in the probabilistic model.

We focus on applying this new nonparametric model to a *classification* problem. In the case of parametric models, Durante (2019) shows that the posterior distribution of a probit likelihood and Gaussian prior has a unified skew-normal distribution. Such a novel result allowed the author to efficiently compute full posterior inferences for Bayesian logistic regression (for small datasets  $n \approx 100$ ). Moreover the author showed that the unified skew-normal distribution is a *conjugate prior* for the probit likelihood (without using this prior model for data analysis).

Here we extend this result to the nonparametric case, we derive a semi-analytical expression for the posterior distribution of the latent function and predictive probabilities for SkewGPs. The term *semi-analytical* is adopted to indicate that posterior inferences require the computation of the cumulative distribution function of a multivariate Gaussian distribution (i.e., the computation of Gaussian orthant probabilities). By using a new formulation (Gessner et al. 2019) of elliptical slice sampling (Murray et al. 2010), *lin-ess*, which permits efficient sampling from a linearly constrained (e.g., orthant) Gaussian domain, we show that we can compute efficiently posterior inferences for SkewGP binary classifiers. *Lin-ess* is a special case of elliptical slice sampling that leverages the analytic tractability of intersections of ellipses and hyperplanes to speed up the elliptical slice algorithm. In particular, this guarantees rejection-free sampling and it is therefore also highly parallelizable.

The main contributions of this paper are

1. we propose a new class of stochastic processes called skew-Gaussian processes (SkewGP) which generalize GP models;
2. we show that a SkewGP prior process is conjugate for the probit likelihood thus deriving for the first time the posterior distribution of a GP classifier in an analytic form;
3. we derive an efficient way to learn the hyperparameters of SkewGP and compute Monte Carlo predictions using *lin-ess*, showing that our model has similar bottleneck computational complexity of GPs;
4. we evaluate the proposed SkewGP classifier against state-of-the-art implementations of the GP classifier which approximate the posterior with the Laplace method or with Expectation propagation;
5. we show on a small image classification dataset that a SkewGP prior can lead to better uncertainty quantification than a GP prior.

## 2 Background

The skew normal distribution is a continuous probability distribution that generalises the normal distribution to allow for non-zero skewness. The probability density function (PDF) of the univariate skew-normal distribution with location  $\xi \in \mathbb{R}$ , scale  $\sigma > 0$  and skew parameter  $\alpha \in \mathbb{R}$  is given by O’Hagan and Leonard (1976):

$$p(z) = \frac{2}{\sigma} \phi\left(\frac{z-\xi}{\sigma}\right) \Phi\left(\alpha\left(\frac{z-\xi}{\sigma}\right)\right), \quad z \in \mathbb{R}$$


where  $\phi$  and  $\Phi$  are the PDF and, respectively, cumulative distribution function (CDF) of the standard univariate normal distribution. This distribution has been generalised in several ways, see Azzalini (2013) for a review. In particular, Arellano and Azzalini (2006) provided a unification of the above generalizations within a single and tractable multivariate *unified skew-normal* distribution that satisfies closure properties for marginals and conditionals and allows more flexibility due the introduction of additional parameters.

### 2.1 Unified skew-normal distribution

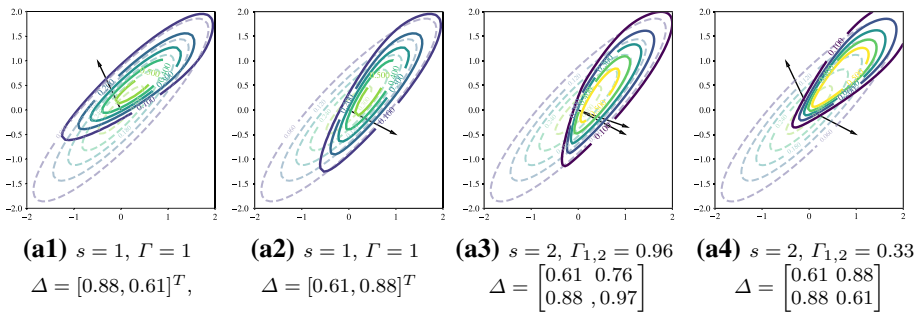
A vector  $z \in \mathbb{R}^p$  is said to have a multivariate unified skew-normal distribution with latent skewness dimension  $s$ ,  $z \sim \text{SUN}_{p,s}(\xi, \Omega, \Delta, \gamma, \Gamma)$ , if its probability density function (Azzalini 2013, Ch.7) is:

$$p(z) = \phi_p(z - \xi; \Omega) \frac{\Phi_s(\gamma + \Delta^T \bar{\Omega}^{-1} D_{\Omega}^{-1}(z - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta)}{\Phi_s(\gamma; \Gamma)} \quad (1)$$

where  $\phi_p(z - \xi; \Omega)$  represents the PDF of a multivariate normal distribution with mean  $\xi \in \mathbb{R}^p$  and covariance  $\Omega = D_{\Omega} \bar{\Omega} D_{\Omega} \in \mathbb{R}^{p \times p}$ , with  $\bar{\Omega}$  being a correlation matrix and  $D_{\Omega}$  a diagonal matrix containing the square root of the diagonal elements in  $\Omega$ . The notation  $\Phi_s(\mathbf{a}; M)$  denotes the CDF of  $N_s(0, M)$  evaluated at  $\mathbf{a} \in \mathbb{R}^s$ . The parameters  $\gamma \in \mathbb{R}^s$ ,  $\Gamma \in \mathbb{R}^{s \times s}$ ,  $\Delta^{p \times s}$  of the SUN distribution are related to a latent variable that controls the skewness, in particular  $\Delta$  is called Skewness matrix. The PDF (1) is well-defined provided that the matrix



**Fig. 1** Density plots for  $SUN_{1,s}(0, 1, \Delta, \gamma, \Gamma)$ . For all plots  $\Gamma$  is a correlation matrix,  $\gamma = 0$ , dashed lines are the contour plots of  $y \sim N_1(0, 1)$



**Fig. 2** Contour density plots for four unified skew-normal. For all plots  $p = 2$ ,  $\xi = [0, 0]^T$ ,  $\Omega$  and  $\Gamma$  are correlation matrices with  $\Omega_{1,2} = 0.88$ ,  $\gamma = 0$ , dashed lines are the contour plots of  $y \sim N_2(\xi, \Omega)$

$$M := \begin{bmatrix} \Gamma & \Delta^T \\ \Delta & \bar{\Omega} \end{bmatrix} \in \mathbb{R}^{(s+p) \times (s+p)} > 0, \tag{2}$$

i.e.,  $M$  is positive definite. Note that when  $\Delta = 0$ , (1) reduces to  $\phi_p(z - \xi; \Omega)$ . Moreover we assume that  $\Phi_0(\cdot) = 1$ , so that, for  $s = 0$ , (1) becomes a multivariate normal distribution.

The rôle of the latent dimension  $s$  can be briefly explained as follows. Consider now a random vector  $\begin{bmatrix} x_0 \\ x_1 \end{bmatrix} \sim N_{s+p}(0, M)$  with  $M$  as in (2) and define  $\mathbf{y}$  as the vector with distribution  $(x_1 | x_0 + \gamma > 0)$ , then it can be shown (Azzalini 2013, Ch. 7) that  $z = \xi + D_{\Omega} \mathbf{y} \sim SUN_{p,s}(\xi, \Omega, \Delta, \gamma, \Gamma)$ . This representation will be used in Sect. 5 to draw samples from the distribution. Figure 1 shows the density of a univariate SUN distribution with latent dimensions  $s = 1$  (a1) and  $s = 2$  (a2). The effect of a higher latent dimension can be better observed in bivariate SUN densities as shown in Fig. 2. The contours of the corresponding bivariate normal are dashed. We also plot the skewness directions given by  $\bar{\Omega}^{-1} \Delta$ .

The skew-normal family has several interesting properties, see Azzalini (2013, Ch.7) for details. Most notably, it is close under marginalization and affine transformations. Specifically, if we partition  $z = [z_1, z_2]^T$ , where  $z_1 \in \mathbb{R}^{p_1}$  and  $z_2 \in \mathbb{R}^{p_2}$  with  $p_1 + p_2 = p$ , then

$$z_1 \sim SUN_{p_1,s}(\xi_1, \Omega_{11}, \Delta_1, \gamma, \Gamma),$$

$$\text{with } \xi = \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix}, \quad \Delta = \begin{bmatrix} \Delta_1 \\ \Delta_2 \end{bmatrix}, \quad \Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}. \tag{3}$$

Moreover, (Azzalini 2013, Ch.7) the conditional distribution is a unified skew-normal, i.e.,  $(Z_2|Z_1 = z_1) \sim SUN_{p_2,s}(\xi_{2|1}, \Omega_{2|1}, \Delta_{2|1}, \gamma_{2|1}, \Gamma_{2|1})$ , where

$$\begin{aligned} \xi_{2|1} &:= \xi_2 + \Omega_{21}\Omega_{11}^{-1}(z_1 - \xi_1), & \Omega_{2|1} &:= \Omega_{22} - \Omega_{21}\Omega_{11}^{-1}\Omega_{12}, \\ \Delta_{2|1} &:= \Delta_2 - \bar{\Omega}_{21}\bar{\Omega}_{11}^{-1}\Delta_1, \\ \gamma_{2|1} &:= \gamma + \Delta_1^T\Omega_{11}^{-1}(z_1 - \xi_1), & \Gamma_{2|1} &:= \Gamma - \Delta_1^T\bar{\Omega}_{11}^{-1}\Delta_1, \end{aligned}$$

and  $\bar{\Omega}_{11}^{-1} := (\bar{\Omega}_{11})^{-1}$ .

