Robust Gaussian Process Regression with Huber Likelihood and Projection Pursuit

Pooja Algikar[1](✉) and Lamine Mili[1]

Virginia Tech, Falls church VA 22043, USA {apooja19,lmili}@vt.edu

Abstract. Outliers in both covariates and output responses pose significant challenges for Gaussian Process (GP) regression models. We present a novel GP regression approach that effectively integrates the Huber likelihood into the GP framework—with additional parameters that can be set before inference. Specifically, we model the likelihood of observed outputs using the Huber probability distribution: this reduces deviations caused by output outliers. For covariate outliers, we introduce projection pursuit weights—attenuating their influence on the model. To address the analytically intractable, yet unimodal, posterior distribution, We employ Laplace approximation and, separately, Gibbs sampling within a Markov Chain Monte Carlo (MCMC) framework. We simplify Gibbs sampling by expressing the likelihood associated with outlying points as normally distributed through the scale mixture representation of the Laplace distribution. This work is particularly important in the field of transmission spectroscopy—where noisy measurements are often neglected in the estimation of planet-to-star radius ratios. We demonstrate the robustness and effectiveness of our method through extensive experiments on synthetic and real-world datasets.

Keywords: Covariate and Response Outliers · Transmission Spectroscopy.

1 Introduction

Bayesian inference which is based on Gaussian likelihood is known to be sensitive to extreme observations and gross errors, called outliers. The estimation of parameters in Gaussian processes (GPs) is affected in non-Gaussian error settings as the predictive uncertainty assigns equal confidence to the measurements, regardless of whether they are outliers or not. We illustrate this problem in a numerical example. Let us consider a 2-d sinc function $y(x) = \text{sinc}(x) + e$, where $x = \sqrt{(x_1^2 + x_2^2)}$ with an additive error that follows the Student's t-distribution with 10 degrees of freedom $e \sim$ Student's-t(2). We add additional large outliers $y^{(l)}$ with magnitude close to 0.8 and $x_1^{(l)}$. Figure 1(a) shows the predicted values at test points $x = [-10, 10]$, obtained from standard.

Existing studies addressing the outlier problem in GP regression use various approaches to define the likelihood. Two common strategies are: (1) using a mixture of two normal distributions or (2) employing heavy-tailed distributions.

Most of these methods assume the error distribution is known a priori—a condition that is often unrealistic in practical applications. Moreover, their robustness is questionable when faced with extreme observations that do not correspond to the non-normal distribution their heavy tailed likelihood is specified to capture. These models typically struggle to handle both general noise patterns and large errors in covariate and response dimensions, often attempting to fit extreme values. We show this shortcoming in Figure 1(b) with the sinc function data for the GP with the Student's t-likelihood and employing the MCMC integration approximation method. We notice that the model overfits when large outliers $x_1^{(l)}$ and $y^{(l)}$ occur simultaneously, as the Student's-t likelihood can effectively compensates only for errors in $y^{(l)}$.

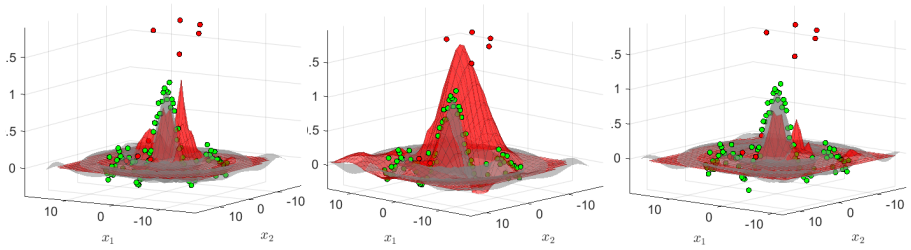

Fig. 1: Predictions for sinc2d($\boldsymbol{x}$): (a) standard GP, (b) Student's t-likelihood via MCMC, (c) GP with Huber likelihood. The red surface shows the mean of the model predictions, the grey surface represents the true sinc2d function, red dots are outliers, and green dots are training points fitting sinc2d. The proposed GP with Huber likelihood and proejction pursuit weights demonstrates robustness to the outliers $\{y^{(l)}, \boldsymbol{x}^{(l)}\}$.

In this paper, we propose a new way of handling extreme outliers in covariate space and output responses that models the likelihood of the observed data using Huber density function. We significantly enhance downweighting of the outliers compared to the earlier work by [1], which was limited to handling outliers only in the output responses with added hyperparameters (β, c).

2 Related Work

Goldberg et al. [11] introduced a dual-model Gaussian process framework to account for covariate-dependent noise. The first Gaussian process model governs the output process y, while the second Gaussian process governs the noise process. [31] investigated heavy-tailed error distributions that are constructed as scale mixtures of normal distributions, which are also used for specifying a

priori distribution based on the earlier ideas suggested by [6,24]. By doing so, the prior distribution discounts any observations highlighting inconsistency between likelihood and prior. Along the same line, [7] assumed a super heavy-tailed error distribution dependent on an explanatory variable to make the estimation of the population mean and ratios robust to outliers. [16] extended a mixture of two normal distributions, one to model small errors in regular observations and a second one to model large errors in outlying observations. However, [22] questioned the adequacy of the two-model approach. They proposed instead twin GP that allow us to choose between the distribution of the regular observations and that of the outliers. [16] suggested a GP with a Laplace likelihood model that utilizes a scale mixture representation of Laplace noise distribution where the variance follows an exponential distribution. [30] proposed a GP model based on the Student's t-likelihood function, where the noise is modeled as a scale mixture of Gaussian distributions. Unfortunately with the non-Gaussian likelihood, the Bayesian inference becomes analytically intractable. Consequently, various advanced approximation methods were proposed [16,30,14,25,5] to overcome the convergence failure of the classical approximation methods such as expectation propagation [21], Markov Chain Monte Carlo [23], variational Bayes [9], and Laplace approximation [32]. More recently, [17,2] presented a robust variants of GPs for datasets with substantial contamination removing the outlier data based on trimming parameters in iterative manner.

In GP regression models with Student's t-likelihoods [16], a scale-mixture representation of the Student's t-distribution is utilized. A variational approximation is devised presuming the Gaussian likelihood whose individual variances are Gamma distributed. Combined with the Kullback-Leibler divergence, $\mathrm{KL}(q||p)$, between the true posterior, p, and the approximation, q, an expectation maximization (EM)-type algorithm is implemented. As for the models with Laplace likelihoods, the scale mixture model yields a unimodal posterior enabling the implementation of the EP approximation and the MCMC sampling. Here, a Laplace approximation is inappropriate because the discontinuous derivatives of the Laplace likelihood at zero may cause the Hessian matrix to be undefined.

