

---

# Efficient Bayesian Inference for a Gaussian Process Density Model

---

**Christian Donner\***

Artificial Intelligence Group  
Technische Universität Berlin  
christian.donner@bccn-berlin.de

**Manfred Opper**

Artificial Intelligence Group  
Technische Universität Berlin  
manfred.opper@tu-berlin.de

## Abstract

We reconsider a nonparametric density model based on Gaussian processes. By augmenting the model with latent Pólya–Gamma random variables and a latent marked Poisson process we obtain a new likelihood which is conjugate to the model’s Gaussian process prior. The augmented posterior allows for efficient inference by Gibbs sampling and an approximate variational mean field approach. For the latter we utilise sparse GP approximations to tackle the infinite dimensionality of the problem. The performance of both algorithms and comparisons with other density estimators are demonstrated on artificial and real datasets with up to several thousand data points.

## 1 INTRODUCTION

Gaussian processes (GP) provide highly flexible non-parametric prior distributions over functions [1]. They have been successfully applied to various statistical problems such as e.g. regression [2], classification [3], point processes [4] or the modelling of dynamical systems [5, 6]. Hence, it would seem natural to apply Gaussian processes also to density estimation which is one of the most basic statistical problems. GP density estimation, however, is a nontrivial task: Typical realisations of a GP do not respect non-negativity and normalisation of a probability density. Hence, functions drawn from a GP prior have to be passed through a nonlinear squashing function and the results have to be normalised subsequently to model a density. These operations make the corresponding posterior distributions non-Gaussian. Moreover, likelihoods depend on all the infinitely many

GP function values in the domain rather than on the finite number of function values at observed data points. Since analytical inference is impossible, [7] introduced an interesting Markov chain Monte–Carlo sampler which allows for (asymptotically) exact inference for a Gaussian process density model, where the GP is passed through a sigmoid link function.<sup>1</sup> The approach is able to deal with the infinite dimensionality of the model, because the sampling of the GP variables is reduced to a finite dimensional problem by a point process representation. However, since the likelihood of the GP variables is not conjugate to the prior, the method has to resort to a time-consuming Metropolis–Hastings approach. In this paper we will use recent results on representing the sigmoidal squashing function as an infinite mixture of Gaussians involving Pólya–Gamma random variables [9] to augment the model in such a way that the model becomes tractable by a simpler Gibbs sampler. The new model structure allows also for a much faster variational Bayesian approximation.

The paper is organised as follows: Sec. 2 introduces the GP density model, followed by an augmentation scheme that makes its likelihood conjugate to the GP prior. With this model representation we derive two efficient Bayesian inference algorithms in Sec. 3, namely an exact Gibbs sampler and an approximate, fast variational Bayes algorithm. The performance of both algorithms is demonstrated in Sec. 4 on artificial and real data. Finally, Sec. 5 discusses potential extensions of the model.

## 2 GAUSSIAN PROCESS DENSITY MODEL

The generative model proposed by [7] constructs densities over some  $d$ -dimensional data space  $\mathcal{X}$  to be of the

---

\*Also affiliated with Bernstein Center for Computational Neuroscience.

---

<sup>1</sup>See [8] for an alternative model allowing, however, only for approximate inference schemes.

form

$$\rho(\mathbf{x}|g) = \frac{\sigma(g(\mathbf{x}))\pi(\mathbf{x})}{\int_{\mathcal{X}} \sigma(g(\mathbf{x}))\pi(\mathbf{x})d\mathbf{x}}. \quad (1)$$

$\pi(\mathbf{x})$  defines a (bounded) base probability measure over  $\mathcal{X}$ , which is usually taken from a fixed parametric family. The denominator ensures normalisation  $\int_{\mathcal{X}} \rho(\mathbf{x}|g)d\mathbf{x} = 1$ . The choice of  $\pi(\mathbf{x})$  is important as will be discussed Sec. 5. A prior distribution over densities is introduced by assuming a Gaussian process prior [1] over the function  $g(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ . The GP is defined by a mean function  $\mu(\mathbf{x})$  (in this paper, we consider only constant mean functions  $\mu(\mathbf{x}) = \mu_0$ ) and covariance kernel  $k(\mathbf{x}, \mathbf{x}')$ . Finally,  $\sigma(z) = \frac{1}{1+e^{-z}}$  is the sigmoid function, which guarantees that the density is non-negative and bounded.

In Bayesian inference, the posterior distribution of  $g$  given observed data  $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$  with  $\mathbf{x} \in \mathcal{X}$  is computed from the GP prior  $p(g)$  and the likelihood as

$$p(g|\mathcal{D}) \propto p(\mathcal{D}|g)p(g).$$

The likelihood is given by

$$p(\mathcal{D}|g) = \frac{\prod_{n=1}^N \sigma(g(\mathbf{x}_n))\pi(\mathbf{x}_n)}{\left(\int_{\mathcal{X}} \sigma(g(\mathbf{x}))\pi(\mathbf{x})d\mathbf{x}\right)^N}. \quad (2)$$

Practical inference for this problem, however, is non-trivial, because (i) the posterior is non-Gaussian and (ii) the likelihood involves an integral of  $g$  over the whole space. Thus, in contrast to simpler problems such as GP regression or classification, it is impossible to reduce inference to finite dimensional integrals. To circumvent the problem that the likelihood is not conjugate to the GP prior, [7] proposed a Metropolis-Hastings MCMC algorithm for this model. We will show in the next sections that one can augment the model with auxiliary latent random variables in such a way that the resulting likelihood is of a conjugate form allowing for a more efficient Gibbs sampler with explicit conditional probabilities.

## 2.1 LIKELIHOOD AUGMENTATION

To obtain a likelihood which is conjugate to the GP  $p(g)$  we require that it assumes a Gaussian form in  $g$ .