### 3 Skew-Gaussian process

In this section, we define a Skew-Gaussian Process (SkewGP). Consider the functions  $\xi : \mathbb{R}^p \rightarrow \mathbb{R}$ , a location function,  $\Omega : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ , a scale function,  $\Delta : \mathbb{R}^p \rightarrow \mathbb{R}^s$ , the Skewness vector function, and  $\gamma \in \mathbb{R}^s, \Gamma \in \mathbb{R}^{s \times s}$ .

We say that a real function  $f : \mathbb{R}^p \rightarrow \mathbb{R}$  is distributed as a Skew-Gaussian process with latent dimension  $s$ , if, for any sequence of  $n$  points  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^p$ , the vector  $f(X) = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^T$  is skew-Gaussian distributed with parameters  $\gamma, \Gamma$ , location, scale and skewness matrices, respectively, given by

$$\xi(X) := \begin{bmatrix} \xi(\mathbf{x}_1) \\ \xi(\mathbf{x}_2) \\ \vdots \\ \xi(\mathbf{x}_n) \end{bmatrix}, \quad \Omega(X, X) := \begin{bmatrix} \Omega(\mathbf{x}_1, \mathbf{x}_1) & \Omega(\mathbf{x}_1, \mathbf{x}_2) & \dots & \Omega(\mathbf{x}_1, \mathbf{x}_n) \\ \Omega(\mathbf{x}_2, \mathbf{x}_1) & \Omega(\mathbf{x}_2, \mathbf{x}_2) & \dots & \Omega(\mathbf{x}_2, \mathbf{x}_n) \\ \vdots & \vdots & \dots & \vdots \\ \Omega(\mathbf{x}_n, \mathbf{x}_1) & \Omega(\mathbf{x}_n, \mathbf{x}_2) & \dots & \Omega(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix}, \tag{4}$$

$$\Delta(X) := \begin{bmatrix} \Delta(\mathbf{x}_1) & \Delta(\mathbf{x}_2) & \dots & \Delta(\mathbf{x}_n) \end{bmatrix}.$$

The skew-Gaussian distribution above is well defined provided that the matrix

$$\begin{bmatrix} \Gamma & \Delta(X) \\ \Delta(X)^T & \Omega(X, X) \end{bmatrix}$$

is positive definite for all  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^p$  and for all  $n > 0$ . In this case we can write

$$f(\mathbf{x}) \sim \text{SkewGP}_s(\xi(\mathbf{x}), \Omega(\mathbf{x}, \mathbf{x}'), \Delta(\mathbf{x}, \mathbf{x}'), \gamma, \Gamma). \tag{5}$$

We detail how to select the parameters in Sect. 4, the proposition below shows that SkewGP is a proper stochastic process.

**Proposition 1** *The construction of a Skew-Gaussian process from a finite-dimensional distribution is well-posed.*

All the proofs are in ‘‘Appendices 1 and 2’’.

### 3.1 Binary classification

Consider the training data  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , where  $\mathbf{x}_i \in \mathbb{R}^p$  and  $y_i \in \{0, 1\}$ . We aim to build a nonparametric binary classifier. We first define a probabilistic model  $\mathcal{M}$  by assuming that  $f \sim \text{SkewGP}(\xi, \Omega, \Delta, \gamma, \Gamma)$  and considering a *probit* model for the likelihood:

$$\begin{aligned}
 p(\mathcal{D}|f) &= \prod_{i=1}^n \Phi(f(\mathbf{x}_i); 1)^{y_i} (1 - \Phi(f(\mathbf{x}_i); 1))^{1-y_i} = \prod_{i=1}^n \Phi((2y_i - 1)f(\mathbf{x}_i); 1) \\
 &= \Phi_n(Wf(X); I_n),
 \end{aligned}
 \tag{6}$$

where  $W = \text{diag}(2y_1 - 1, \dots, 2y_n - 1)$ . A SkewGP prior combined with a probit likelihood gives rise to a posterior SkewGP over functions, this because skew-Gaussian distributions are conjugate priors for probit models. In the finite dimensional parametric case, this property was shown by Durante (2019), hereafter we extend it to the nonparametric one.

**Theorem 1** *The posterior of  $f(X)$  is a skew-Gaussian distribution, i.e.*

$$p(f(X)|\mathcal{D}) = \text{SUN}_{n,s+n}(\tilde{\xi}, \tilde{\Omega}, \tilde{\Delta}, \tilde{\gamma}, \tilde{\Gamma})
 \tag{7}$$

$$\tilde{\xi} = \xi, \quad \tilde{\Omega} = \Omega,
 \tag{8}$$

$$\tilde{\Delta} = [\Delta, \Omega D_\Omega W^T],
 \tag{9}$$

$$\tilde{\gamma} = [\gamma, W\xi]^T,
 \tag{10}$$

$$\tilde{\Gamma} = \begin{bmatrix} \Gamma & \Delta^T D_\Omega W^T \\ W D_\Omega \Delta & (W \Omega W^T + I_n) \end{bmatrix},
 \tag{11}$$

where, for simplicity of notation, we have denoted  $\xi(X), \Omega(X, X), \Delta(X)$  as  $\xi, \Omega, \Delta$  and  $\Omega = D_\Omega \tilde{\Omega} D_\Omega$ .

From Theorem 1 we can immediately derive the following result.

**Corollary 1** *The marginal likelihood of the observations  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$  given the probabilistic model  $\mathcal{M}$ , that is the prior (5) and likelihood (6), is*

$$p(\mathcal{D}|\mathcal{M}) = \frac{\Phi_{s+n}(\tilde{\gamma}; \tilde{\Gamma})}{\Phi_s(\gamma; \Gamma)},
 \tag{12}$$

with  $\tilde{\gamma}, \tilde{\Gamma}$  defined in Theorem 1.

In classification, based on the training data  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , and given test inputs  $\mathbf{x}^*$ , we aim to predict the probability that  $y^* = 1$ .

**Corollary 2** *The posterior predictive probability of  $y^* = 1$  given the test input  $\mathbf{x}^* \in \mathbb{R}^p$  and the training data  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$  is*

$$p(y^* = 1 | \mathbf{x}^*, X, y) = \frac{\Phi_{s+n+1}(\tilde{\gamma}^*; \tilde{\Gamma}^*)}{\Phi_{s+n}(\tilde{\gamma}; \tilde{\Gamma})}, \quad (13)$$

where  $\tilde{\gamma}^*, \tilde{\Gamma}^*$  are the corresponding matrices of the posterior computed for the augmented dataset  $\hat{X} = [X^T, \mathbf{x}^{*T}]^T$ ,  $\hat{y} = [y^T, 1/2]^T$ .

Note that the dummy value  $\frac{1}{2}$  in  $\hat{y}$  does not influence the value of  $p(y^* = 1 | \mathbf{x}^*, X, y)$  and it was chosen only for mathematical convenience, as it allows for marginalization over  $f(\mathbf{x}^*)$  and to derive the expression of  $\tilde{\gamma}^*, \tilde{\Gamma}^*$  similarly to the ones in Theorem 1.<sup>1</sup>

#### 4 Prior functions, parameters and hyperparameters

A SkewGP prior is completely defined by the location function  $\xi(\mathbf{x})$ , the scale function  $\Omega(\mathbf{x}, \mathbf{x}')$ , the latent dimension  $s \in \mathbb{N}$ , the skewness vector function  $\Delta(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^s$  and  $\gamma \in \mathbb{R}^s, \Gamma \in \mathbb{R}^{s \times s}$ . As it is common for GPs, we will take the location function  $\xi(\mathbf{x})$  to be zero, although this need not be done. Let  $K(\mathbf{x}, \mathbf{x}')$  be a positive definite covariance function (kernel) and let  $\Omega = K(X, X)$  be the covariance matrix obtained by applying  $K(\mathbf{x}, \mathbf{x}')$  elementwise to the training data  $X$ . In this paper, we propose the following way to define the location, scale and skewness functions of a SkewGP:

$$M = \begin{bmatrix} \Gamma & \Delta(X, R)^T \\ \Delta(X, R) & \tilde{\Omega}(X, X) \end{bmatrix} := \begin{bmatrix} L\bar{K}(R, R)L & L\bar{K}(R, X) \\ \bar{K}(X, R)L & \bar{K}(X, X) \end{bmatrix}, \quad (14)$$

with  $L \in \mathbb{R}^{s \times s}$  is a diagonal matrix whose elements  $L_{ii} \in \{-1, 1\}$  (a phase),  $R = [r_1, \dots, r_s]^T \in \mathbb{R}^{s \times p}$  is a vector of  $s$  pseudo-points and  $\bar{K}(\mathbf{x}, \mathbf{x}') = \frac{1}{\sigma^2} K(\mathbf{x}, \mathbf{x}')$  (for stationary kernels) with  $K(\mathbf{x}, \mathbf{x}')$  being the kernel function and  $\sigma^2$  the variance parameter of the kernel, e.g., for the RBF kernel

$$K(\mathbf{x}, \mathbf{x}') := \sigma^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right).$$

It can easily be proven that  $M > 0$  and, therefore, (2) holds. We select the parameters of the kernel  $\sigma, \ell$ , the locations  $r_i$  of the pseudo-points and the phase diagonal matrix  $L$  by maximizing the marginal likelihood. In particular we exploit the lower bound (16) explained in Sect. 5.