3 Contributions

[1] proposed a robust Gaussian Process (GP) regression method that leverages generalized Bayesian inference to preserve computational conjugacy. Their method handles outliers in the output responses through weighting mechanism J in the noise term: $\sigma^2 J_{ii} = \sigma^2 \left(1 + r_i^2/c^2\right)$, where r_i is the residual associated with i^{th} data point $r_i = y_i - m(\boldsymbol{x}_i)$ and c is the threshold parameter. However, a potential limitation of this approach is that it may not adequately account for outliers in the output response, $y_i^{(l)}, y_j^{(c)}$, when they occur alongside outliers in the covariate dimensions, $\boldsymbol{x}_k^{(l)} = [x_1^{(l)}, x_2^{(l)}, \ldots, x_d]$.

Our approach first addresses covariate outliers $\boldsymbol{x}_k^{(l)}$ by introducing projection pursuit weights $w(\boldsymbol{x}_k)$. These weights are then applied to scale the residuals r, ensuring that the influence of an outlier is adjusted based on the presence

of extreme covariate outliers $\boldsymbol{x}_k^{(l)}$. This method enables the model to transform contaminated data points $\{y_i^{(l)}, y_j^{(c)}, \boldsymbol{x}_k^{(l)}\}$ into a more reliable dataset. Notably, the projection pursuit weighting operates independently and can be applied to various likelihoods, as shown in our experiments.

To further handle extreme outliers in output response $y_i^{(l)}$, we employ a Huber density function—derived from the exponential of the Huber loss- giving robust L_1 norm treatment for the residuals having over-limit magnitude. The combination of projection pursuit weighting and Huber likelihood can handle cases where the locations i, j, and k coincide. Additionally, when extreme outliers are detected in the covariate dimensions $\boldsymbol{x}_k^{(l)}$, the model selectively retains the corresponding output $y_k^{(l)}$ if it improves the regression fit.

4 The Model

Let us consider a regression setting $y_i = f(\boldsymbol{x}_i) + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ is a homoscedastic i.i.d. random variable with constant variance. In GP models, the systematic dependency between the covariates $\boldsymbol{x} \in \mathcal{X}$, where $\mathcal{X} \subseteq \mathbb{R}^d$, and the response $y \in \mathcal{Y}$ is given by a latent function, $f(\boldsymbol{x}) : \mathbb{R}^d \to \mathbb{R}$. In a truly non-parametric sense, the latent vector function at n covariates, $\mathbf{f} = [f(\mathbf{x}_1), \ldots, f(\boldsymbol{x}_n)]^\top$, is assumed to have a priori probability distribution. This distribution is a joint multivariate normal distribution with zero mean vector and covariance matrix, $\mathbf{K}$, that is,

$$\mathbf{f}|\mathbf{X}, \boldsymbol{\theta} \sim \mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K}). \tag{1}$$

The covariance matrix, $\mathbf{K}$, is a positive semi-definite matrix that captures residual spatial association with elements $K_{i,j} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$, $i, j = 1, \ldots, n$. The function $k(\cdot, \cdot)$, chosen from a parametric kernel family such as the Gaussian or the Matérn kernel, is characterized by hyperparameters denoted by $\boldsymbol{\theta}$. The likelihood of the data is expressed as $\mathbf{y}|\mathbf{f}, \sigma \sim \mathcal{N}(\mathbf{y}|\mathbf{f}, \boldsymbol{\Sigma})$, and the resulting posterior distribution on $\mathbf{f}$ as where $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \ldots, \sigma_n^2)$.

Next, we develop three aspects of the proposed GP-Huber model: Huber likelihood, projection pursuit weights, and the resulting unimodal posterior distribution. Following that, we discuss the hyperparametric settings of the GP-Huber.

4.1 Huber likelihood

We propose to use the Huber density function based on the Huber loss proposed by [12] to model the likelihood of the observed data. The Huber loss function $\rho(\cdot)$ is a truncated mixture of two commonly used loss functions: squared loss, $l(r) = r^2$ for residuals below threshold b, and absolute loss, $l(r) = |r|$ for residuals $r_i = y_i - f(\mathbf{x}_i)$ below threshold b, given by

$$\rho(r) = \begin{cases} \frac{1}{2}r^2, & \text{if } |r| \leq b \\ b|r| - \frac{1}{2}b^2. & \text{otherwise} \end{cases} \tag{2}$$

[12] considered the contamination model $(1-\varepsilon)G(r)+\varepsilon H(r)$, where $G(r)$ is the Gaussian cumulative density function and $H(r)$ is the unknown cumulative density function. The associated least favorable Huber density function with a fraction of contamination ε is defined as

$$p_H(\mathbf{y}|\mathbf{f},\phi)=\prod_{i=1}^{n}\frac{1-\varepsilon}{\sqrt{2\pi\sigma}}\exp\left(-\rho(r_i)\right). \tag{3}$$

The parameter ε, symbolizing the fraction of the dataset presumed to deviate from the underlying model, can be computed utilizing the minimum covariance determinant estimator [13]. The threshold b is selected to protect estimation of the model parameters and hyperparameters against the fraction of contamination ε. The Huber likelihood provides a balance between sensitivity to inliers and robustness to outliers, controlled by the threshold b, which has a theoretical interpretation and can be set based on domain knowledge or easily tuned empirically. Student's-t likelihood, while also robust to outliers, may give undue influence to extreme observations because of its heavy tails. The Laplace likelihood's uniform linear loss may underweight small residuals—potentially leading to less efficient estimates when the data contains mostly inliers.

4.2 Projection Pursuit Weighting

The idea is to scale the residual r_i associated with the i^{th} data point with projection pursuit weight $w(\boldsymbol{x}_i)$ based on robust variant of Mahalanobis distances, called projection statistics $\text{PS}(\boldsymbol{x}_i):\mathbb{R}^d\to\mathbb{R}^d$. This scaling highlights the impact of outliers in single or multiple dimensions masking each other in the covariate space. Residual larger than the threshold b gets robust $L1$ norm treatment, while those smaller than b are treated with an efficient $L2$ norm within the Huber loss $\rho(r)$.