**Representing the denominator** As a starting point, we follow [10] and use the representation

$$\frac{1}{z^N} = \frac{\int_0^\infty \lambda^{N-1} e^{-\lambda z} d\lambda}{\Gamma(N)},$$

where  $\Gamma(\cdot)$  is the gamma function. Identifying  $z = \int_{\mathcal{X}} \sigma(g(\mathbf{x}))\pi(\mathbf{x})d\mathbf{x}$  in Eq. (2) we can rewrite the

likelihood as  $p(\mathcal{D}|g) = \int_0^\infty p(\mathcal{D}, \lambda|g)d\lambda$  where

$$p(\mathcal{D}, \lambda|g) \propto \exp\left(-\int_{\mathcal{X}} \lambda \sigma(g(\mathbf{x}))\pi(\mathbf{x})d\mathbf{x}\right) \times p(\lambda) \prod_{n=1}^N \lambda \sigma(g(\mathbf{x}_n))\pi(\mathbf{x}_n), \quad (3)$$

with the improper prior  $p(\lambda) = \lambda^{-1}$  over the auxiliary latent variable  $\lambda$ . To transform the likelihood further into a form which is Gaussian in  $g$ , we utilise a representation of the sigmoid function as a scale mixture of Gaussians.

### Pólya-Gamma representation of sigmoid function

As discovered by [9], the inverse hyperbolic cosine can be represented as an infinite mixture of scaled Gaussians

$$\cosh^{-b}(z/2) = \int_0^\infty e^{-\frac{z^2}{2}\omega} p_{\text{PG}}(\omega|b, 0)d\omega,$$

where  $p_{\text{PG}}(\omega|b, 0)$  is the *Pólya-Gamma density* of random variable  $\omega \in \mathbb{R}^+$ . Moments of those densities can be easily computed [9]. Later, we will also use the *tilted* Pólya-Gamma densities defined as

$$p_{\text{PG}}(\omega|b, c) \propto \exp\left(-\frac{c^2}{2}\omega\right) p_{\text{PG}}(\omega|b, 0). \quad (4)$$

These definitions allows for a Gaussian representation of the sigmoid function as

$$\sigma(z) = \frac{e^{z/2}}{2 \cosh(z/2)} = \int_0^\infty e^{f(\omega, z)} p_{\text{PG}}(\omega|1, 0)d\omega \quad (5)$$

with  $f(\omega, z) = \frac{z}{2} - \frac{z^2}{2}\omega - \ln 2$ . This result will be used to transform the products over observations  $\sigma(g(\mathbf{x}_n))$  in the likelihood (3) into a Gaussian form.

We will next deal with the first term in the likelihood (3) which contains the integral over  $\mathbf{x}$ . For this part of the model we will derive a point process representation which can be understood as a generalisation of the approach of [7].

**Marked-Poisson representation** Utilising the sigmoid property  $\sigma(z) = 1 - \sigma(-z)$  and the Pólya-Gamma representation (5) the integral in the exponent of Eq. (3) can be written as a double integral

$$\begin{aligned} -\int_{\mathcal{X}} \lambda \sigma(g(\mathbf{x}))\pi(\mathbf{x})d\mathbf{x} &= \\ \int_{\mathcal{X}} (\sigma(-g(\mathbf{x})) - 1)\lambda\pi(\mathbf{x})d\mathbf{x} &= \\ \int_{\mathcal{X}} \int_{\mathbb{R}^+} \left(e^{f(\omega, -g(\mathbf{x}))} - 1\right) \lambda\pi(\mathbf{x})p_{\text{PG}}(\omega|1, 0)d\omega d\mathbf{x} \end{aligned}$$

Next we will use a result for the characteristic function of a Poisson process. Following [11, chap. 3] one has

$$\mathbb{E}_\phi \left[ \prod_{z \in \Pi} h(z) \right] = \exp \left( \int_{\mathcal{Z}} (h(z) - 1) \phi(z) dz \right). \quad (6)$$

$h(\cdot)$  is a function on a space  $\mathcal{Z}$  and the expectation is over a Poisson process  $\Pi$  with rate function  $\phi(z)$ .  $\Pi = \{z_m\}_{m=1}^M$  denotes a random set of points on the space  $\mathcal{Z}$ . To apply this result to our problem, we identify  $\mathcal{Z} = \mathcal{X} \times \mathbb{R}^+$ ,  $z = (\mathbf{x}, \omega)$  and  $\phi_\lambda(\mathbf{x}, \omega) = \lambda \pi(\mathbf{x}) p_{\text{PG}}(\omega|1, 0)$  and finally  $h(z) = e^{f(\omega, -g(\mathbf{x}))}$  to rewrite the exponential in Eq. (3) as

$$e^{-\int_{\mathcal{X}} \lambda \sigma(g(\mathbf{x})) \pi(\mathbf{x}) d\mathbf{x}} = \mathbb{E}_{\phi_\lambda} \left[ \prod_{(\omega, \mathbf{x}) \in \Pi} e^{f(\omega, -g(\mathbf{x}))} \right]. \quad (7)$$

By substituting Eq. (5) and (7) into Eq.(3) we obtain the final augmented form of the likelihood of Eq. (2) which is one of the main results of our paper.

$$p(\mathcal{D}, \lambda, \Pi, \omega_N | g) \propto \prod_{n=1}^N \phi_\lambda(\mathbf{x}_n, \omega_n) e^{f(\omega_n, g(\mathbf{x}_n))} \times p_{\phi_\lambda}(\Pi | \lambda) p(\lambda) \prod_{(\omega, \mathbf{x}) \in \Pi} e^{f(\omega, -g(\mathbf{x}))}, \quad (8)$$

with  $p_\phi(\Pi | \lambda)$  being the density over a Poisson process  $\Pi = \{(\mathbf{x}_m, \omega_m)\}_{m=1}^M$  in the augmented space  $\mathcal{X} \times \mathbb{R}^+$  with intensity  $\phi_\lambda(\mathbf{x}, \omega)$ .<sup>2</sup> This new process can be identified as a *marked Poisson process* [11, chap. 5], where the events  $\{\mathbf{x}_m\}_{m=1}^M$  in the original data space  $\mathcal{X}$  follow a Poisson process with rate  $\lambda \pi(\mathbf{x})$ . Then, on each event  $\mathbf{x}_m$  an independent *mark*  $\omega_m \sim p_{\text{PG}}(\omega_m | b, 0)$  is drawn at random from the Pólya–Gamma density. Finally,  $\omega_N = \{\omega_n\}_{n=1}^N$  is the set of latent Pólya–Gamma variables which result from the sigmoid augmentation at the observations  $\mathbf{x}_n$ .