Similarly to the inducing points in the sparse approximation of GPs (Quiñonero-Candela and Rasmussen 2005; Snelson and Ghahramani 2006; Titsias 2009; Bauer et al. 2016), the points  $r_i$  can be viewed as a set of  $s$  latent variables. However, their rôle is completely different from that of the inducing points, they allow us to locally modulate the skewness of SkewGP. Figure 3 shows latent functions (in gray, second column) drawn from a SkewGP with latent dimension 2 and the result of squashing these sample functions through the probit logistic function (first column). In all cases, we have considered  $\xi(\mathbf{x}) = 0$  and a RBF kernel with  $\ell = 0.3$  and  $\sigma^2 = 1$ . The values of the other parameters of the SkewGP are reported at the top of the plots in the first column. The green line is the mean function

<sup>1</sup> Note in fact that, for  $y = 1/2$ , the likelihood  $\Phi((2y - 1)f(\mathbf{x}^*), 1) = 0.5$  and it does not depend on  $f(\mathbf{x}^*)$ , and so it is marginalised out.

**Fig. 3** The second column shows random latent functions (in gray) drawn from a SkewGP with latent dimension 2 and the first column shows the result of squashing these sample functions through the probit logistic function. The values of the parameters of the SkewGP are reported at the top of the plots in the first column. The green line is the mean function and the red dots represent the location of the  $s = 2$  latent pseudo-points

and the red dots represent the location of the  $s = 2$  latent pseudo-points. For large positive values of  $\gamma_1, \gamma_2$ , SkewGP is equivalent to a GP (plots (a1)–(a2)). At the decreasing of  $\gamma_i$ ,  $i = 1, 2$  (plots (b1)–(b2)), the mean shifts up and the mass of the distribution is concentrated on the top of the figure. By changing the phase (the sign of)  $L_{22}$  (plots (c1)–(c2)), the mean and the mass of the distribution shift down at the location of the second pseudo-observation  $r_2$ . We can magnifying this effect by decreasing both  $\gamma_i$  (plots (d1)–(d2)). It is also possible to introduce skewness without changing the mean (plots (e1)–(e2)). In this latter case,  $r_1 = r_2$  and the mass of the distribution is shifted up.

## 5 Computational complexity

Corollaries 1 and 2 provide two straightforward ways to compute the marginal likelihood and the predictive posterior probability however Eqs. (12) and (13) both require the accurate computation of  $\Phi_{s+n}$ . Quasi-randomized Monte Carlo methods (Genz 1992; Genz and Bretz 2009; Botev 2017) allows calculation of  $\Phi_{s+n}$  for small  $n$  (few hundreds observations). Therefore, these procedures are not in general suitable for medium and large  $n$  [apart from special cases Phinikettos and Gandy (2011), Genton et al. (2018), Azzimonti and Ginsbourger (2018)]. We overcome this issue with an effective use of sampling for the predictive posterior and a mini-batch approach to marginal likelihood.


### 5.1 Posterior predictive distribution

In order to compute the predictive distribution we generate samples from the posterior distribution at training points and then exploit the closure properties of the SUN distribution to obtain samples at test points. The following result from Azzalini (2013) allows us to draw independent samples from the posterior in (7):


$$\begin{aligned} f(X) &\sim \tilde{\xi} + D_{\Omega}(U_0 + \tilde{\Delta}\tilde{T}^{-1}U_1), \\ U_0 &\sim \mathcal{N}(0; \tilde{\Omega} - \tilde{\Delta}\tilde{T}^{-1}\tilde{\Delta}^T), \quad U_1 \sim \mathcal{T}_{\tilde{Y}}(0; \tilde{T}), \end{aligned} \quad (15)$$

where  $\mathcal{T}_{\tilde{Y}}(0; \tilde{T})$  is the pdf of a multivariate Gaussian distribution truncated component-wise below  $-\tilde{Y}$ . Equation (15) is a consequence of the additive representation of skew-normal distributions, see Azzalini (2013, Ch. 7.1.2 and 5.1.3) for more details. Note that sampling  $U_0$  can be achieved with standard methods, however using standard rejection sampling for the variable  $U_1$  would incur in exponentially growing sampling time as the dimension increases. A commonly used sampling technique for this type of problems is Elliptical Slice Sampling (*ess*) (Murray et al. 2010) which is a Markov Chain Monte Carlo algorithm that performs inference in models with Gaussian priors. This method looks for acceptable samples along elliptical slices and by doing so drastically reduces the number of rejected samples. Recently, Gessner et al. (2019) proposed an extension of *ess*, called linear elliptical slice sampling (*lin-ess*), for multivariate Gaussian distribution truncated on a region defined by linear constraints. In particular, this approach analytically derives the acceptable







(a1)




(a2)




(b1)




(b2)




(c1)




(c2)




(d1)



(d2)



(e1)



(e2)

regions on the elliptical slices used in *ess* and thus guarantees rejection-free sampling. This leads to a large speed up over *ess*, especially in high dimensions.

Given posterior samples at the training points it is possible to compute the predictive posterior at a new test points  $\mathbf{x}^*$  thanks to the following result.

**Theorem 2** *The posterior samples of the latent function computed at the test point  $\mathbf{x}^*$  can be obtained by sampling  $f(\mathbf{x}^*)$  from:*

$$\begin{aligned} f(\mathbf{x}^*) &\sim \text{SUN}_{1,s}(\xi_*, \Omega_*, \Delta_*, \gamma_*, \Gamma_*), \\ \xi_* &= \xi(\mathbf{x}^*) + K(\mathbf{x}^*, X)K(X, X)^{-1}(f(X) - \xi(X)), \\ \Omega_* &= K(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, X)K(X, X)^{-1}K(X, \mathbf{x}^*), \\ \Delta_* &= \Delta(\mathbf{x}^*) - \bar{K}(\mathbf{x}^*, X)\bar{K}(X, X)^{-1}\Delta(X), \\ \gamma_* &= \gamma + \Delta(X)^T \bar{K}(X, X)^{-1}D_\Omega(X, X)^{-1}(f(X) - \xi(X)), \\ \Gamma_* &= \Gamma - \Delta(X)^T \bar{K}^{-1}\Delta(X), \end{aligned}$$

with  $f(X)$  sampled from the posterior  $\text{SUN}_{n,s+n}(\tilde{\xi}, \tilde{\Omega}, \tilde{\Delta}, \tilde{\gamma}, \tilde{\Gamma})$  in Theorem 1,  $K$  is a kernel that defines the matrices  $\Gamma, \Delta, \Omega$  as in eq. (14) and where

$$\begin{bmatrix} \bar{K}(X, X) & \bar{K}(X, \mathbf{x}^*) \\ \bar{K}(\mathbf{x}^*, X) & \bar{K}(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} := D_\Omega^{-1} \begin{bmatrix} K(X, X) & K(X, \mathbf{x}^*) \\ K(\mathbf{x}^*, X) & K(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} D_\Omega^{-1}$$

and  $D_\Omega = \text{diag}[D_\Omega(X), D_\Omega(\mathbf{x}^*)]$  is a diagonal matrix containing the square root of the diagonal elements of the inner matrix in the r.h.s..

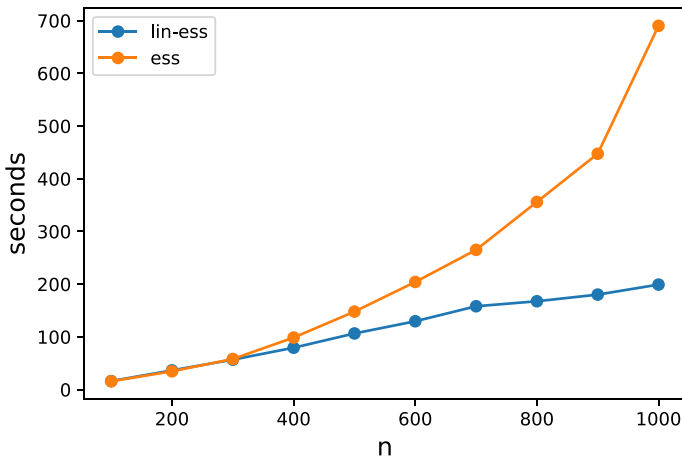
Observe that the computation of the predictive posterior requires the inversion of a  $n \times n$  matrix ( $\tilde{\Omega}_{11}$ ), which has complexity  $O(n^3)$  with storage demands of  $O(n^2)$ . SkewGPs then have a bottleneck in computational complexity similar to that of GPs. Moreover, note that sampling from  $\text{SUN}_{1,s}$  is extremely efficient when the latent dimension  $s$  is small (in the experiments we use  $s = 2$ ).