We obtain standardized the residual $r_{S_i}=r_i/(w_i\sigma s)$ by scaling r_i by its corresponding projection pursuit weight w_i and using a scaling factor $s=b_d\,\text{med}|\boldsymbol{r}|$, where $b_d=1+5/(n-d)$ is the dimensionality correction factor. When the error distribution is unknown, s accounts for its spread parameter. The projection pursuit weights $\boldsymbol{w}$ limit the influence of outliers simultaneously arising in multiple covariate dimensions at multiple locations on the loss function, are based on projection statistics PS_i, calculated as

$$w_i=\begin{cases}1, & \text{for } \text{PS}_i^2\le c_i,\\ \frac{c_i}{\text{PS}_i^2}, & \text{for } \text{PS}_i^2>c_i.\end{cases} \tag{4}$$

The projection statistics [28,8] are a robust version of Mahalanobis distances based on the median absolute distance from the median. Formally defined as the maxima of the standardized projection distances obtained by projecting the point cloud in the directions that originate from the co-ordinate wise median

and that pass through each of the data points, $\boldsymbol{x}_i$ [20]. They're easy to calculate:

$$\mathrm{PS}_i = \underset{||\boldsymbol{u}_j||=1}{max} \frac{|\boldsymbol{x}_i^T \boldsymbol{u}_j - \underset{k}{\mathrm{median}}(\boldsymbol{x}_k^T \boldsymbol{u}_j)|}{1.4826 \underset{i}{\mathrm{median}} |\boldsymbol{x}_i^T \boldsymbol{u}_j - \underset{k}{\mathrm{median}}(\boldsymbol{x}_k^T \boldsymbol{u}_j)|}, \tag{5}$$

where $\boldsymbol{u}_j = \frac{\boldsymbol{x}_j - \mathbf{M}}{||\boldsymbol{x}_j - \mathbf{M}||}$; $j, k = 1, \ldots, n$. The co-ordinate wise median $\mathbf{M}$ is given by $\mathbf{M} = \{\underset{j=1,\ldots,n}{\mathrm{med}} \boldsymbol{x}_{j1}, \ldots, \underset{j=1,\ldots,n}{\mathrm{med}} \boldsymbol{x}_{jd}\}$. The projection statistics attain the maximum breakdown point given by $[(n-d-1)/2]/n$ [19].

[27] and [20] showed that, when $n > 5d$, the squared projection statistics PS_i^2 roughly follow a χ^2 distribution with a degree of freedom equal to the number of non-zero elements ν_i in the row vector of the associated regressor, $\mathbf{x}_i$, i.e., $\mathrm{PS}_i^2 \sim \chi^2_{\nu_i}$. However, when $n \leq 5d$, it is the PS that roughly follow a χ^2 distribution, that is, $\mathrm{PS}_i \sim \chi^2_{\nu_i}$. Consequently, the threshold c_i is chosen as the 97.5 percentile of the chi-square distribution with ν_i degrees of freedom while defining weights in (4). Throughout the inference process (as detailed in Section 5), we use standardized residuals r_{S_i} within the Huber likelihood.

$$p_H(\mathbf{y}|\mathbf{f}, \phi) = \prod_{i=1}^{n} \frac{1-\varepsilon}{\sqrt{2\pi}\sigma} \exp\left(-\rho(r_{S_i})\right). \tag{6}$$

4.3 GP-Huber posterior

The posterior distribution resulting from our model, which incorporates a non-conjugate prior, is given as:

$$p(\mathbf{f}|\mathcal{D}, \boldsymbol{\theta}, \sigma) = \frac{p_G(\mathbf{f}|\mathbf{0}, \mathbf{K})}{p(\mathcal{D}|\boldsymbol{\theta}, \sigma)} p_H(\mathbf{y}|\mathbf{f}, \sigma), \tag{7}$$

where where $p_G(\mathbf{f}|\mathbf{0}, \mathbf{K})$ is the Gaussian prior $\mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K})$ and $p_H(\mathbf{y}|\mathbf{f}, \sigma)$ is the likelihood modeled using the Huber density. This formulation leads to a posterior that does not have a closed-form expression due to the non-conjugate nature of the Huber likelihood. The marginal likelihood (or evidence) of the data, which plays a crucial role in model selection and hyperparameter optimization, is expressed as:

$$p(\mathcal{D}|\sigma, \boldsymbol{\theta}) = \int p_G(\mathbf{f}|\mathbf{0}, \mathbf{K}) p_H(\mathbf{y}|\mathbf{f}, \sigma) d\mathbf{f}. \tag{8}$$

Theorem 1. *Let $\mathcal{D} = (\boldsymbol{x}_i, y_i)_{i=1}^n$ be a dataset with distinct covariates $\boldsymbol{x}_i \in \mathcal{X}$ and response $y_i \in \mathcal{Y}$, where $n < \infty$. The kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ is positive definite, with elements $K_{ij} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$ defined by a continuous kernel function $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$. Assume the Huber likelihood function $p_H(\mathbf{y}|\mathbf{f}, \boldsymbol{\sigma})$ based on strictly convex and continuous Huber loss $\rho(r_i) : \mathbb{R} \to \mathbb{R}$. Then the posterior distribution $p(\mathbf{f}|\mathcal{D}, \boldsymbol{\theta}, \sigma)$ is unimodal.*

The proof is presented in appendix 2.1. This theorem shows that despite the non-Gaussian and potentially complex nature of the Huber likelihood, the posterior retains a single peak. This simplifies both parameter inference and hyperparameter optimization.

We can set the threshold $b = 1.5$ to achieve high efficiency at the Gaussian distribution (see appendix 2.3). This would make our model robust to 10% outliers (since fraction of contamination is $\varepsilon = 0.1$). Note that, in the context of our work, "efficiency" refers to the estimator's ability to achieve low variance when the noise follows a Gaussian distribution. Specifically, a highly efficient estimator can make the best use of data that is predominantly Gaussian, leading to more accurate parameter estimation. The contamination fraction ε defines the model's tolerance to deviations from the Gaussian assumption, allowing it to handle a proportion of outlier points without being overly influenced by them. The parameter b controls the threshold for identifying outliers and thus influences the transition between $L2$ and $L1$ norm treatment. By setting $b = 0.45$, we get $\varepsilon = 0.45$ for heavy-tailed and Gaussian error distributions, we aim to accommodate up to 45% outliers while maintaining reasonable efficiency. The only hyperparameter of the likelihood function requiring estimation is $\phi = \sigma^2$.

5 Approximate Bayesian Inference

By retaining the optimization-friendly properties of convex problems ensured by to unimodality (see Theorem 1), our method enables the use of the Laplace approximation [29] for the posterior. To facilitate predictions f^*, we develop Gibbs sampling and Laplace's method. The key requirement for the latter is the continuity of the Huber density function. In Gibbs sampling, the joint posterior distribution $p(\mathbf{f}, \boldsymbol{\theta}, \sigma^2)$ can be simplified using the scale mixture model of the Laplace distribution for data points with residuals $r \geq b$: this representation expresses the likelihood of these points as a normal distribution—making the sampling process more efficient.