**Augmented posterior over GP density** With Eq. (8) we obtain the joint posterior over the GP  $g$ , the rate scaling  $\lambda$ , the marked Poisson process  $\Pi$ , and the Pólya–Gamma variables at the observations  $\omega_N$  as

$$p(\omega_N, \Pi, \lambda, g | \mathcal{D}) \propto p(\mathcal{D}, \omega_N, \Pi, \lambda | g) p(g). \quad (9)$$

In the following, this new representation will be used to derive two inference algorithms.

<sup>2</sup>Densities such as  $p_{\phi_\lambda}(\Pi | \lambda)$  could be understood as the Radon–Nykodym derivative [12] of the corresponding probability measure with respect to some fixed dominating measure. However, we will not need an explicit form here.

### 3 INFERENCE

We will first derive an efficient Gibbs sampler which (asymptotically) solves the inference problem exactly, and then a variational mean-field algorithm, which only finds an approximate solution, but in a much faster time.

#### 3.1 GIBBS SAMPLER

Gibbs sampling [13] generates samples from the posterior by creating a Markov chain, where at each time, a block of variables is drawn from the conditional posterior given all the other variables. Hence, to perform Gibbs sampling, we have to derive these conditional distributions for each set of variables from Eq. (9). Most of the following results are easily obtained by direct inspection. The only non-trivial case is the conditional distribution over the latent point process  $\Pi$ .

**Pólya-Gamma variables at observations** The conditional posterior over the set of Pólya–Gamma variables  $\omega_N$  depends only on the function  $g$  at the observations  $\{g(\mathbf{x}_n)\}_{n=1}^N$  and turns out to be

$$p(\omega_N | g) = \prod_{n=1}^N p_{\text{PG}}(\omega_n | 1, g(\mathbf{x}_n)), \quad (10)$$

where we have used the definition of a tilted Pólya–Gamma density in Eq. (4). This density can be efficiently sampled by methods developed by [9]<sup>3</sup>.

**Rate scaling** The rate scaling  $\lambda$  has a conditional Gamma density given by

$$\text{Gamma}(\lambda | \alpha, 1) = \frac{(\lambda)^{\alpha-1} e^{-\lambda}}{\Gamma(\alpha)}. \quad (11)$$

with  $\alpha = |\Pi| + N = M + N$ . Hence, the posterior is dependent on the number of observations and the number on events of the marked Poisson process  $\Pi$ .

**Posterior Gaussian process** Due to the form of the augmented likelihood the conditional posterior for the GP  $g_{N+M}$  at the observations  $\{\mathbf{x}_n\}_{n=1}^N$  and the latent events  $\{\mathbf{x}_m\}_{m=1}^M$  is a multivariate Gaussian density

$$p(g_{N+M} | \Pi, \omega_N) = \mathcal{N}(\boldsymbol{\mu}_{N+M}, \Sigma_{N+M}), \quad (12)$$

with covariance matrix  $\Sigma_{N+M} = [D + K_{N+M}^{-1}]^{-1}$ . The diagonal matrix  $D$  has its first  $N$  entries given by  $\omega_N$  followed by  $M$  entries being  $\{\omega_m\}_{m=1}^M$ . The mean is  $\boldsymbol{\mu}_{N+M} = \Sigma_{N+M} \left[ \mathbf{u} + K_{N+M}^{-1} \boldsymbol{\mu}_0^{(N+M)} \right]$ , where the

<sup>3</sup>The sampler implemented by [14] is used for this work.

first  $N$  entries of  $N + M$  dimensional vector  $\mathbf{u}$  are  $1/2$  and the rest are  $-1/2$ .  $K_{N+M}$  is the prior covariance kernel matrix of the GP evaluated at the observed points  $\mathbf{x}_n$  and the latent events  $\mathbf{x}_m$ , and  $\boldsymbol{\mu}_0^{(N+M)}$  is an  $N + M$  dimensional vector with all entries being  $\mu_0$ .

The predictive conditional posterior for the GP for any set of points in  $\mathcal{X}$  is simply given via the conditional prior  $p(g|\mathbf{g}_{N+M})$ , which has a well known form and can be found in [1].

**Sampling the latent marked point process** We easily find that the conditional posterior of the marked point process is given by

$$p(\Pi|g, \lambda) = \frac{\prod_{\omega, \mathbf{x} \in \Pi} e^{f(\omega, -g(\mathbf{x}))} p_{\phi_\lambda}(\Pi|\lambda)}{\exp\left(\int_{\mathcal{X} \times \mathbb{R}^+} (e^{f(\omega, -g(\mathbf{x}))} - 1) \phi_\lambda(\mathbf{x}, \omega) d\omega d\mathbf{x}\right)}, \quad (13)$$

where the form of the normalising denominator is obtained using Eq. (6). By computing the characteristic function of this conditional point process (see App. A) we can show that it is again a marked Poisson process with intensity

$$\Lambda(\mathbf{x}, \omega) = \lambda \pi(\mathbf{x}) \sigma(-g(\mathbf{x})) p_{\text{PG}}(\omega|1, g(\mathbf{x})). \quad (14)$$

To sample from this process we first draw Poisson events  $\mathbf{x}_m$  in the original data space  $\mathcal{X}$  using the rate  $\int_{\mathbb{R}^+} \Lambda(\mathbf{x}, \omega) d\omega = \lambda \pi(\mathbf{x}) \sigma(-g(\mathbf{x}))$  [11, chap. 5]. Subsequently for each event  $\mathbf{x}_m$  a mark  $\omega_m$  is generated from the conditional density  $\omega_m \sim p_{\text{PG}}(\omega|1, g(\mathbf{x}_m))$ .

To sample the events  $\{\mathbf{x}_m\}_{m=1}^M$ , we use the well known approach of *thinning* [4]. We note, that the rate is upper bounded by the base measure  $\lambda \pi(\mathbf{x})$ . Hence, we first generate points  $\tilde{\mathbf{x}}_m$  from a Poisson process with intensity  $\lambda \pi(\mathbf{x})$ . This is easily achieved by noting that the required number  $M_{\text{max}}$  of such events is Poisson distributed with mean parameter  $\int_{\mathcal{X}} \lambda \pi(\mathbf{x}) d\mathbf{x} = \lambda$ . The position of the events can then be obtained by sampling  $\{\tilde{\mathbf{x}}_m\}_{m=1}^{M_{\text{max}}}$  independent points from the base density  $\tilde{\mathbf{x}}_m \sim \pi(\mathbf{x})$ . These events are *thinned* by keeping each point  $\tilde{\mathbf{x}}_m$  with probability  $\sigma(-g(\tilde{\mathbf{x}}_m))$ . The kept events constitute the final set  $\{\mathbf{x}_m\}_{m=1}^M$ .