### 5.2 Marginal likelihood

As discussed in the previous section, in practical application of SkewGP, the (hyper-) parameters of the scale function  $\Omega(\mathbf{x}, \mathbf{x}')$ , of the skewness vector function  $\Delta(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^s$  and the parameters  $\gamma \in \mathbb{R}^s, \Gamma \in \mathbb{R}^{s \times s}$  have to be selected. As for GPs, we use Bayesian model selection to consistently set such parameters. This requires the maximization of the marginal likelihood with respect to these parameters and, therefore, it is essential to provide a fast and accurate way to evaluate the marginal likelihood. In this paper, we propose a simple approximation of the marginal likelihood that allows us to efficiently compute a lower bound.

**Proposition 2** *Consider the marginal likelihood  $p(\mathcal{D}|\mathcal{M})$  in Corollary 1, then it holds*

$$p(\mathcal{D}|\mathcal{M}) = \frac{\Phi_{s+n}(\tilde{\gamma}; \tilde{\Gamma})}{\Phi_s(\gamma; \Gamma)} \geq \frac{\sum_{i=1}^b \Phi_{s+|B_i|}(\tilde{\gamma}_{B_i}; \tilde{\Gamma}_{B_i}) - (b-1)}{\Phi_s(\gamma; \Gamma)}, \tag{16}$$



**Fig. 4** Computational time for sampling from the posterior of SkewGP0 using ess versus lin-ess on a standard laptop

where  $B_1, \dots, B_b$  denotes a partition of the training dataset into  $b$  random disjoint subsets,  $|B_i|$  denotes the number of observations in the  $i$ th element of the partition,  $\tilde{\gamma}_{B_i}, \tilde{\Gamma}_{B_i}$  are the parameters of the posterior computed using only the subset  $B_i$  of the data.

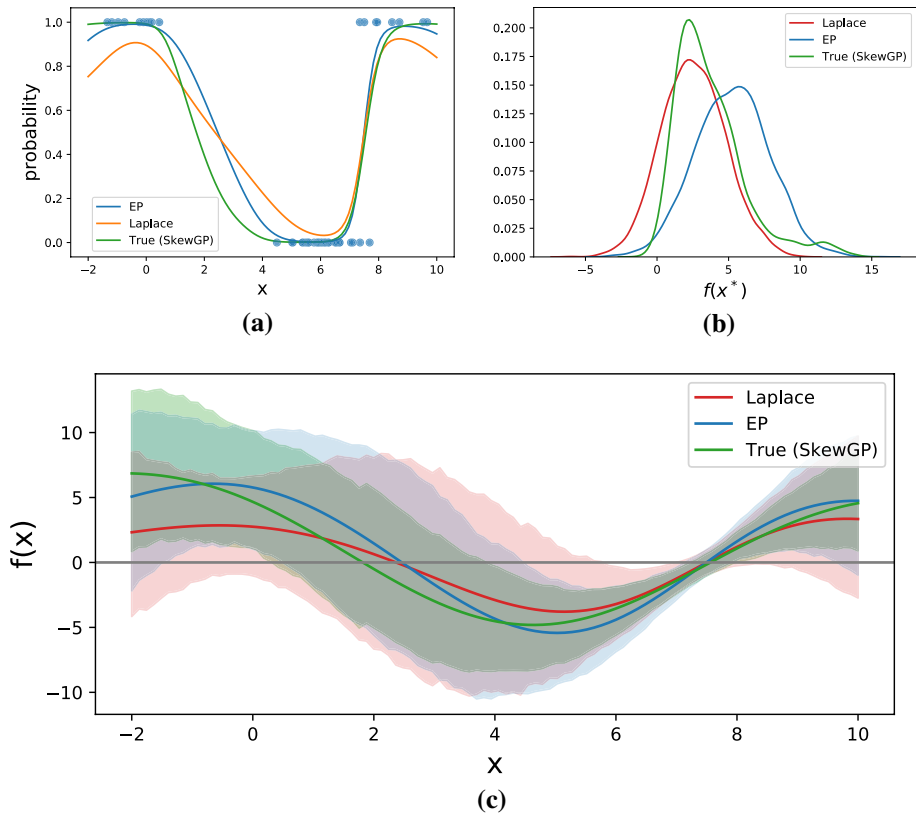
If the batch-size is low (in the experiments we have used  $|B_i| = 30$ ), then we can efficiently compute each term  $\Phi_{s+|B_i|}(\tilde{\gamma}_{B_i}; \tilde{\Gamma}_{B_i})$  by using a quasi-randomised Monte Carlo method. We can then optimize the hyper-parameter of SkewGP by maximizing the lower bound in (16).

### 5.3 Computational load and parallelization

To evaluate the computational load, we have generated artificial classification data using a probit likelihood model and drawing the latent function  $f(X) = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]$ , with  $\mathbf{x}_i \sim N(0, 1)$ , from a GP with zero mean and radial basis function kernel (lengthscale 0.5 and variance 2). We have then computed the full posterior latent function from Theorem 1, that is a SkewGP posterior. Figure 4 compares the CPU time for sampling 1000 instances of  $f(X)$  from the SkewGP posterior as a function of  $n$  for lin-ess versus a standard Elliptical Slice Sampler (ess) (5000 burn in).<sup>2</sup> It can be observed the computational advantage of lin-ess with respect to ess.

The average CPU time required to compute  $\Phi_{s+|B_i|}(\tilde{\gamma}_{B_i}; \tilde{\Gamma}_{B_i})$  with  $|B_i| = 30$ , using the randomized lattice routine with 5000 points (Genz 1992), is 0.5 seconds on a standard laptop. Since the above method is randomized, we use the *simultaneous perturbation*

<sup>2</sup> Sampling is performed according to (15) and, therefore, lin-ess and ess are applied to sample  $U_1$  from  $\mathcal{T}_{\tilde{\gamma}}(0; \tilde{\Gamma})$ . To increase the probability of acceptance for ess, we have replaced the indicator function that defines the truncation  $I_{u_1 > \tilde{\gamma}}$  with a logistic function  $\text{sigmoid}(80(u_1 - \tilde{\gamma}))$ . We have verified that, using 5000 samples for burn in, the posterior first moment of lin-ess and ess are close for all considered values of  $n$ .

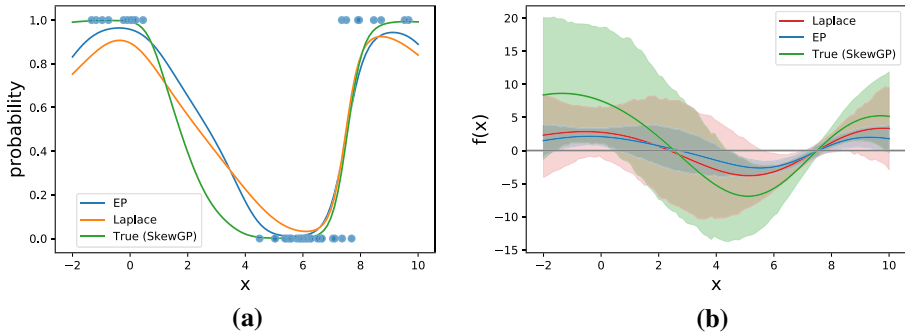


**Fig. 5** Plot **a** shows the training dataset and compares the true posterior mean probability of class 1, that is a SkewGP, versus the approximations computed via Laplace’s method and EP. Plot **c** shows the posterior mean latent function and corresponding 95% credible intervals for the three methods and Plot **b** reports the density-plot of the latent function posterior prediction at  $x^* = 0.42$

*stochastic approximation* algorithm (Spall 1998) for optimizing the maximum lower bound (16).<sup>3</sup>

Finally notice that, both the computation of the lower bound of the marginal likelihood and the sampling from the posterior via lin-ess are highly parallelizable. In fact, each term  $\Phi_{s+|B_i|}(\tilde{\gamma}_{B_i}; \tilde{\Gamma}_{B_i})$  can be computed independently as well as each sample in lin-ess can be sample independently (because lin-ess is rejection-free and, therefore, no “burn in” is necessary).

<sup>3</sup> The SkewGP classifier is implemented in Python, but we call Matlab to use existing implementations of both the simultaneous perturbation stochastic approximation algorithms and the randomized lattice routine. We plan to re-implement everything in Python and then release the source code for SkewGP.