5.1 Gibbs Sampling

The Huber density function is a mixture of a truncated normal and a Laplace density function for an absolute standardized residual respectively lying within and outside the threshold b. This yields

$$p_H(y|f, \boldsymbol{\sigma}) = \begin{cases} \frac{C_1}{\sqrt{2\pi} w_i \sigma_g s} \exp\left(-\frac{r_i^2}{2 w_i^2 \sigma_g^2 s^2}\right) & |r_{S_i}| \leq b, \\ \frac{C_2}{2 w_i a s} \exp\left(-\frac{b|r_i|}{w_i a s}\right) & |r_{S_i}| > b, \end{cases} \tag{9}$$

where C_1 and C_2 are the constants respectively, defined as $C_1 = 1 - \varepsilon$ and $C_2 = \sqrt{\frac{\pi}{2}} \exp(b^2/2)$. The Laplace distribution $p_L(y_i|f(\mathbf{x}_i), a)$ with location parameter a can be represented as a scale mixture of normal distributions $\mathcal{N}(y_i|f(\mathbf{x}_i), \sigma_i^2)$ where σ_i^2 follows an exponential distribution $p_E(\sigma_i^2|\beta)$ [3] and $i = 1, \ldots, n_l$ are

the indices of the points associated with the standardized residuals larger than the threshold b hereafter referred to as outlying points. Formally, we have

$$p_L(y_i|f(\mathbf{x}_i), a) = \int p_G(y_i|f(\mathbf{x}_i), \sigma_i^2) p_E(\sigma_i^2|\beta) d\sigma_i^2. \tag{10}$$

Using this property, we represent the individual standard deviations corresponding to n_l outlying training points as $\{\sigma_{l_1}, \ldots, \sigma_{l_{n_l}}\}$, which are elements of the vector $\boldsymbol{\sigma_l}$. The variance associated with n_g inlying points is denoted as σ_g^2. Conclusively, the Huber probability density function takes the form

$$\mathbf{y}|\mathbf{f}, \sigma_g^2, \boldsymbol{\sigma}_l^2, \beta \sim \begin{cases} \prod_{i=1}^{n_g} C_1 \mathcal{N}(y_i|f(\mathbf{x}_i), \sigma_g^2) & |r_{S_i}| \leq b, \\ \prod_{i=1}^{n_l} C_2 \mathcal{N}(y_i|f(\mathbf{x}_i), \sigma_{l_i}^2) \text{Exponential}(\sigma_{l_i}^2, \beta) & |r_{S_i}| > b, \end{cases} \tag{11}$$

where $n_g + n_l = n$ is the total number of points in the training dataset. An alternative representation of the likelihood function is given by

$$\mathbf{y}_g, \mathbf{y}_l|\mathbf{f}_g, \mathbf{f}_l, \sigma_g^2, \boldsymbol{\sigma}_l^2 \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{y}_g|\mathbf{f}_g \\ \mathbf{y}_l|\mathbf{f}_l \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{gg} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{ll} \end{bmatrix} \right), \tag{12}$$

where $\boldsymbol{\Sigma}_{gg}$ and $\boldsymbol{\Sigma}_{ll}$ both are diagonal matrices, the former with constant diagonal elements equal to σ_g^2 and the latter with diagonal entries $\{\sigma_{l_1}^2, \ldots, \sigma_{l_{n_l}}^2\}$. Let the hyperparameter vector $\boldsymbol{\sigma^2}$ consist of the diagonal entries of the matrix $\boldsymbol{\Sigma}_{gg}$, which are σ_g^2 and $\boldsymbol{\sigma}_l^2$. The joint posterior probability density function of $\mathbf{f}$, $\boldsymbol{\sigma}^2$, and $\boldsymbol{\theta}$ is given by

$$p(\mathbf{f}, \boldsymbol{\sigma}^2, \boldsymbol{\theta}) \propto p(\mathbf{y}|\mathbf{f}, \boldsymbol{\sigma}^2) p_G(\mathbf{f}|\mathbf{0}, \mathbf{K}) p(\boldsymbol{\sigma}^2|\boldsymbol{\beta}) p(\boldsymbol{\beta}|\boldsymbol{\zeta}) p(\boldsymbol{\theta}|\boldsymbol{\zeta}). \tag{13}$$

We assume that the hyper-hyperparameter vector $\boldsymbol{\beta}$ and the hyperparameter vector $\boldsymbol{\theta}$ follow the log-uniform distribution with parameters contained in $\boldsymbol{\zeta}$. Since the distribution of the variance parameter σ_g^2 of n_g inlying training points is degenerate, the hyper-hyperparameter vector $\boldsymbol{\beta} = [\beta_g, \beta_l]^T$ corresponding to the n_g points follows a degenerate distribution as well. Therefore, $p(\sigma_g^2|\beta_g)$ is a Dirac impulse while $\boldsymbol{\sigma}_l^2|\beta_l \sim \text{Exponential}(\boldsymbol{\sigma}_l^2|\beta_l)$. The samples generated from this distribution are highly correlated. Therefore, in order to better mix the Monte Carlo chains, we follow the trick used by [16] as follows:

$$p(\boldsymbol{\sigma}^2, \boldsymbol{\beta}, \boldsymbol{\theta}) \propto \left[\int p_G(\mathbf{y}|\mathbf{f}, \boldsymbol{\Sigma}) p_G(\mathbf{f}|\mathbf{0}, \mathbf{K}) d\mathbf{f} \right] p(\boldsymbol{\sigma}^2|\boldsymbol{\beta}) p(\boldsymbol{\beta}|\boldsymbol{\zeta}) p(\boldsymbol{\theta}|\boldsymbol{\zeta}), \tag{14}$$

where the covariance matrix of the n_g inlying samples and the n_l outlying samples is given by $\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{gg} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{ll} \end{bmatrix}$. The samples can be used to obtain the approximated probability density functions of the latent vector function, $p(\mathbf{f}^*|\mathcal{D}, \mathbf{X}^*)$, at the new test covariates contained in $\mathbf{X}^*$ by averaging over all unknowns. Formally, we have

$$p(\mathbf{f}^*|\mathcal{D}, \mathbf{X}^*) = \int p(\mathbf{f}^*|\mathbf{f}, \boldsymbol{\sigma}^2, \boldsymbol{\theta}, \mathbf{X}^*, \mathcal{D}) p(\mathbf{f}, \boldsymbol{\sigma}^2, \boldsymbol{\theta}|\mathcal{D}) d\mathbf{f} d\boldsymbol{\sigma}^2 d\boldsymbol{\theta}. \tag{15}$$

For T samples, it can be evaluated as

$$p(\mathbf{f}^*|\mathcal{D}, \mathbf{X}^*, \boldsymbol{\zeta}) = \frac{1}{T} \sum_{t=1}^{T} \int p(\mathbf{f}^*|\mathbf{f}, \mathbf{X}, \mathbf{X}^*, \boldsymbol{\theta}_t) p(\mathbf{f}|\mathcal{D}, \boldsymbol{\sigma}_t^2, \boldsymbol{\theta}_t) d\mathbf{f}. \tag{16}$$

Table 1: RMSE and MAE values on the Neal dataset for the Case 1. Values in parentheses represent the performance for Case 3. Bold values highlight the best performance with the lowest RMSE and MAE.