**Sampling hyperparameters** In this work we will consider specific functional forms for the kernel  $k(\mathbf{x}, \mathbf{x}')$  and the base measure  $\pi(\mathbf{x})$  which are parametrised by hyperparameters  $\boldsymbol{\theta}_k$  and  $\boldsymbol{\theta}_\pi$ . These will be sampled by a Metropolis-Hastings method [15]. The GP prior mean  $\mu_0$  can be directly sampled from the conditional posterior given  $\mathbf{g}_{M+N}$ . In this work, the hyperparameters are sampled every  $v = 10$  step. Different choices of  $v$  might yield faster convergence of the Markov Chain. Pseudo code for the Gibbs sampler is provided in Alg. 1.

---

### Algorithm 1: Gibbs sampler for GP density model.

---

**Init:**  $\{\mathbf{x}_m\}_{m=1}^M$ ,  $\mathbf{g}_{N+M}$ ,  $\lambda$ , and  $\boldsymbol{\theta}_k, \boldsymbol{\theta}_\pi, \mu_0$   
**1 for** *Length of Markov chain* **do**  
**2**     **Sample PG variables at**  $\{\mathbf{x}_m\}$ :  $\omega_N \sim \text{Eq. (10)}$   
**3**     **Sample latent Poisson process:**  $\Pi \sim \text{Eq. (13)}$   
**4**     **Sample rate scaling:**  $\lambda \sim \text{Eq. (11)}$   
**5**     **Sample GP:**  $\mathbf{g}_{N+M} \sim \text{Eq. (12)}$   
**6**     **Sample hyperparameters:** Every  $v^{\text{th}}$  sample with Metropolis–Hastings  
**7 end**

---

## 3.2 VARIATIONAL BAYES

While expected to be more efficient than a Metropolis-Hastings sampler based on the unaugmented likelihood [7], the Gibbs sampler is practically still limited. The main computational bottleneck comes from the sampling of the conditional Gaussian over function values of  $g$ . The computation of the covariances requires the inversion of matrices of dimensions  $N + M$ , with a complexity  $\mathcal{O}((N + M)^3)$ . Hence the algorithm does not only become infeasible, when we have many observations, i.e when  $N$  is large, but also if the sampler requires many thinned events, i.e. if  $M$  is large. This can happen in particular for bad choices of the base measure  $\pi(\mathbf{x})$ . In the following, we introduce a variational Bayes algorithm [16], which solves the inference problem approximately, but with a complexity which scales linearly in the data size and is independent of structure.

**Structured mean–field approach** The idea of variational inference [16] is to approximate an intractable posterior  $p(Z|\mathcal{D})$  by a simpler distribution  $q(Z)$  from a tractable family.  $q(Z)$  is optimised by minimising the Kullback-Leibler divergence between  $q(Z)$  and  $p(Z|\mathcal{D})$  which is equivalent to maximising the so called *variational lower bound* (sometimes also called ELBO for evidence lower bound) given by

$$\mathcal{L}(q(Z)) = \mathbb{E}_Q \left[ \ln \frac{p(Z, \mathcal{D})}{q(Z)} \right] \leq \ln p(\mathcal{D}), \quad (15)$$

where  $Q$  denotes the probability measure with density  $q(Z)$ . A common approach for variational inference is a structured mean–field method, where dependencies between sets of variables are neglected. For the problem at hand we assume that

$$q(\omega_N, \Pi, g, \lambda) = q_1(\omega_N, \Pi) q_2(g, \lambda). \quad (16)$$

A standard result for the variational mean–field approach shows that the optimal independent factors, which max-

imise the lower bound in Eq. (15) are given by

$$\ln q_1(\boldsymbol{\omega}_N, \Pi) = \mathbb{E}_{Q_2} [\ln p(\mathcal{D}, \boldsymbol{\omega}_N, \Pi, \lambda, g)] + \text{const.}, \quad (17)$$

$$\ln q_2(g, \lambda) = \mathbb{E}_{Q_1} [\ln p(\mathcal{D}, \boldsymbol{\omega}_N, \Pi, \lambda, g)] + \text{const.} \quad (18)$$

By inspecting Eq. (9), (17), and (18) it turns out that the densities of all four sets of variables factorise as

$$\begin{aligned} q_1(\boldsymbol{\omega}_N, \Pi) &= q_1(\boldsymbol{\omega}_N)q_1(\Pi), \\ q_2(g, \lambda) &= q_2(g)q_2(\lambda). \end{aligned}$$

We will optimise the factors by a straightforward iterative algorithm, where each factor is updated given expectations over the others based on the previous step. Hence, the lower bound in Eq. (15) is increased in each step. Again we will see that the augmented likelihood in Eq. (8) allows for analytic solutions of all required factors.

**Pólya–Gamma variables at the observations** Similar to the Gibbs sampler, the variational posterior of the Pólya–Gamma variables at the observations is a product of tilted Pólya–Gamma densities given by

$$q_1(\boldsymbol{\omega}_N) = \prod_{n=1}^N p_{\text{PG}}(\omega_n | 1, c_n), \quad (19)$$

with  $c_n = \sqrt{\mathbb{E}_{Q_2} [g(\mathbf{x}_n)^2]}$ . The only difference is, that the second argument of  $p_{\text{PG}}$  depends on the expectation of the square of  $g(\mathbf{x}_n)$ .

**Posterior marked Poisson process** Similar to the corresponding result for the Gibbs sampler we can show<sup>4</sup> that the optimal latent point process  $\Pi$  is a Poisson process with rate given by

$$\begin{aligned} \Lambda_1(\mathbf{x}, \omega) &= \lambda_1 \pi(\mathbf{x}) \sigma(-c(\mathbf{x})) p_{\text{PG}}(\omega | 1, c(\mathbf{x})) \\ &\times e^{(c(\mathbf{x}) - g_1(\mathbf{x}))/2} \end{aligned} \quad (20)$$

with  $\lambda_1 = e^{\mathbb{E}_{Q_2} [\ln \lambda]}$ ,  $c(\mathbf{x}) = \sqrt{\mathbb{E}_{Q_2} [f(\mathbf{x})^2]}$ , and  $g_1(\mathbf{x}) = \mathbb{E}_{Q_2} [g(\mathbf{x})]$ . Note also the similarity to the Gibbs sampler in Eq. (14).