**Fig. 6** Plot **a** shows the training dataset and compares the true posterior mean probability of class 1, that is a SkewGP, versus the approximations computed via Laplace's method and EP. Plot **b** shows the posterior mean latent function and corresponding 95% credible intervals for the three methods

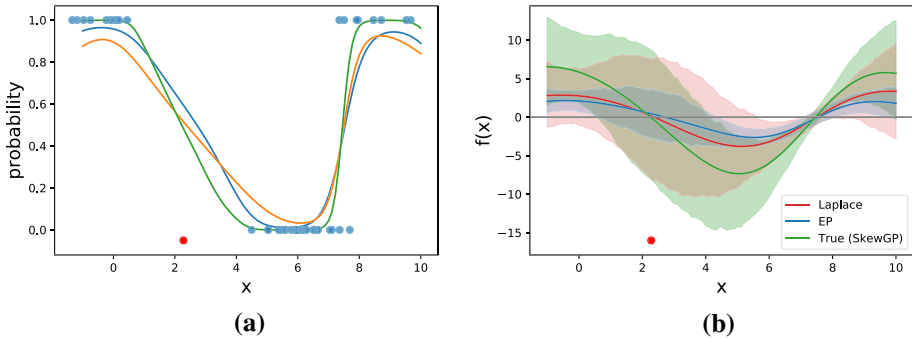
## 6 Properties of the posterior

In the above sections, we have shown how to compute the posterior distribution of a SkewGP process when the likelihood is a probit model. The full conjugacy of the model allows us to prove that the posterior is again a SkewGP process. This section provides more details on the properties of the posterior and compares it with two approximations. For GP classification, there are two main alternative approximation schemes for finding a Gaussian approximation to the posterior: the Laplace's method and the Expectation-Propagation (EP) method, see, e.g. Rasmussen and Williams (2006) chapter 3.

Figure 5 provides a one-dimensional illustration using a synthetic classification problem with 50 observations and scalar inputs taken from (Kuss and Rasmussen 2005). Figure 5a shows the dataset and the predictive posterior probability for the Laplace and EP approximations. Moreover, by using a SkewGP prior with latent dimension  $s = 0$  (that coincides with a GP prior), we have computed the exact SkewGP predictive posterior probability. Therefore, all three methods have the same prior: a GP with zero mean and RBF covariance function (the lengthscale and variance of the kernel are the same for all the three methods and have been set equal to the values that maximise the Laplace's approximation to the marginal likelihood). Figure 5c shows the posterior mean latent function and corresponding 95% credible intervals. It is evident that the true posterior (SkewGP) of the latent function is skew (see for instance for  $x \in [0, 2]$  and the slice plot in Fig. 5b)). Laplace's approximation peaks at the posterior mode, but places too much mass over positive values of  $f$ . The EP approximation aims to match the first two moments of the posterior and, therefore, usually obtains a better coverage of the posterior mass. That is why EP is usually the method of choice for approximate inference in GP classification.

Figure 6 shows the true posterior and the two approximations for the same dataset, but now the lengthscale and variance of the kernel are the optimal values for the three methods. It is evident that the skewness of the posterior provides a better model fit to the data.

Figure 7 shows the posteriors corresponding to a prior SkewGP process with latent-dimension  $s = 2$ . The red dot denotes the optimal location of the pseudo-point  $r_1$ , while  $r_2 = 13.5$  (their initial location were 5.8 and, respectively, 6). The additional degrees of freedom of the SkewGP prior process gave a much more satisfactory answer than that obtained from a GP prior model. By comparing Figs. 6 and 7, it can be noticed that the skew-point allows us to locally modulate the skewness. Moreover, the additional degrees



**Fig. 7** Plot **a** shows the training dataset and compares the true posterior mean probability of class 1, that is a SkewGP2, versus the approximations computed via Laplace’s method and EP. Plot **b** shows the posterior mean latent function and corresponding 95% credible intervals for the three methods. The red dot denotes the optimal location of the pseudo-point  $r_1$

of freedom do not lead to overfitting, even with small data, as highlighted by the optimized location of  $r_2$  (far away) that has not effect on the skewness of the posterior SkewGP.

### 7 Results

We have evaluated the proposed SkewGP classifier on a number of benchmark classification datasets and compared its classification accuracy with the accuracy of a Gaussian process classifier that uses either Laplace’s method (GP-L) or Expectation Propagation (GP-EP) for approximate Bayesian inference. For GP-L and GP-EP, we have used the implementation available in GPy (GPy, since 2012).


#### 7.1 Penn machine learning benchmarks

From the Penn Machine Learning Benchmarks (Olson et al. 2017), we have selected 124 datasets (number of features up to 500). Since this pool includes non-binary class datasets, we have defined a binary sub-problem by considering the first (class 0) and second (class 1) class. The resulting binarised subset of datasets includes datasets with number of instances between 100 and 7400. We have scaled the inputs to zero mean and unit standard deviation and used, as performance measure, the average information in bits of the predictions about the test targets (Kuss and Rasmussen 2005):

$$I(p_i^*, y_i) = \frac{y_i + 1}{2} \log_2(p_i^*) + \frac{1 - y_i}{2} \log_2(1 - p_i^*) + 1,$$

where  $p_i^*$  is the predicted probability for class 1. This score equals 1 bit if the true label is predicted with absolute certainty, 0 bits for random guessing and takes negative values if the prediction gives higher probability to the wrong class. We have assessed the above performance measure for the three classifiers by using 5-fold cross-validation.

While we could use any kernel for GP-L, GP-EP and SkewGP, in this experiment we have chosen the RBF kernel with a lengthscale for each dimension. Figure 8 contrasts GP-L and GP-EP with SkewGP0 (SkewGP with  $s = 0$ ) and SkewGP2 (SkewGP with

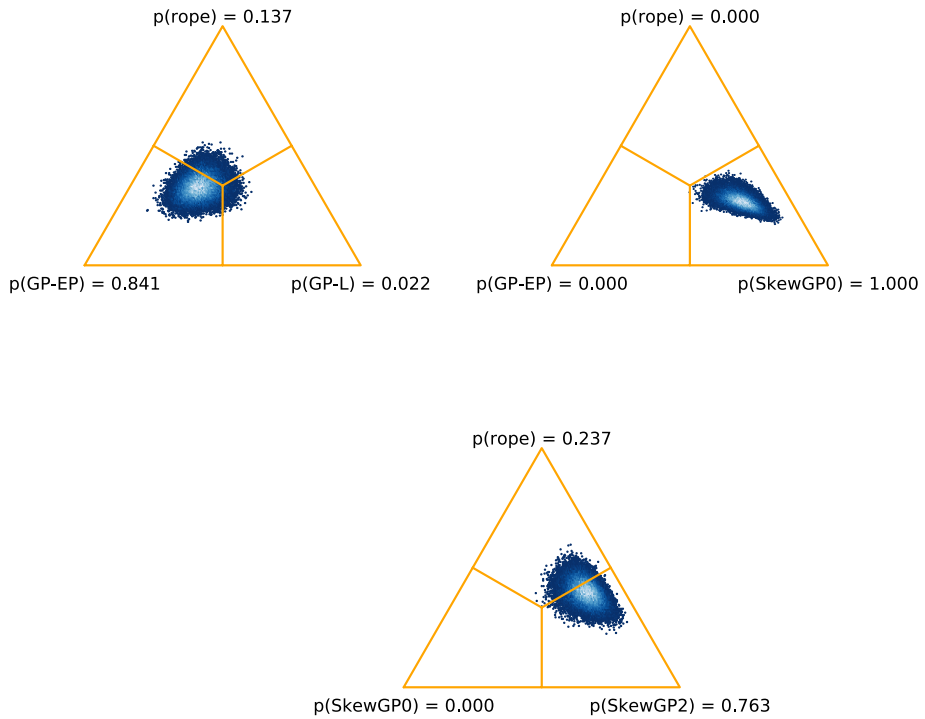


**Fig. 8** Average information score on 124 datasets of the Penn Machine Learning Benchmark dataset

$s = 2$ ). We selected  $s = 2$  because we decided to use the same dimension for all datasets and, since there are several datasets where the ratio between the number of features and the number of instances is high, a latent dimension  $s > 2$  leads to a number of parameters that exceeds the number of instances affecting the convergence of the maximization of the marginal likelihood. The proposed SkewGP2 and SkewGP0 outperform the other two models for most data sets. The average information score of SkewGP2 is 0.573 (average accuracy 0.904), SkewGP0 is 0.557 (acc. 0.882), GP-EP is 0.542 (acc. 0.859) and GP-L is 0.512 (acc. 0.863).

This claim is supported by a statistical analysis. We have compared the three classifiers using the (nonparametric) Bayesian signed-rank test (Benavoli et al. 2014, 2017). This test declares two classifiers practically equivalent when the difference of average information is less than 0.01 (1%). The interval  $[-0.01, 0.01]$  thus defines a region of practical equivalence (rope) for classifiers. The test returns the posterior probability of the vector  $[p(C_1 > C_2), p(C_1 \approx C_2), p(C_1 < C_2)]$  and, therefore, this posterior can be visualised in the probability simplex (Fig. 9). For the comparison GP-L versus GP-EP, as expected it can be seen that GP-EP is better than GP-L.<sup>4</sup> Conversely, for GP-EP versus SkewGP0, 100% of the posterior mass is in the region in favor of SkewGP0, which is the region at the right bottom of the triangle. This confirms that SkewGP0 is practically significantly better than GP-L and GP-EP. The comparison SkewGP2 versus SkewGP0 shows that SkewGP2 has

<sup>4</sup> It is well known that GP-EP usually achieves a more calibrated estimate of the class probability. Laplace’s method gives over-conservative predictive probabilities (Kuss and Rasmussen 2005).



**Fig. 9** Bayesian Wilcoxon signed-rank test


surely an average information score that is not worse than that of SkewGP0 and better with probability about 0.76.