	SCtMCMC	tLA	HuberMCMC^{+pw}	HuberLA^{+pw}	RCGP	GP	LaplaceMCMC
			$\varepsilon \sim \mathcal{N}(0.01, 0.08)$				
RMSE	0.74 (0.52)	0.75 (1.31)	0.37 (0.42)	**0.25** (**0.25**)	1.84 (0.82)	1.44 (0.90)	0.43 (0.46)
MAE	0.47 (0.25)	0.48 (0.61)	0.31 (0.25)	**0.14** (**0.14**)	1.28 (0.54)	1.24 (0.68)	0.33 (0.26)
			$\varepsilon \sim$ Student-$t(10)$				
RMSE	4.86 (11.56)	1.22 (1.31)	**0.50** (0.81)	1.17 (**0.37**)	1.89 (0.88)	1.52 (0.98)	0.59 (0.93)
MAE	1.67 (1.25)	0.77 (0.65)	**0.41** (0.39)	0.79 (**0.18**)	1.71 (0.85)	1.34 (0.22)	0.43 (0.35)
			$\varepsilon \sim$ Laplace$(0, 0.1)$				
RMSE	4.76 (0.48)	1.23 (1.31)	**0.58** (0.42)	1.17 (**0.35**)	1.95 (0.86)	1.51 (0.89)	1.06 (0.82)
MAE	1.64 (0.23)	0.76 (0.61)	**0.41** (0.24)	0.68 (**0.18**)	1.27 (0.46)	1.23 (0.41)	0.75 (0.34)
			$\varepsilon \sim$ Student-$t(1)$ (Cauchy)				
RMSE	4.75 (0.57)	1.25 (1.32)	0.61 (0.49)	1.20 (**0.17**)	1.97 (0.62)	1.50 (0.89)	**0.42** (0.75)
MAE	1.65 (0.27)	0.78 (0.67)	**0.47** (0.27)	0.81 (**0.11**)	1.78 (0.42)	1.32 (0.66)	0.66 (0.38)

5.2 Laplace Approximation

To ensure the continuity of the derivative of the Huber density function with respect to the latent vector function $\mathbf{f}$, we utilize the pseudo-Huber loss function [4], which is defined as

$$\rho(r_S) = b^2 \left(\sqrt{(1 + \left(\frac{r_S}{b}\right)^2} - 1 \right). \tag{17}$$

Laplace approximation of the posterior requires the likelihood to be log-concave in order for it to be represented by a unimodal multivariate normal distribution. It is executed by approximating the posterior distribution of $\mathbf{f}$ with a normal distribution [26], that is,

$$\mathbf{f}|\mathcal{D}, \sigma, \boldsymbol{\theta} \sim \mathcal{N}(\hat{\mathbf{f}}|\mathbf{f}, \mathbf{A}). \tag{18}$$

The remainder of the method is detailed in appendix 1. Finally, we present the following theorem which guarantees the robustness of GP-Huber to outliers.

Theorem 2. *Under the same assumptions as Theorem 1, the influence of an individual observation y on the posterior mean $\mathbb{E}[f \mid y]$ is bounded:*

$$\left| \frac{\partial}{\partial y} \mathbb{E}[f \mid y] \right| \leq \frac{b}{\sigma}.$$

Proof is provided in appendix 2.2

6 Experiments

Through our experiments, we aim to address the following questions:
($\mathcal{Q}1$) When is HuberLA^{+pw} (GP-Huber with Laplace's method with pursuit

weights) preferable, and under which outlier scenarios is HuberMCMC^{+pw} (GP-Huber with Gibbs sampling with pursuit weights) more suitable?
($\mathcal{Q}2$) Does GP-Huber show a significant performance improvement over standard GP regression and the RCGP method proposed by [1] under their experimental settings?
($\mathcal{Q}3$) Does projection pursuit weighting give GP-Huber an edge over baselines with the same weighting?
($\mathcal{Q}4$) Does GP-Huber provide more accurate estimates of the planet-to-star radius ratio compared to the standard GP method used by [10] in the transmission spectroscopy experiment?

We conducted experiments on benchmark datasets with extreme outliers in location, magnitude, and error distribution. The threshold b was 1.5 for Gaussian errors and 0.45 for Student's-t, Laplace, and Cauchy errors. An anisotropic squared exponential kernel was used, with a zero mean function except in the spectroscopy experiment. Performance was measured using RMSE and MAE. Our implementation is available online [1].

.

6.1 Neal Dataset

We evaluate the proposed GP-Huber on the Neal dataset [23] for the following cases of extreme outliers:

Case 1: Extreme outliers $y_i^{(l)}, \boldsymbol{x}_k^{(l)}$ in added in output and covariate dimensions, respectively.
Case 2: Only output dimensions $y_i^{(l)}$ were contaminated with extreme data points.
Case 3: Bad data points $y_j^{(c)}$, $\boldsymbol{x}_k^{(l)}$in added to both output and covariate dimensions, respectively, with the former being relatively close to the main data cluster compared to Case 1.
Case 4: Only output dimensions were contaminated with data points $y_j^{(c)}$ relatively close to the data cloud compared to Case 1.

In all the cases above, the locations i, j and k may differ or coincide (refer to appendix 3.1 for the location and magnitude details on outliers). For each case, we considered four different error distributions: $\mathcal{N}(0.01, 0.08)$, Student-t(10), Laplace$(0, 0.1)$, Student's-t(1).