**Optimal posterior for rate scaling** The posterior for the rate scaling  $\lambda$  is a Gamma distribution given by

$$q_2(\lambda) = \text{Gamma}(\lambda | \alpha_2, 1) = \frac{\lambda^{\alpha_2 - 1} e^{-\lambda}}{\Gamma(\alpha_2)}, \quad (21)$$

where  $\alpha_2 = N + \mathbb{E}_{Q_1} [\sum_{\mathbf{x}' \in \Pi} \delta(\mathbf{x} - \mathbf{x}')]$ , and  $\mathbb{E}_{Q_1} [\sum_{\mathbf{x}' \in \Pi} \delta(\mathbf{x} - \mathbf{x}')] = \int_{\mathcal{X}} \int_{\mathbb{R}^+} \Lambda_1(\mathbf{x}, \omega) d\omega d\mathbf{x}$ , and  $\delta(\cdot)$  is the Dirac delta function. The integral is solved by importance sampling as will be explained (see Eq. (25)).

<sup>4</sup>The proof is similar to the one from App. A.

**Approximation of GP via sparse GP** The optimal variational form for the posterior  $g$  is a GP given by

$$q_2(g) \propto e^{U(g)} p(g),$$

where  $U(g) = \mathbb{E}_{Q_1} [\ln p(\mathcal{D}, \boldsymbol{\omega}_N, \Pi, \lambda | g)]$  results in the Gaussian log-likelihood

$$U(g) = -\frac{1}{2} \int_{\mathcal{X}} A(\mathbf{x}) g(\mathbf{x})^2 d\mathbf{x} + \int_{\mathcal{X}} B(\mathbf{x}) g(\mathbf{x}) d\mathbf{x} + \text{const.}$$

with

$$\begin{aligned} A(\mathbf{x}) &= \sum_{n=1}^N \mathbb{E}_{Q_1} [\omega_n] \delta(\mathbf{x} - \mathbf{x}_n) + \int_{\mathbb{R}^+} \omega \Lambda_1(\mathbf{x}, \omega) d\omega, \\ B(\mathbf{x}) &= \frac{1}{2} \sum_{n=1}^N \delta(\mathbf{x} - \mathbf{x}_n) - \frac{1}{2} \int_{\mathbb{R}^+} \Lambda_1(\mathbf{x}, \omega) d\omega. \end{aligned}$$

For general GP priors, this free form optimum is intractable by the fact that the likelihood depends on  $g$  at infinitely many points. Hence, we resort to an additional approximation which makes the dimensionality of the problem again finite. The well known framework of *sparse* GPs [17, 18, 19] turns out to be useful in this case. This has been introduced for likelihoods with large, but finite dimensional likelihoods [19, 20] and later generalised to infinite dimensional problems [21, 22]. The sparse approximation assumes a variational posterior of the form

$$q_2(g) = p(g | \mathbf{g}_s) q_2(\mathbf{g}_s),$$

where  $\mathbf{g}_s$  is the GP evaluated at a finite set of *inducing points*  $\{\mathbf{x}_l\}_{l=1}^L$  and  $p(g | \mathbf{g}_s)$  is the conditional prior. A variational optimisation yields

$$q_2(\mathbf{g}_s) \propto e^{U^s(\mathbf{g}_s)} p(\mathbf{g}_s), \quad (22)$$

where the first term can be seen as a new ‘effective’ likelihood only depending on the inducing points. This new (log) likelihood is given by

$$\begin{aligned} U^s(\mathbf{g}_s) &= \mathbb{E}_P [U(g) | \mathbf{g}_s] = \\ &= -\frac{1}{2} \int_{\mathcal{X}} A(\mathbf{x}) \tilde{g}_s(\mathbf{x})^2 d\mathbf{x} + \int_{\mathcal{X}} B(\mathbf{x}) \tilde{g}_s(\mathbf{x}) d\mathbf{x} + \text{const.}, \end{aligned}$$

with  $\tilde{g}_s(\mathbf{x}) = \mu_0 + \mathbf{k}_s(\mathbf{x})^\top K_s^{-1}(\mathbf{g}_s - \boldsymbol{\mu}_0^{(L)})$ ,  $\mathbf{k}_s(\mathbf{x})$  being an  $L$  dimensional vector, where the  $l^{\text{th}}$  entry is  $k(\mathbf{x}, \mathbf{x}_l)$  and  $K_s$  being the prior covariance matrix for all inducing points. The expectation is computed with respect to the GP prior conditioned on the sparse GP  $\mathbf{g}_s$ . We identify Eq. (22) being a multivariate normal distribution with covariance matrix

$$\Sigma_2^s = \left[ K_s^{-1} \int_{\mathcal{X}} A(\mathbf{x}) \mathbf{k}_s(\mathbf{x})^\top \mathbf{k}_s(\mathbf{x}) d\mathbf{x} K_s^{-1} + K_s^{-1} \right]^{-1}, \quad (23)$$

---

**Algorithm 2:** Variational Bayes algorithm for GP density model

---

**Init:** Inducing points,  $q_2(\mathbf{g}_s)$ ,  $q_2(\lambda)$ , and  $\theta_k, \theta_\pi, \mu_0$

1 **while**  $\mathcal{L}$  not converged **do**

2     **Update**  $q_1$

3     |     **PG distributions at observations:**  $q_1^*(\omega_N)$

       |     with Eq. (19)

4     |     **Rate of latent process:**  $\Lambda_1(\mathbf{x}, \omega)$  with Eq. (20)

5     |     **Update**  $q_2$

6     |     |     **Rate scaling:**  $\alpha_2$  with Eq. (21)

7     |     |     **Sparse GP:**  $\Sigma_2^s, \mu_2^s$  with Eq. (23), (24)

8     |     **Update**  $\theta_k, \theta_\pi, \mu_0$  **with gradient update**

9 **end**

---

and mean

$$\mu_2^s = \Sigma_2^s \left( K_s^{-1} \int_{\mathcal{X}} \mathbf{k}_s(\mathbf{x}) \tilde{B}(\mathbf{x}) d\mathbf{x} + K_s^{-1} \mu_0^{(L)} \right), \quad (24)$$

with  $\tilde{B}(\mathbf{x}) = B(\mathbf{x}) - A(\mathbf{x})(\mu_0 - \mathbf{k}_s(\mathbf{x})^\top K_s^{-1} \mu_0^{(L)})$ .