The difference between GP-L and GP-EP, and SkewGP is that the posterior of SkewGP can be skewed. Therefore, we expect SkewGP to outperform GP-L and GP-EP on the datasets for which the posterior is far from Normal (e.g., highly skewed). To verify that we have computed the sample skewness statistics (SS) for each test input  $\mathbf{x}_i^*$ :

$$SS(\mathbf{x}_i^*) = \frac{E[(f(\mathbf{x}_i^*) - \mu)^3]}{(E[(f(\mathbf{x}_i^*) - \mu)^2])^{3/2}}$$

with  $\mu = E[f(\mathbf{x}_i^*)]$  and the expectation  $E[\cdot]$  can be approximated using the posterior samples drawn as in Theorem 2. Note that  $SS(\mathbf{x}_i^*) = 0$  for symmetric distributions. Figure 10(left) shows, for each of the 124 datasets, the difference between the average information score of SkewGP0 and GP-EP in the y-axis, and  $\max_{\mathbf{x}_i^*} SS(\mathbf{x}_i^*)$  for SkewGP0 in the x-axis. We used a regression tree (green line) to detect structural changes in the mean of these data points. It is evident that, for large values of the maximum skewness statistics, SkewGP0 outperforms GP-EP (the average difference is positive). Figure 10(right) reports a similar plot for SkewGP2 and the difference is even more evident. This confirms that SkewGP on average outperforms GP-EP in those datasets where the posterior is skewed and has a similar performance otherwise.





**Fig. 10** Maximum skewness statistics versus difference in average information score. For visualisation purpose only, we bounded the y-axis to  $[-0.11, 0.16]$

**Table 1** Image classification

Class	RBF kernel		NN kernel		Time	
	Accuracy		Accuracy		GP-EP	SkewGP2
	GP-EP	SkewGP2	GP-EP	SkewGP2		
0	0.945	<b>0.955</b>	0.956	<b>0.958</b>	80.7	67.8
1	0.988	<b>0.990</b>	0.988	<b>0.991</b>	125.7	68.0
2	0.923	<b>0.934</b>	0.943	0.943	130.9	70.7
3	0.951	<b>0.961</b>	0.953	<b>0.961</b>	79.2	56.9
4	0.932	<b>0.946</b>	<b>0.946</b>	0.944	121.1	49.3
5	0.967	<b>0.980</b>	0.975	<b>0.983</b>	110.5	53.0
6	0.917	<b>0.922</b>	0.924	<b>0.933</b>	115.7	65.1
7	0.970	<b>0.974</b>	0.981	<b>0.983</b>	113.6	66.3
8	0.969	<b>0.979</b>	0.980	<b>0.981</b>	121.5	51.4
9	0.980	<b>0.982</b>	0.983	<b>0.986</b>	137.0	55.0
One-vs-rest	0.785	<b>0.821</b>	0.823	<b>0.844</b>		

Bold denotes the highest accuracy in the corresponding classification task

### 7.2 Image classification

We have also considered an image classification task: *Fashion MNIST* dataset (each image is  $28 \times 28 = 784$  pixels and there are 10 classes: 0 T-shirt/top, 1 Trouser, 2 Pullover, 3 Dress, 4 Coat, 5 Sandal, 6 Shirt, 7 Sneaker, 8 Bag, 9 Ankle boot). We randomly pooled 10000 images from the dataset and divided them into two sets, with 5000 cases for training and 5000 for testing. For each one of the 10 classes, we have defined a binary classification sub-problem by considering one class against all the other classes. We have compared GP-EP and SkewGP2, that is a SkewGP with latent dimension  $s = 2$  (for the same reason outlined in the previous section). We have initialised  $r_i$  by taking 2 random samples from

the training data. We have also considered two different kernels: RBF and the Neural Network kernel (Williams 1998). Table 1 reports the accuracy for each of the 10 binary classification sub-problems. For the RBF kernel, it can be noticed that SkewGP2 outperforms GP-EP in all sub-problems. For the NN kernel, the differences between the two models are less substantial (due to the higher performance of the NN kernel on this dataset) but in any case in favor of SkewGP2. We have also reported, for both the models, the computational time<sup>5</sup> (in minutes) needed to optimize the hyperparameters, to compute the posterior and to compute the predictions for all instances in the test set. This shows that SkewGP2 is also faster than GP-EP.<sup>6</sup> The last row reports the accuracy on the original multi-class classification problem obtained by using the one-vs-rest heuristic, with the only goal of showing that the more accurate estimate of the probability by SkewGP leads also to an increase in accuracy for one-vs-rest. A multi-class Laplace's approximation for GP classification was developed by Williams and Barber (1998) and other implementations are for instance discussed by Hernández-Lobato et al. (2011) and Chai (2012), we plan to address multi-class classification in future work.

Our goal is assessing the accuracy but also the quality of the probabilistic predictions. Figure 11, plot (a1), shows, for the RBF kernel case and for each instance in the test set of the binary sub-problem relative to class 8 versus rest, the value of the mean predicted probability of class *rest* for SkewGP2 (x-axis) and GP-EP (y-axis). Each instance is represented as a blue point. The red points highlight the instances that were misclassified by GP-EP. Figure (a2) shows the same plot, but the red points are now the instances that were misclassified by SkewGP2. By comparing (a1) and (a2), it is clear that SkewGP2 provides a higher quality of the probabilistic predictions.

SkewGP2 also returns a better estimate of its own uncertainty, as shown in plots (b1) versus (b2). For each instance in the test set and for each sample from the posterior, we have computed the predicted class (the class that has probability greater than 0.5). For each test instance, we have then computed the standard deviation of all these predictions and used it to color the scatter plot of the mean predicted probability. In this way, we have a visual representation of first order (mean predicted probability) and second order (standard deviation of the predictions) uncertainty. Fig. 11(b1) is relative to GP-EP and shows that GP-EP confidence is low only for the instances whose mean predicted probability is close to 0.5. This is not reflected in the value of the mean predicted probability for the misclassified instances (compare plot (a1) and (b1)). We have also computed the histogram of the standard deviation of the predictions for those instances that were misclassified by GP-EP in Figure (c1). Note that, the peak of the histogram corresponds to very low standard deviation, that means GP-EP has misclassified instances that have low second order uncertainty. This implies that the model is overestimating its confidence. Conversely, the second order uncertainty of SkewGP2 is clearly consistent, see plot (a2) and (b2), and in particular the histogram in (c2)—the peak is in correspondence of high values of the standard deviation of the predictions. In other words, SkewGP2 has mainly misclassified instances with high second order uncertainty, that is what we expect from a calibrated probabilistic model. We have reported additional examples of the better calibration of SkewGP2 for the MNIST and German-road sign dataset in "Appendices 1 and 2".

<sup>5</sup> More precisely, the table reports the average computational time for the RBF and NN kernel case.

<sup>6</sup> This is due to both the efficiency of lin-ess and the batch approximation of the marginal likelihood.