The baseline models considered for comparison on the Neal dataset, along with RCGP, include: GP with a Student's t error model solved using MCMC integration (SCtMCMC), GP with a Student's t error model using Laplace approximation (tLA), and GP with a Laplace likelihood solved via MCMC integration (LaplaceMCMC). Table 1 presents the RMSE and MAE values comparing GP-Huber against these baselines for Cases 1 and 3. Refer to appendix 3.1 for the Tables A3, A4 for the Cases 2, 4 and appendix 2.4 for the implementation details of the baselines. Furthermore, pursuit weighting is incorporated into all

[1] https://anonymous.4open.science/r/GpHuber-C9D6

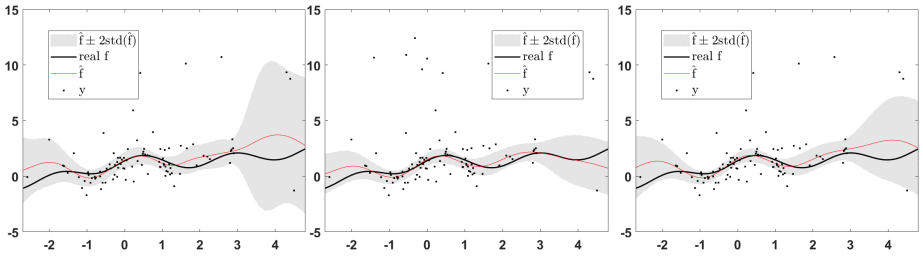

(a) (b) (c)

Fig. 2: Predicted values for the Case 1 of the Student's t-error distribution for the Neal dataset obtained from the eight considered GP regression models: (a) HuberMCMC ; (b) HuberLA; (c) RCGP.

baseline models, and their performances are compared in Tables A5 to A8 (provided in appendix 3.2). Now, we are in position to answer $\mathcal{Q}1$.

When is HuberMCMC better?

In scenarios with $y^{(l)}, \boldsymbol{x}^{(l)}$ (Case 1), HuberMCMC performed better than HuberLA (see, Tables 1 and A3). HuberMCMC also outperformed tLA in predictive accuracy, demonstrating a more robust fit that is less influenced by $\boldsymbol{x}^{(l)}$ (Figure 2). HuberLA generally provided better uncertainty quantification compared to HuberMCMC (see Figures 2 and A3), while maintaining competitive predictive performance. In outlier scenarios with $y^{(l)}$ (Case 2), HuberMCMC exhibited superior performance across Student's-t, Laplace, and Cauchy error distributions (see, Table A3). This suggests that HuberMCMC is a robust choice for datasets containing extreme output outliers i.e. outlier scenarios similar to Cases 1 and 2.

When is HuberLA better?

HuberLA exhibited superior performance in handling closer output outliers $y^{(c)}$ compared to HuberMCMC (values in parenthesis in the Table 1 and Table A4). Figure A4 highlights HuberLA's robustness to $x^{(l)}$, in contrast to tLA which is influenced by such points. While HuberLA generally provided more accurate predictions and reliable uncertainty quantification than both HuberMCMC and tLA, HuberMCMC performed competitively for the Cases 3 and 4.

From Tables A5 and A6 in appendix 3.2 (Cases 1 and 3), where projection pursuit weights were added to other baselines, we observe that HuberLA and HuberMCMC benefit the most from these weights. While Student's-t likelihood also scales residuals by pursuit weights, its logarithmic penalty $\propto \log(1 + r^2/\nu)$ (with ν controlling tail heaviness) is less sensitive to large residuals than Huber likelihood's linearized penalty $\propto |r|$. Laplace likelihood similarly penalizes $|r|$ but lacks a quadratic center, making Huber likelihood the optimal balance of robustness and efficiency. Tables A7 and A8 show the results for outlier Cases 2 and 4, where projection pursuit weighting is added to other baselines. The weights equal 1 for all data points due to the absence of covariate outliers $x^{(l)}$. GP-Huber

performs comparably to other baselines in these cases. This demonstrates that the weighting mechanism enhances GP-Huber's accuracy, addressing $\mathcal{Q}3$.

6.2 UCI datasets

In this set of experiments, we compared the performance of GP-Huber on the UCI datasets, Energy and Yacht, against RCGP and other baselines: t-GP, m-GP, and standard GP, using the outlier settings from [1]. We specifically focused on the "focused outlier" and "asymmetrical outlier" scenarios, as they closely resemble our extreme and close outlier cases. MAE values of the comparison are

Table 2: MAE values for energy and yacht. Bold values indicate the best performance for each row.

	GP	RCGP	t-GP	m-GP	HuberMCMC	HuberLA
			Focused Outliers			
Energy	0.03 (0.04)	**0.02** (0.00)	0.03 (0.05)	0.24 (0.00)	0.12 (0.01)	0.04 (0.01)
Yacht	0.26 (0.15)	**0.10** (0.14)	0.20 (0.04)	0.24 (0.00)	0.24 (0.02)	0.18 (0.00)
			Asymmetric Outliers			
Energy	0.54 (0.02)	0.44 (0.04)	0.42 (0.02)	0.41 (0.00)	0.47 (0.02)	**0.11 (0.00)**
Yacht	0.54 (0.06)	0.35 (0.02)	0.41 (0.00)	0.40 (0.00)	0.51 (0.01)	**0.12 (0.00)**

presented in Table 2. As expected, HuberLA demonstrates to be more robust than HuberMCMC since the asymmetrical and focused outliers cases considered in the study of [1] broadly fall under the Cases 3 and 4 in our study. On the Energy dataset, HuberLA outperformed both tLA and RCGP.

In our experiments, HuberLA outperformed RCGP and other baselines significantly in asymmetric outlier case and also showing the good computational efficiency, thus answering $\mathcal{Q}2$. Note that the outliers are present only in the response and not in the covariate dimensions, the projection pursuit weights are equal to 1 for these datasets.

Computational costs for the experiments on Neal and UCI datasets are presented in Table A10 and A11 in appendix 4. HuberLA—similar to RCGP and tLA—requires less computational time than HuberMCMC, as expected. The models converge faster for unidimensional data: HuberMCMC performs comparably to MCMC techniques with Student-t likelihood.For multidimensional cases, HuberMCMC, as expected for sampling-based methods, requires more time to converge, while HuberLA achieves faster convergence (between 5 to 10 s).

6.3 Transmission Spectroscopy

Transmission spectroscopy records the relative change in the stellar flux, which is the incident photons per unit area, as a planet travels in front of the star.

The sources of error, such as photon noise and instrumental and astrophysical systematics, raise many potential challenges for precise planet's atmosphere characterization. The goal is to infer the planet to star radius ratio ρ_{radius} from the observed flux as the planet passes in front of the star. The optical state parameters are metered via auxiliary measurements of the spectral trace such as position, width, angle, or other parameters, indicating the state of the detector and optics, which are thought to be the cause of instrumental systematics. Instead of modeling the latter as a linear function of the optical state parameters, [10] proposed a non-parametric model by leveraging GPs.