**Integrals over  $\mathbf{x}$**  The sparse GP approximation and the posterior over  $\lambda$  in Eq. (21) requires the computation of integrals of the form

$$I \doteq \int_{\mathcal{X}} \int_{\mathbb{R}^+} y(\mathbf{x}, \omega) \Lambda_1(\mathbf{x}, \omega) d\omega d\mathbf{x},$$

with specific functions  $y(\mathbf{x}, \omega)$ . For these functions, the inner integral over  $\omega$  can be computed analytically, but the outer one over the space  $\mathcal{X}$  has to be treated numerically. We approximate it via importance sampling

$$I \approx \frac{1}{R} \sum_{r=1}^R \int_{\mathbb{R}^+} y(\mathbf{x}_r, \omega_r) \frac{\Lambda_1(\mathbf{x}_r, \omega_r)}{\pi(\mathbf{x}_r)} d\omega_r, \quad (25)$$

where every sample point  $\mathbf{x}_r$  is independently drawn from the base measure  $\pi(\mathbf{x})$ .

**Updating hyperparameters** Having an analytic solution for every factor of the variational posterior in Eq. (16) we further require the optimisation of hyperparameters.  $\theta_k, \theta_\pi$  and  $\mu_0$  are optimised by maximising the lower bound in Eq. (15) (see App. B for explicit form) with a gradient ascent algorithm having an adaptive learning rate (Adam) [23]. Additional hyperparameters are the locations of inducing points  $\{\mathbf{x}_l\}_{l=1}^L$ . Half of them are drawn randomly from the initial base measure, while half of them are positioned on regions with a high density of observations found by a k-means algorithm. Pseudo code for the complete variational algorithm is provided in Alg. 2.

Python code for Alg. 1 and 2 is provided at [24].

## 4 RESULTS

To test our two inference algorithms, the Gibbs sampler and the variational Bayes algorithm (VB), we will first evaluate them on data drawn from the generative model. Then we compare both on an artificial dataset and several real datasets. We will only consider cases with  $\mathcal{X} = \mathbb{R}^d$ . To evaluate the quality of inference we consider always the logarithm of the expected test likelihood

$$\ell_{\text{test}}(\tilde{\mathcal{D}}) \doteq \ln \left( \mathbb{E} \left[ \prod_{\mathbf{x} \in \tilde{\mathcal{D}}} \rho(\mathbf{x}) \right] \right),$$

where  $\tilde{\mathcal{D}}$  is test data unknown to the inference algorithm and the expectation is over the inferred posterior measure. In practice we sample this expectation from the inferred posterior over  $g$ . Since this quantity involves an integral, that is again approximated by Eq. (25), we check that the standard deviation  $\text{std}(I)$  is less than 1% of the value of the estimated value  $I$ .

**Data from generative model.** We generate datasets according to Eq. (1), where  $g$  is drawn from the GP prior with  $\mu_0 = 0$ . As covariance kernel we assume a squared exponential throughout this work

$$k(\mathbf{x}, \mathbf{x}') = \theta_k^{(0)} \prod_{i=1}^d \exp \left( -\frac{(x_i - x'_i)^2}{2(\theta_k^{(i)})^2} \right).$$

The base measure  $\pi(\mathbf{x})$  is a standard normal density. We use the algorithm described in [7] to generate exact samples. In this section, the hyperparameters  $\theta_k, \theta_\pi$  and  $\mu_0$  are fixed to the true values for inference. Unless stated otherwise for the VB the number of inducing points is fixed to 200 and the number of integration points for importance sampling to  $5 \times 10^3$ . For the Gibbs sampler, we sample a Markov chain of  $5 \times 10^3$  samples after a burn-in period of  $2 \times 10^3$  samples.

In Fig. 1 we see a 1 dimensional example dataset, where both inference algorithms recover well the structure of the underlying density. The inferred posterior means are barely distinguishable. However, evaluating the inferred densities on an unseen test set, we note that the Gibbs sampler performs slightly better. Of course, this is expected since the sampler provides exact inference for the generative model and should (on average) not be outperformed by the approximate VB as long as the sampled Markov chain is long enough. In Fig. 1 (bottom left) we see that only 13 iterations of the VB are required to meet the convergence criterion. For Markov chain samplers to be efficient, correlations between samples should decay quickly. Fig. 1 (bottom middle) shows the autocorrelation of  $\ell_{\text{test}}$ , which was evaluated at each sample of the

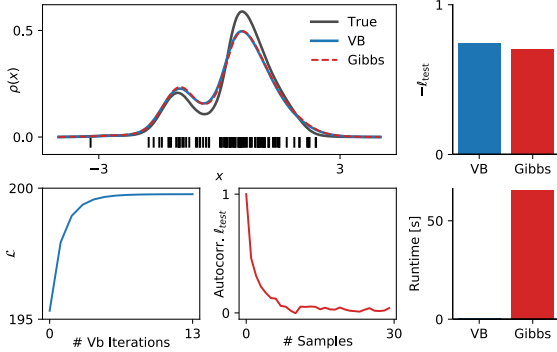


Figure 1: **1D data from the generative model.** Data consist of 100 samples from the underlying density sampled from the GP density model. **Upper left:** True density (black line), data (black vertical bars), mean posterior density inferred by Gibbs sampler (red dashed line) and VB algorithm (blue line). **Upper right:** Negative log expected test likelihood of Gibbs and VB inferred posterior. **Lower left:** Variational lower bound as function of iterations of the VB algorithm. **Lower middle:** Autocorrelation of test likelihood as function of Markov chain samples obtained from Gibbs sampler. **Lower right:** Runtime of the two algorithms (VB took 0.3 s).

Dim	# points	Gibbs		VB	
		$\ell_{\text{test}}$	$T$ [s]	$\ell_{\text{test}}$	$T$ [s]
1	50	-146.9	30.1	-149.2	1.13
2	100	-257.0	649.9	-260.2	2.03
2	200	-285.3	546.1	-289.6	1.41
6	400	-823.9	4667	-822.2	0.89

Table 1: **Performance of Gibbs sampler and VB** on different datasets sampled from generative model.  $\ell_{\text{test}}$  was evaluated on a unknown test set including 50 samples. In addition, runtime  $T$  is reported in seconds.