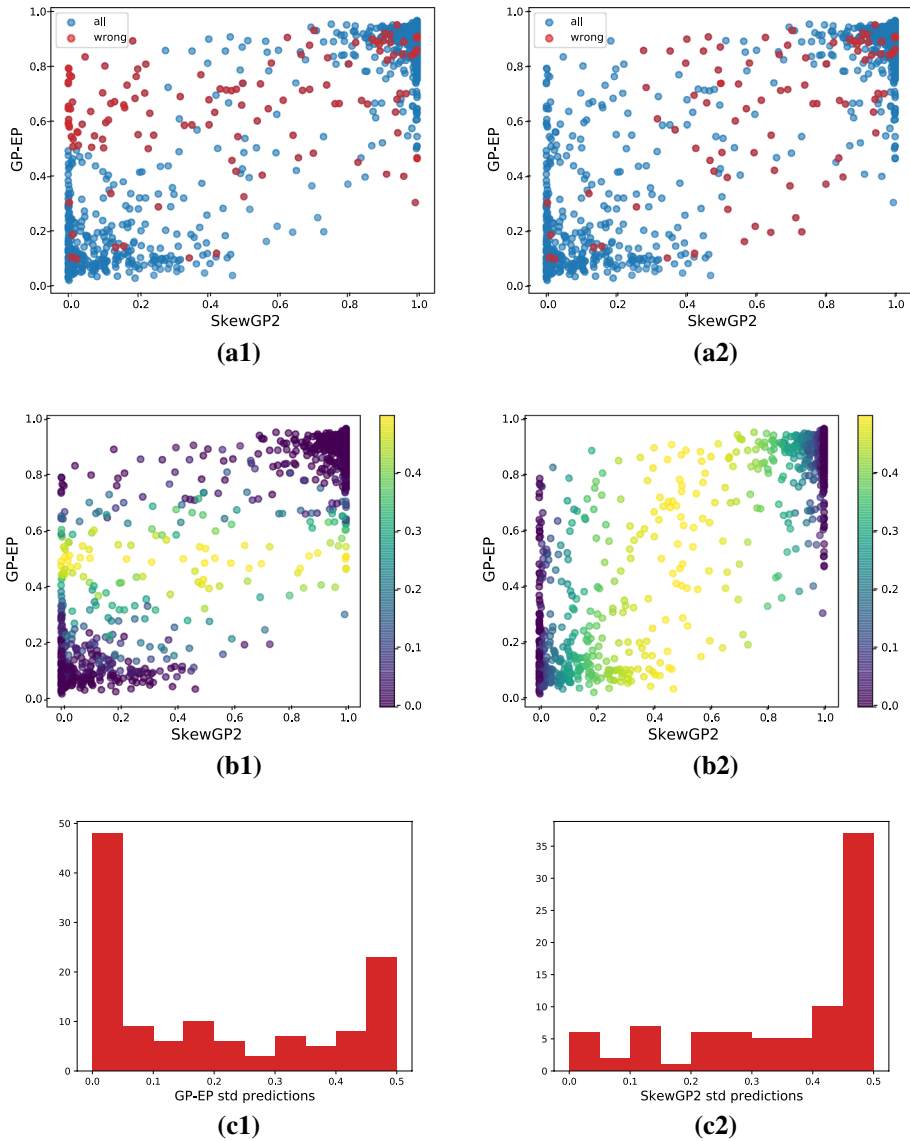


Fig. 11 Fashion MNIST dataset

## 8 Conclusions

We have introduced the Skew Gaussian process (SkewGP) as an alternative to Gaussian processes for classification. We have shown that SkewGP and the probit likelihood are conjugate and provided marginals and closed form conditionals. We have also shown that SkewGP contains the GP as a special case and, therefore, SkewGPs inherit all good properties of GPs and increase their flexibility. The SkewGP prior was applied in classification

showing improved performance over GPs (Laplace’s method and Expectation Propagation approximations).

As future work, we plan to study other more native ways to parametrize the skewness matrix  $\Delta$  that do not rely on an underlying kernel. Moreover, we plan to investigate the possibility of using inducing points, as for sparse GPs, to reduce the computational load for matrix operations (complexity  $O(n^3)$  with storage demands of  $O(n^2)$ ) as well as deriving the posterior for the multi-class classification problem.

**Acknowledgements** D. Azzimonti gratefully acknowledges support from the Swiss National Research Programme 75 “Big Data” Grant No. 407540\_167199 / 1.

**Funding** Open access funding provided by SUPSI - University of Applied Sciences and Arts of Southern Switzerland.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

### Appendix 1: Proofs

*Proposition 1* To prove that, we exploit Kolmogorov’s extension theorem (Orbanz 2009). Suppose that a family  $\mathcal{F}^I$  of probability measures are the  $I$ -finite-dimensional marginals of an infinite-dimensional measure  $\mathcal{F}$  (a “stochastic process”). Each measure  $\mathcal{F}^I$  belongs to the finite-dimensional subspace of dimension  $I$ . Given two marginals  $\mathcal{F}^I, \mathcal{F}^J$ , as marginals of the same measure  $\mathcal{F}$ , they must be marginals of each other, that is

$$\mathcal{F}^I = \mathcal{F}^{I \downarrow J} \text{ whether } I \subset J, \tag{17}$$

where  $\cdot \downarrow I$  denotes the projection onto the subspace of dimension  $I$ . A family of probability measures that satisfies (17) is called a projective family. The Kolmogorov’s extension theorem states that any projective family on the finite-dimensional subspaces of an infinite-dimensional product space uniquely defines a stochastic process on the space. This means that we can define a nonparametric Bayesian model from a finite-dimensional distribution by simply verifying that (17) holds. From (3), it then immediately follows that the definition (5) uniquely defines a stochastic process and, therefore, SkewGP is a well-defined stochastic process.

*Theorem 1* We aim to derive the posterior of  $f(X)$ . The joint distribution of  $f(X), D$  is

$$\begin{aligned} & p(D|f(X))p(f(X)) \\ &= \Phi_n(Wf; I_n) \text{SUN}_{n,s}(\xi, \Omega, \Delta, \gamma, \Gamma) \\ &= \Phi_n(Wf; I_n) \phi_n(f - \xi; \Omega) \frac{\Phi_s(\gamma + \Delta^T \bar{\Omega}^{-1} D_{\Omega}^{-1} (f - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta)}{\Phi_s(\gamma; \Gamma)} \\ &\propto \phi_n(f - \xi; \Omega) \Phi_n(Wf; I_n) \Phi_s(\gamma + \Delta^T \bar{\Omega}^{-1} D_{\Omega}^{-1} (f - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta) \end{aligned} \tag{18}$$

where we denoted  $f = f(X) \in \mathbb{R}^n$  and omitted the dependence on  $X$ . First, note that

$$\begin{aligned} &\Phi_n(Wf; I_n) \\ &= \Phi_n(W\xi + (\bar{\Omega}D_\Omega W^T)^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi); (W\Omega W^T + I_n) - WD_\Omega \bar{\Omega} D_\Omega W^T) \end{aligned}$$

Therefore, we can write

$$\begin{aligned} &\Phi_n(Wf; I_n) \Phi_s(\gamma + \Delta^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta) \\ &= \Phi_n(W\xi + (\bar{\Omega}D_\Omega W^T)^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi); (W\Omega W^T + I_n) - WD_\Omega \bar{\Omega} D_\Omega W^T) \\ &\quad \cdot \Phi_s(\gamma + \Delta^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta) \\ &= \Phi_{s+n}(m; M) \end{aligned} \tag{19}$$

with

$$m = \begin{bmatrix} \gamma + \Delta^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi) \\ W\xi + (\bar{\Omega}D_\Omega W^T)^T \bar{\Omega}^{-1} D_\Omega^{-1} (f - \xi) \end{bmatrix},$$

and

$$M = \begin{bmatrix} \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta & 0 \\ 0 & (W\Omega W^T + I_n) - WD_\Omega \bar{\Omega} D_\Omega W^T \end{bmatrix}.$$

From (18)–(19) and the definition of the PDF of the SUN distribution (1), we can easily show that we can rewrite (18) as a SUN distribution with updated parameters:

$$\begin{aligned} \tilde{\xi} &= \xi, \\ \tilde{\Omega} &= \Omega, \\ \tilde{\Delta} &= [\Delta, \bar{\Omega}D_\Omega W^T], \\ \tilde{\gamma} &= [\gamma, W\xi]^T, \\ \tilde{\Gamma} &= \begin{bmatrix} \Gamma & \Delta^T D_\Omega W^T \\ WD_\Omega \Delta & (W\Omega W^T + I_n) \end{bmatrix}. \end{aligned}$$

*Corollary 1* This follows directly from the above proof by observing that  $\Phi_{s+n}(\tilde{\gamma}, \tilde{\Gamma})$  is the normalization constant of the posterior and, therefore, the marginal likelihood is

$$\frac{\Phi_{s+n}(\tilde{\gamma}, \tilde{\Gamma})}{\Phi_s(\gamma, \Gamma)}.$$

*Corollary 2* Denote with  $\hat{f}(\hat{X}) = [f(X)^T, f(x^*)^T]^T$  and observe that the predictive distribution is by definition

$$\int \int \Phi(f^*; 1) p(f^* | f) p(f | \mathcal{D}) df df^*,$$

with  $f^* := f(x^*)$  and  $f = f(X)$ . Note that we have omitted the dependence on  $x^*, X$  for easier notation ( $p(f^* | f)$  corresponds to  $p(f^* | x^*, X, f)$ ). We can write the posterior as

$$p(f | \mathcal{D}) = \frac{p(\mathcal{D} | f) p(f)}{p(\mathcal{D})},$$

and so

$$\begin{aligned} & \int \int \Phi(f^*; 1) p(f^* | f) p(f | \mathcal{D}) df df^* \\ &= \frac{1}{p(\mathcal{D})} \int \int \Phi(f^*; 1) p(f^* | f) p(\mathcal{D} | f) p(f) df df^* \\ &= \frac{1}{p(\mathcal{D})} \int \int \Phi(f^*; 1) p(\mathcal{D} | f) p(f, f^*) df df^*. \end{aligned}$$

Observe that

$$\Phi(f^*; 1) p(\mathcal{D} | f) = \Phi(f^*; 1) \Phi_n(Wf; I_n) = \Phi_{n+1}(\hat{W}\hat{f}; I_{n+1})$$

with  $\hat{W} = \text{diag}(2y_1 - 1, \dots, 2y_{n+1} - 1)$  and  $y_{n+1} = 0.5$ . Note that  $2y_{n+1} - 1 = 2 \cdot 0.5 - 1 = 0$  and this is the reason why we have introduced the dummy class value  $1/2$ .

Observe that

$$\int \int \Phi_{n+1}(\hat{W}\hat{f}; I_{n+1}) p(f, f^*) df df^*$$

is the marginal likelihood of a SkewGP posterior corresponding to the augmented dataset  $\hat{X} = [X^T, \mathbf{x}^{*T}]^T$ ,  $\hat{y} = [y^T, 1/2]^T$ . Therefore, we have that

$$\frac{1}{p(\mathcal{D})} \int \int \Phi_{n+1}(\hat{W}\hat{f}; I_{n+1}) p(f, f^*) df df^* = \frac{\Phi_{s+n+1}(\tilde{\gamma}^*, \tilde{\Gamma}^*)}{\Phi_s(\gamma, \Gamma)},$$

where  $\tilde{\gamma}^*, \tilde{\Gamma}^*$  are the corresponding matrices of the posterior computed for the augmented dataset.