The observation set obtained from HST-NICMOS includes the light curves for 18 wavelength channels extracted from $n = 638$ spectra of the planetary system HD-189733. The flux measurements contained in the vector, $\boldsymbol{f} = [f_1, f_2, \ldots, f_n]^T$, are recorded at n time instants, $\{t_1, t_2, \ldots, t_n\}$ and the optical state parameters $\mathbf{x}_{t_i}$ collected in the matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ constitute the training dataset. We extend the work of [10] by using the GP-Huber model to estimate the planet-to-star radius ratio ρ_{radius}. As demonstrated earlier, the robustness to outliers of GP-Huber allows us to utilize 517 measurements associated with four out-of-transit orbits, namely orbit numbers, $\{2, 3, 4, 5\}$, and 137 measurements associated with one in-transit orbit, namely orbit number 1. The latter was excluded from the analysis performed by [10] as it constitutes much larger systematics effects attributed to the spacecraft settling. The observed transit flux modeled in the GP framework follows a normal distribution, that is,

$$\boldsymbol{f}(\boldsymbol{t}, \mathbf{X}) \sim \mathcal{N}(\boldsymbol{T}(\boldsymbol{t}, \boldsymbol{\phi}), \mathbf{K}), \tag{19}$$

where the parameter vector, $\boldsymbol{\phi}$, include the parameter of interest, ρ_{radius}, and other parameters. We consider the analytical quadratic limb darkening transit function proposed by [18]. Analogous with (11), we assume that the observed transit flux vector, $\boldsymbol{f} = \boldsymbol{f}(\boldsymbol{t}, \mathbf{X})$, in the GP-Huber framework follows a normal distribution, that is,

$$\boldsymbol{f} | \boldsymbol{T}(\boldsymbol{t}, \phi), \mathbf{X}, \boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2 \sim \mathcal{N}\left(\boldsymbol{T}(\boldsymbol{t}, \mathbf{X}), \boldsymbol{\Sigma} + \mathbf{K}\right). \tag{20}$$

The joint un-normalized log-posterior function of $\boldsymbol{\phi}$, $\boldsymbol{\beta}$, and $\boldsymbol{\theta}$ with the gamma aprior probability density function, $p(\boldsymbol{\theta}) = \frac{1}{\boldsymbol{l}} \exp\left(\frac{-\boldsymbol{\theta}}{\boldsymbol{l}}\right)$, over the covariance function hyperparameters is given by

$$\begin{aligned} \log P(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta} | \boldsymbol{f}, \mathbf{X}, \boldsymbol{\zeta}) &= \log \left(\mathcal{L}(\mathbf{r}_S | \mathbf{X}, \boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2)\right) \\ &- \tfrac{\tau}{l_\tau} - \sum_{i=1}^{d} \left(\frac{1}{s_i l_i}\right) + \log(\boldsymbol{\beta}) - \boldsymbol{\beta}^T \boldsymbol{\sigma}^2 + \log(p(\boldsymbol{\beta} | \boldsymbol{\zeta})) + \mathrm{C}. \end{aligned} \tag{21}$$

The challenging task now is to infer the parameter ρ_{radius} from the joint posterior distribution of $(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta})$. The log-likelihood $\mathcal{L}$ term is expressed as

$$\log \mathcal{L}(\boldsymbol{r}_S | \mathbf{X}, \boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2) = \tfrac{-1}{2} \mathbf{r}_S^T (\boldsymbol{\Sigma} + \mathbf{K})^{-1} \mathbf{r} - \tfrac{1}{2} \log |\boldsymbol{\Sigma} + \mathbf{K}| - \tfrac{n}{2} \log(2\pi) + \log(1 - \varepsilon), \tag{22}$$

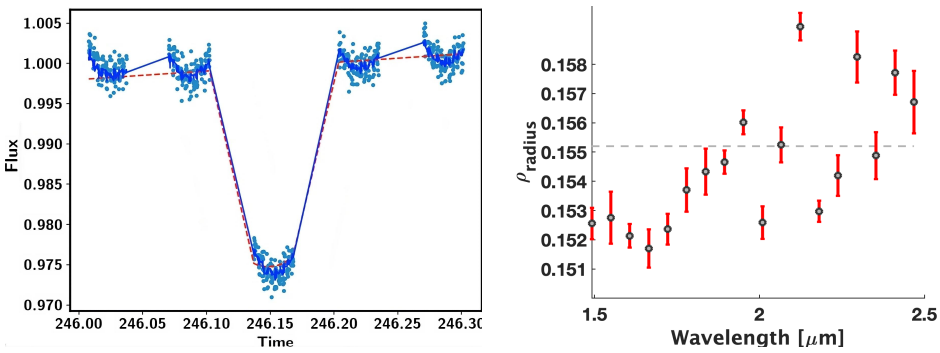

(a) (b)

Fig. 3: Transit curve fit and estimated ρ_{radius}. (*a*) Transit curve mean function $T(t, \boldsymbol{\theta})$ (in red dotted line) and GP-Huber model fit (in blue solid line); (*b*) results of planet-to-star radius ratios (ρ_{radius}) obtained from GP-Huber with error-bars. The dashed grey line represents the sanity check values.

where $\mathbf{r} = \boldsymbol{f} - \boldsymbol{T}(\boldsymbol{t}, \mathbf{X})$. One of the approaches is to use the Bayesian method that seeks the posterior distribution of ρ_{radius} by marginalizing over the other parameters of the mean function parameters $\boldsymbol{\phi}$ and the covariance function hyperparameters, $\boldsymbol{\theta}$ using MCMC methods. The other method proposed as the type-II maximum likelihood method by [10], where the hyperparameters, $\boldsymbol{\theta}$ and $\boldsymbol{\sigma}^2$. Formally, we have

$$(\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\sigma}}^2, \hat{\boldsymbol{\beta}}) = \underset{\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta}}{\arg\max} \log P(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta} | \boldsymbol{f}, \mathbf{X}, \boldsymbol{\zeta}). \tag{23}$$

And the posterior distribution of the parameter of interest ρ_{radius} is obtained by marginalizing the joint posterior distribution $p(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta})$ over the hyperparameters and the rest of the mean function parameters. In the standard type II maximum likelihood method, the hyperparameters are fixed to their maximum likelihood estimates i.e. by maximizing the evidence $p(\mathcal{D}|\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\sigma}^2)$.

Figure 3(a) shows the transit fit obtained for one wavelength channel. Figure 3(b) shows the estimated ρ_{radius} obtained using MCMC integration over the rest of the mean function parameters $\boldsymbol{\phi}$ and hyperparameters $\boldsymbol{\theta}$ along with the values estimated from the white light curve represented as the white dashed line. Note that the estimated ρ_{radius} values are very close to the white light curve value of 0.155. Most of our results agree with the results obtained from the Gibson model except for wavelength channels 1.665μm and 2.124μm (see, appendix 3.3), which effectively answers $\mathcal{Q}4$. We retain the noisy orbit-1 observations—and still, GP-Huber delivers estimates on par with, or better than, existing methods. This highlights its robustness and accuracy in the face of outliers in both the flux and the optical state parameters, data that spectroscopy models typically discard.