Markov chain. After about 10 samples the correlations reach a plateau close to 0, demonstrating excellent mixing properties of the sampler. Comparing the run time of both algorithms, VB (0.3 s) outperforms the sampler  $\sim 1$  min by more than 2 orders of magnitude.

To demonstrate the inference for more complicated problems, 2 dimensional data are generated with 200 samples (Fig. 2). The posterior mean densities inferred by both algorithms capture the structure well. As before, the log expected test likelihood is larger for the Gibbs sampler ( $\ell_{\text{test}} = -296.2$ ) compared to VB ( $\ell_{\text{test}} = -306.0$ ). However, the Gibbs sampler took  $> 20$  min while the VB required only 1.8 s to obtain the result.

In Tab. 1 we show results for datasets with different size and different dimensionality. The results confirm that the




Figure 2: **2D data from generative model.** **Right:** 200 samples from the underlying two dimensional density. **Middle:** Posterior mean of Gibbs sampler inferred density. **Right:** Posterior mean of VB inferred density.




Figure 3: **Comparison to other density estimation methods on artificial 2D data.** Training data consist of 100 data points uniformly distributed on a circle (1.5 radius) and additional Gaussian noise (0.2 std.). **From left to right:** The posterior mean inferred by Gibbs sampler and VB algorithm, followed by density estimation using KDE and GMM.

run time for the Gibbs sampler scales strongly with size and dimensionality of a problem, while the VB algorithm seems relatively unaffected in this regard. However, the VB is in general outperformed by the sampler in terms of expected test likelihood or in the same range. Note, that the runtime of the Gibbs sampler does not solely depend on the number of observed data points  $N$  (compare data set 2 and 3 in Tab. 1). As discussed earlier this can happen, when the base measure  $\pi(\mathbf{x})$  is very different from the target density  $\rho(\mathbf{x})$  resulting in many latent Poisson events (i.e.  $M$  is large).

**Circle data** In the following, we compare the GP density model and its two inference algorithms with two alternative density estimation methods. These are given by a kernel density estimator (KDE) with a Gaussian kernel and a Gaussian mixture model (GMM) [25]. The free parameters of these models (kernel bandwidth for KDE and number of components for GMM) are optimised by 10-fold cross-validation. Furthermore, GMM is initialised 10 times and the best result is reported. For the GP density model a Gaussian density is assumed as base measure  $\pi(\mathbf{x})$ , and hyperparameters  $\theta_{\pi}$ ,  $\theta_k$ , and  $\mu_0$  are now optimised. Similar to [7] we consider 100 samples uniformly drawn from a circle with additional Gaussian

	Gibbs	VB	KDE	GMM
$\ell_{\text{test}}$	-220.31	-230.53	-228.43	-237.34

Table 2: Log expected test likelihood for circle data.




Figure 4: **Performance on ‘Egyptian Skulls’ dataset [26].** 100 training points and 4 dimensions. Bar height shows average negative log test likelihood obtained by five random permutations of training and test set and points mark single permutation results.

noise. The inferred densities (only the mean of the posterior for Gibbs and VB) are shown in Fig. 3. Both GP density methods recover well the structure of the data, but the VB seems to overestimate the width of the Gaussian noise compared to the Gibbs sampler. While the KDE also recovers relatively well the data structure the GMM fails in this case. This is also reflected on the log expected test likelihoods (Tab. 2).

**Real data sets** The ‘Egyptian Skulls’ dataset [26] contains 150 data points in 4 dimensions. 100 training points are randomly selected and performance is evaluated on the remaining ones. Before fitting data is whitened. Base measure and fitting procedure for all algorithms are the same as for the circular data. Furthermore, fitting is done for 5 random permutations of training and test set. The results in Fig. 4 show that both algorithms for the GP density model outperform the two other ones on this dataset.

Often practical problems may consist of many more data points and dimensions. As discussed, the Gibbs sampler is not practical for such kind of problems, while the VB could handle larger amounts of data. Unfortunately, the sparsity assumption and the integration via importance sampling is expected to become poorer with increasing number of dimensions. Noting, however, that the ‘effective’ dimensionality in our model is determined by the base measure  $\pi(\mathbf{x})$ , one can circumvent this problem by an educated choice of  $\pi(\mathbf{x})$  if data  $\mathcal{D}$  lie in a submanifold of the high dimensional space  $\mathcal{X}$ .

We employ this strategy by first fitting a GMM to the problem and then utilising the fit as base measure. In Fig. 5 we consider 3 different datasets<sup>5</sup> to test this pro-

<sup>5</sup>Only real valued dimensions are considered and for the




Figure 5: **Application on higher dimensional data with many data points.** The improvement on log expected test likelihood  $\ell_{\text{test}}$  per test point compared to GMM, when using same as base measure  $\pi(\mathbf{x})$  for the VB inference. **From top to bottom:** ‘Forest Fire’ dataset [27, 28] (400 training points, 117 test points, 5 dim.), ‘Thyroid’ dataset [29] ( $3 \times 10^3$ , 772, 6), ‘Wine’ dataset [27] ( $6 \times 10^3$ , 498, 9). Bars mark improvement on average of random permutations of training and test set while points mark single runs.

cedure. As in Fig. 4, fitting is repeated 5 times for random permutations if training and test set. For the ‘Thyroid’ dataset, one of the 5 fits is excluded, because the importance sampling yielded poor approximation  $\text{std}(I) > I \times 10^{-2}$ . The training sets contain 400 to 6000 data points with 5 to 9 dimensions. The results for KDE are not reported, since it is always outperformed by the GMM. Fig. 5 demonstrates combining the GMM and VB algorithm results in an improvement of the log test likelihood  $\ell_{\text{test}}$  compared to using only GMM. Average relative improvements of  $\ell_{\text{test}}$  are 8.9% for ‘Forest Fire’, 4.1% for ‘Thyroid’, and 1.1% for ‘Wine’ dataset.

## 5 DISCUSSION

We have shown how inference for a nonparametric, GP based, density model can be made efficient. In the following we would like to discuss various possible extensions but also limitations of our approach.