*Proposition 2* This follows by the Fréchet inequality:

$$Pr(A_1, A_2, \dots, A_b) \geq \max \left( 0, \sum_{i=1}^b Pr(A_i) - (b - 1) \right),$$

where  $A_i$  are events. In fact, note that

$$\Phi_{s+n}(\tilde{\gamma}; \tilde{\Gamma}) = Pr(u_1 \geq \tilde{\gamma}_1, \dots, u_{s+n} \geq \tilde{\gamma}_{s+n})$$

where  $Pr$  is computed w.r.t. the PDF of a multivariate distribution with zero mean and covariance  $\tilde{\Gamma}$ .

*Theorem 2* The proof follows straightforwardly from that of Corollary 2.




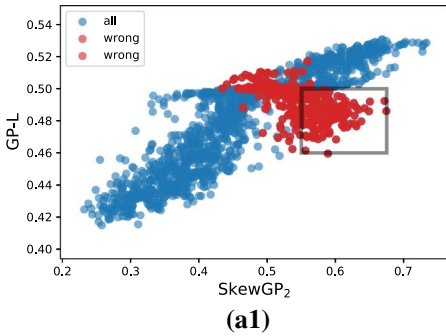
## Appendix 2: Additional image classification examples

We have defined a binary sub-problem from the German Traffic Sign data by considering Speed-Limit 30 versus 50 and from MNIST digit dataset by considering 3 versus 5. We have compared GP-L versus SkewGP<sub>2</sub> (both with RBF kernel). Table 2 reports the average accuracy and shows again SkewGP<sub>2</sub> outperforms GP-L (GP-EP achieves lower accuracy than GP-L in both cases).

Our goal is assessing the accuracy but also the quality of the probabilistic predictions. Figure 12, plot (a1), shows, for each instance in the test set (one of the fold) of the


**Table 2** Image classification

Dataset	Classes	GP-L	SkewGP <sub>2</sub>
German traffic sign	Speed-limit 30 versus 50	0.96	0.98
MNIST	3 versus 5	0.80	0.90



**Fig. 12** MNIST dataset

MNIST dataset, the value of the mean predicted probability of class 5 for SkewGP<sub>2</sub> (x-axis) and GP-L (y-axis). Each instance is represented as a blue point. The mean predicted probability ranges in [0.41, 0.53] for GP-L and in [0.25, 0.8] for SkewGP<sub>2</sub>. The red points highlight the instances that were misclassified by GP-L (plot (a2) reports the



**Fig. 13** German-road sign dataset

images of some of the misclassified instances included in the rectangle). Figure (b1) shows the same plot, but the red points are now the instances that were misclassified by  $SkewGP_2$ . By comparing (a) and (b), it is evident that  $SkewGP_2$  provides a higher quality of the probabilistic predictions.

$SkewGP_2$  also returns a better estimate of its own uncertainty. This is showed in plots (c1) and (c2). For each instance in the test set and for each sample from the posterior, we have computed the predicted class (the class that has probability greater than 0.5). For each test instance, we have then computed the standard deviation of all these predictions and used it to color the scatter plot of the mean predicted probability. In this way, we have a visual representation of first order (mean predicted probability) and second



order (standard deviation of the predictions) uncertainty. Plot 12(c1) is relative to GP-L and shows that GP-L confidence is low only for the instances whose mean predicted probability is close to 0.5. This is not reflected in the value of the mean predicted probability for the misclassified instances (compare plot (a1) and (c1) and note that the red spot in (a1) is outside the yellow area in (c1)). Conversely, the second order uncertainty of SkewGP<sub>2</sub> is clearly consistent with plot (b1).

Figure 13 shows a similar plot for the “German road-sign” dataset. Plot 13(c1) is relative to GP-L and shows that GP-L confidence is low only for the instances whose mean predicted probability is in [0.3, 0.7]. This is not reflected in the value of the mean predicted probability for the misclassified instances (compare plot (a1) and (c1) and note that some of the red points in (a1) are outside the yellow area in (c1)). Conversely, 13(c2) shows that the second order uncertainty of SkewGP<sub>2</sub> is clearly consistent with plot (b1).

## References

- Arellano, R. B., & Azzalini, A. (2006). On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics*, 33(3), 561–574.
- Azzalini, A. (2013). *The skew-normal and related families* (Vol. 3). Cambridge: Cambridge University Press.
- Azzimonti, D., & Ginsbourger, D. (2018). Estimating orthant probabilities of high-dimensional Gaussian vectors with an application to set estimation. *Journal of Computational and Graphical Statistics*, 27(2), 255–267.
- Bauer, M., van der Wilk, M., & Rasmussen, C.E. (2016). Understanding probabilistic sparse Gaussian process approximations. In *Advances in neural information processing systems* (pp 1533–1541).
- Benavoli, A., Mangili, F., Corani, G., Zaffalon, M., & Ruggeri, F. (2014). A Bayesian Wilcoxon signed-rank test based on the Dirichlet process. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2014)* (pp 1–9).
- Benavoli, A., Corani, G., Demser, J., & Zaffalon, M. (2017). Time for a change: A tutorial for comparing multiple classifiers through Bayesian analysis. *Journal of Machine Learning Research*, 18(77), 1–36.
- Botev, Z. I. (2017). The normal law under linear restrictions: Simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(1), 125–148.
- Chai, K. M. A. (2012). Variational multinomial logit Gaussian process. *Journal of Machine Learning Research*, 13, 1745–1808.
- Durante, D. (2019). Conjugate Bayes for probit regression via unified skew-normal distributions. *Biometrika*, 106(4), 765–779.
- Fang, K. W. (2018). *Symmetric multivariate and related distributions*. Boca Raton: Chapman and Hall/CRC.
- Genton, M. G., Keyes, D. E., & Turkiyyah, G. (2018). Hierarchical decompositions for the computation of high-dimensional multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 27(2), 268–277.
- Genz, A. (1992). Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 1(2), 141–149.
- Genz, A., & Bretz, F. (2009). *Computation of multivariate normal and t probabilities* (Vol. 195). Berlin: Springer.
- Gessner, A., Kanjilal, O., & Hennig, P. (2019) Integrals over Gaussians under linear domain constraints. arXiv preprint [arXiv:1910.09328](https://arxiv.org/abs/1910.09328).
- GPpy (since 2012) GPpy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPpy>
- Hernández-Lobato, D., Hernández-Lobato, J.M., & Dupont, P. (2011). Robust multi-class Gaussian process classification. In *Advances in neural information processing systems* (pp 280–288).
- Kuss, M., & Rasmussen, C. E. (2005). Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Research*, 6, 1679–1704.
- Murray, I., Adams, R., & MacKay, D. (2010). Elliptical slice sampling. In: Teh YW, Titterton M (eds) *Proceedings of the thirteenth international conference on artificial intelligence and statistics, PMLR, Chia Laguna Resort, Sardinia, Italy, proceedings of machine learning research* (Vol. 9, pp 541–548).

- Nickisch, H., & Rasmussen, C. E. (2008). Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9, 2035–2078.
- O'Hagan, A. (1991). Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 29(3), 245–260.
- O'Hagan, A., & Leonard, T. (1976). Bayes estimation subject to uncertainty about parameter constraints. *Biometrika*, 63(1), 201–203.
- Olson, R. S., La Cava, W., Orzechowski, P., Urbanowicz, R. J., & Moore, J. H. (2017). PMLB: A large benchmark suite for machine learning evaluation and comparison. *BioData Mining*, 10(1), 36.
- Orbanz, P. (2009). Construction of nonparametric Bayesian models from parametric Bayes equations. In: *Advances in neural information processing systems* (pp. 1392–1400).
- Phinikettos, I., & Gandy, A. (2011). Fast computation of high-dimensional multivariate normal probabilities. *Computational Statistics & Data Analysis*, 55(4), 1521–1529.
- Quiñero-Candela, J., & Rasmussen, C. E. (2005). A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6, 939–1959.
- Rasmussen, C. E., & Williams, C. K. (2006). *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
- Snelson, E., & Ghahramani, Z. (2006). Sparse Gaussian processes using pseudo-inputs. In: *Advances in neural information processing systems* (pp 1257–1264).
- Spall, J. C. (1998). Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3), 817–823.
- Titsias, M. (2009). Variational learning of inducing variables in sparse Gaussian processes. In: van Dyk D, Welling M (eds) *Proceedings of the twelfth international conference on artificial intelligence and statistics, PMLR, Hilton clearwater beach resort, Clearwater Beach, Florida USA, proceedings of machine learning research* (Vol. 5, pp. 567–574).
- Williams, C. K. (1998). Computation with infinite neural networks. *Neural Computation*, 10(5), 1203–1216.
- Williams, C. K., & Barber, D. (1998). Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12), 1342–1351.
- Zhang, Z., Wu, G., & Chang, E. Y. (2007). Semiparametric Regression Using Student  $t$  Processes. *IEEE Transactions on Neural Networks*, 18(6), 1572–1588.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.