7 Conclusions

The proposed GP-Huber model shows promise for handling a variety of heavy-tailed and Gaussian error distributions with extreme outliers in both covariate and output dimensions. Notably, it introduces additional parameters, b and ε, which can be heuristically set prior to parameter inference. The model's unimodal posterior simplifies Gibbs sampling and allows for an efficient Laplace approximation. We prove the bounded influence of observations on the posterior mean. From our experiments on the Neal and UCI datasets, we found that HuberMCMC$^{+\text{pw}}$ shows superior robustness against extreme outliers, while HuberLA$^{+\text{pw}}$ performs better with near outliers, compared to RCGP$^{+\text{pw}}$ and other baselines$^{+\text{pw}}$. Additionally, the transmission spectroscopy experiment demonstrates their potential in real-world applications.

Future work involves extending the scalability of GP-Huber to handle large datasets by implementing sparse inference techniques.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Altamirano, M., Briol, F.X., Knoblauch, J.: Robust and conjugate gaussian process regression. In: Forty-first International Conference on Machine Learning (2024)
2. Andrade, D., Takeda, A.: Robust gaussian process regression with the trimmed marginal likelihood. In: Uncertainty in Artificial Intelligence. pp. 67–76. PMLR (2023)
3. Andrews, D.F., Mallows, C.L.: Scale mixtures of normal distributions. Journal of the Royal Statistical Society: Series B (Methodological) **36**(1), 99–102 (1974)
4. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. IEEE Transactions on image processing **6**(2), 298–311 (1997)
5. Daemi, A., Alipouri, Y., Huang, B.: Identification of robust gaussian process regression with noisy input using em algorithm. Chemometrics and Intelligent Laboratory Systems **191**, 1–11 (2019)
6. De Finetti, B.: The bayesian approach to the rejection of outliers. In: Proceedings of the fourth Berkeley Symposium on Probability and Statistics. vol. 1, pp. 199–210. University Press Berkeley, Calif. (1961)
7. Desgagné, A., Gagnon, P.: Bayesian robustness to outliers in linear regression and ratio estimation. Brazilian Journal of Probability and Statistics **33**(2), 205–221 (2019)
8. Donoho, D.L.: Breakdown properties of multivariate location estimators. Tech. rep., Technical report, Harvard University, Boston. URL http://www-stat. stanford ... (1982)
9. Ghahramani, Z., Beal, M.: Propagation algorithms for variational bayesian learning. Advances in neural information processing systems **13** (2000)
10. Gibson, N., Aigrain, S., Roberts, S., Evans, T., Osborne, M., Pont, F.: A gaussian process framework for modelling instrumental systematics: application to transmission spectroscopy. Monthly notices of the royal astronomical society **419**(3), 2683–2694 (2012)

11. Goldberg, P., Williams, C., Bishop, C.: Regression with input-dependent noise: A gaussian process treatment. Advances in neural information processing systems **10** (1997)
12. Huber, P.J.: Robust estimation of a location parameter. In: Breakthroughs in statistics: Methodology and distribution, pp. 492–518. Springer (1992)
13. Hubert, M., Debruyne, M.: Minimum covariance determinant. Wiley interdisciplinary reviews: Computational statistics **2**(1), 36–43 (2010)
14. Jylänki, P., Vanhatalo, J., Vehtari, A.: Robust gaussian process regression with a student-t likelihood. Journal of Machine Learning Research **12**(11) (2011)
15. Kreidberg, L.: Exoplanet atmosphere measurements from transmission spectroscopy and other planet-star combined light observations. arXiv preprint arXiv:1709.05941 (2017)
16. Kuss, M.: Gaussian process models for robust regression, classification, and reinforcement learning. Ph.D. thesis, echnische Universität Darmstadt Darmstadt, Germany (2006)
17. Li, Z.Z., Li, L., Shao, Z.: Robust gaussian process regression based on iterative trimming. Astronomy and Computing **36**, 100483 (2021)
18. Mandel, K., Agol, E.: Analytic light curves for planetary transit searches. The Astrophysical Journal **580**(2), L171 (2002)
19. Maronna, R.A., Yohai, V.J.: The behavior of the stahel-donoho robust multivariate estimator. Journal of the American Statistical Association **90**(429), 330–341 (1995)
20. Mili, L., Cheniae, M., Vichare, N., Rousseeuw, P.J.: Robust state estimation based on projection statistics [of power systems]. IEEE Transactions on Power Systems **11**(2), 1118–1127 (1996)
21. Minka, T.P.: Expectation propagation for approximate bayesian inference. arXiv preprint arXiv:1301.2294 (2013)
22. Naish-Guzman, A., Holden, S.: Robust regression with twinned gaussian processes. Advances in neural information processing systems **20** (2007)
23. Neal, R.M.: Monte carlo implementation of gaussian process models for bayesian regression and classification. arXiv preprint physics/9701026 (1997)
24. Ramsey, J., Novick, M.: Plu robust bayesian decision theory: point estimation. J. Amer. Statist. Assoc. **75**, 401–407 (1980)
25. Ranjan, R., Huang, B., Fatehi, A.: Robust gaussian process modeling using em algorithm. Journal of Process Control **42**, 125–136 (2016)
26. Rue, H., Martino, S., Chopin, N.: Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. Journal of the Royal Statistical Society Series B: Statistical Methodology **71**(2), 319–392 (2009)
27. Stahel, W., Weisberg, S., Rousseeuw, P.J., van Zomeren, B.C.: Robust distances: simulations and cutoff values. In: Directions in Robust Statistics and Diagnostics: Part II. pp. 195–203. Springer (1991)
28. Stahel, W.A.: Robuste schätzungen: infinitesimale optimalität und schätzungen von kovarianzmatrizen. Ph.D. thesis, ETH Zurich (1981)
29. Tierney, L., Kadane, J.B.: Accurate approximations for posterior moments and marginal densities. Journal of the american statistical association **81**(393), 82–86 (1986)
30. Vanhatalo, J., Jylänki, P., Vehtari, A.: Gaussian process regression with student-t likelihood. Advances in neural information processing systems **22** (2009)
31. West, M.: Outlier models and prior distributions in bayesian linear regression. Journal of the Royal Statistical Society Series B: Statistical Methodology **46**(3), 431–439 (1984)

32. Williams, C.: Computing with infinite networks. Advances in neural information processing systems **9** (1996)