**Choice of base measure** As we have shown for applications to real data, the choice of the base measure is quite important, especially for the sampler and for high dimensional problems. While many datasets might favour a normal distribution as base measure, problems with outliers might favour fat tailed densities. In general, any density which can be evaluated on the data space  $\mathcal{X}$  and which allows for efficient sampling, is a valid choice as base measure  $\pi(\mathbf{x})$  in our inference approach for the GP density model. Any powerful density estima-

‘forest fire’ dataset dimensions are excluded, where data have more than half 0 entries.



tor which fulfils this condition could provide a base measure which could then potentially be improved by the GP model. It would e.g. be interesting to apply this idea to neural networks [30, 31] based estimators. Other generalisations of our model could consider alternative data spaces  $\mathcal{X}$ . One might e.g. think of specific discrete and structured sets  $\mathcal{X}$  for which appropriate Gaussian processes could be defined by suitable Mercer kernels.

**Big data & high dimensionality** Our proposed Gibbs sampler suffers from cubic scaling in the number of data points and is found to be already impractical for problems with hundreds of observations. This could potentially be tackled by using sparse (approximate) GP methods for the sampler (see [32] for a potential approach). On the other hand, the proposed VB algorithm scales only linearly with the training set size and can be applied to problems with several thousands of observations. The integration of stochastic variational inference into our method could potentially increase this limit [33].

Potential limitations of the GP density model are given by high dimensional problems. If approached naively, the combination of the sparse GP approximation and the numerical integration using importance sampling is expected to yield bad approximations in such cases.<sup>6</sup> If the data is concentrated on a low dimensional submanifold of the high-dimensional space, one could still try to combine our method with other density estimators providing a base measure  $\pi(\mathbf{x})$  that is adapted to this submanifold, to allow for tractable GP inference.

## Acknowledgements

CD was supported by the Deutsche Forschungsgemeinschaft (GRK1589/2) and partially funded by Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1294 “Data Assimilation”, Project (A06) “Approximative Bayesian inference and model selection for stochastic differential equations (SDEs)”.

## References

- [1] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [2] Christopher KI Williams and Carl Edward Rasmussen. Gaussian processes for regression. In *Advances in neural information processing systems*, pages 514–520, 1996.
- [3] Hannes Nickisch and Carl Edward Rasmussen. Approximations for binary gaussian process classification. *Journal of Machine Learning Research*, 9(Oct):2035–2078, 2008.
- [4] Ryan Prescott Adams, Iain Murray, and David JC MacKay. Tractable nonparametric bayesian inference in poisson processes with gaussian process intensities. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 9–16. ACM, 2009.
- [5] Cedric Archambeau, Dan Cornford, Manfred Opper, and John Shawe-Taylor. Gaussian process approximations of stochastic differential equations. In *Gaussian Processes in Practice*, pages 1–16, 2007.
- [6] Andreas Damianou, Michalis K Titsias, and Neil D Lawrence. Variational gaussian process dynamical systems. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2011.
- [7] Iain Murray, David MacKay, and Ryan P Adams. The gaussian process density sampler. In *Advances in Neural Information Processing Systems*, pages 9–16, 2009.
- [8] Jaakko Riihimäki, Aki Vehtari, et al. Laplace approximation for logistic gaussian process density estimation and regression. *Bayesian analysis*, 9(2):425–448, 2014.
- [9] Nicholas G Polson, James G Scott, and Jesse Windle. Bayesian inference for logistic models using pólya-gamma latent variables. *Journal of the American statistical Association*, 108(504):1339–1349, 2013.
- [10] Stephen G Walker. Posterior sampling when the normalizing constant is unknown. *Communications in Statistics—Simulation and Computation*, 40(5):784–792, 2011.
- [11] John Frank Charles Kingman. *Poisson processes*. Wiley Online Library, 1993.
- [12] Takis Konstantopoulos, Zurab Zerakidze, and Grigol Sokhadze. *Radon–Nikodým Theorem*, pages 1161–1164. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [13] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. In *Readings in Computer Vision*, pages 564–584. Elsevier, 1987.
- [14] Scott Linderman. `pypolyagamma`. <https://github.com/slinderman/pypolyagamma>, 2017.

<sup>6</sup>Potentially in such cases other sparsity methods [34] might be more favourable.

- [15] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [16] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [17] Lehel Csató. Gaussian Processes -Iterative Sparse Approximations. *PhD Thesis*, 2002.
- [18] Lehel Csató, Manfred Opper, and Ole Winther. Tap gibbs free energy, belief propagation and sparsity. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 657–663. MIT Press, 2002.
- [19] Michalis K Titsias. Variational learning of inducing variables in sparse gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 567–574, 2009.
- [20] Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, pages 1257–1264, 2006.
- [21] Alexander G de G Matthews, James Hensman, Richard Turner, and Zoubin Ghahramani. On sparse variational methods and the kullback-leibler divergence between stochastic processes. In *Artificial Intelligence and Statistics*, pages 231–239, 2016.
- [22] Philipp Batz, Andreas Ruttner, and Manfred Opper. Approximate bayes learning of stochastic differential equations. *arXiv preprint arXiv:1702.05390*, 2017.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] Christian Donner. Sgpd\_inference. [https://github.com/christiando/SGPD\\_Inference](https://github.com/christiando/SGPD_Inference), 2018.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python . *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [26] David J Hand, Fergus Daly, K McConway, D Lunn, and E Ostrowski. *A handbook of small data sets*, volume 1. cRc Press, 1993.
- [27] Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017.
- [28] Paulo Cortez and Aníbal de Jesus Raimundo Morais. A data mining approach to predict forest fires using meteorological data. 2007.
- [29] Fabian Keller, Emmanuel Muller, and Klemens Bohm. Hics: high contrast subspaces for density-based outlier ranking. In *Data Engineering (ICDE), 2012 IEEE 28th International Conference on*, pages 1037–1048. IEEE, 2012.
- [30] Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 29–37, 2011.
- [31] Benigno Uribe, Iain Murray, and Hugo Larochelle. A deep and tractable density estimator. In *International Conference on Machine Learning*, pages 467–475, 2014.
- [32] Yves-Laurent Kom Samo and Stephen Roberts. Scalable nonparametric bayesian inference on point processes with gaussian processes. In *International Conference on Machine Learning*, pages 2227–2236, 2015.
- [33] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [34] Yarin Gal and Richard Turner. Improving the gaussian process sparse spectrum approximation by representing uncertainty in frequency inputs. In *International Conference on Machine Learning*, pages 655–664, 2015.