



Universidad Autónoma  
de Madrid

**Biblos-e Archivo**  
Repositorio Institucional UAM

**Repositorio Institucional de la Universidad Autónoma de Madrid**  
<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:  
This is an **author produced version** of a paper published in:

Neurocomputing 378 (2020): 210-227

**DOI:** <https://doi.org/10.1016/j.neucom.2019.09.090>

**Copyright:** © 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 licence <http://creativecommons.org/licenses/by-nc-nd/4.0/>

El acceso a la versión del editor puede requerir la suscripción del recurso  
Access to the published version may require subscription

1 Alpha Divergence Minimization in Multi-Class Gaussian  
2 Process Classification

3 Carlos Villacampa-Calvo<sup>a,\*</sup>, Daniel Hernández-Lobato<sup>a</sup>

4 <sup>a</sup>*Computer Science Department, Escuela Politécnica Superior, Universidad Autónoma de*  
5 *Madrid, C/ Francisco Tomás y Valiente, 11, Madrid 28049, Spain*

---

6 **Abstract**

This paper analyzes the minimization of  $\alpha$ -divergences in the context of multi-class Gaussian process classification. For this task, several methods are explored, including memory and computationally efficient variants of the Power Expectation Propagation algorithm, which allow for efficient training using stochastic gradients and mini-batches. When these methods are used for training, very large datasets (several millions of instances) can be considered. The proposed methods are also very general as they can interpolate between other popular approaches for approximate inference based on Expectation Propagation (EP) ( $\alpha \rightarrow 1$ ) and Variational Bayes (VB) ( $\alpha \rightarrow 0$ ) simply by varying the  $\alpha$  parameter. An exhaustive empirical evaluation analyzes the generalization properties of each of the proposed methods for different values of the  $\alpha$  parameter. The results obtained show that one can do better than EP and VB by considering intermediate values of  $\alpha$ .

7 *Keywords:* Gaussian Processes, Expectation Propagation,  $\alpha$ -divergences,  
8 Approximate Inference, Variational Inference

---

\*Corresponding author

*Email address:* `carlos.villacampa@uam.es` (Carlos Villacampa-Calvo)

## 9 1. Introduction

10 Gaussian Processes (GPs) are non-parametric models that can be used  
11 to address machine learning problems, including multi-class classification  
12 [1]. In these models, the expressiveness grows with the training set size  
13  $N$ . Furthermore, they are probabilistic models in which prior knowledge  
14 can be easily specified, and they readily provide a predictive distribution  
15 which accounts for prediction uncertainty. In spite of these advantages, using  
16 Gaussian process in practice is difficult because often the likelihood is not  
17 Gaussian. Therefore, exact inference in these models is usually intractable  
18 and approximate methods need to be employed. A challenging example is  
19 multi-class classification because in this case there is one latent function  
20 (GP) per class, and the likelihood factors are more complicated than, for  
21 example, in binary classification models. An extra difficulty is that standard  
22 approaches for multi-class GP classification require, at least, the inversion  
23 of one  $N \times N$  matrix per class. This is an expensive operation that limits  
24 the applicability of these models to large problems. Notwithstanding, there  
25 are several methods that have been proposed for multi-class GP classification  
26 [2, 3, 4, 5, 6]. Most of them, however, do not scale well with the size of the  
27 training set.

28 The use of sparse approximations allows to scale-up GPs. These techniques  
29 introduce  $M \ll N$  inducing points, whose location is learnt during the training  
30 process. These points lead to an approximate prior with a low-rank covariance  
31 structure [7], reducing the training cost to  $\mathcal{O}(NM^2)$ . This improved cost has  
32 been pushed forward by Hensman et al. [8, 9], which employs a variational  
33 Bayes (VB) approximation combined with stochastic optimization techniques

34 that allows to address datasets with millions of instances. Recent work in the  
35 literature also combines stochastic optimization techniques with alternative  
36 methods for approximate inference based on expectation propagation (EP)  
37 [10, 11]. The results obtained indicate that EP and VB have similar training  
38 costs, but EP may provide better predictive distributions in terms of the test  
39 log-likelihood.

40 While VB minimizes the Kullback-Leibler (KL) divergence between the  
41 approximate and the target distribution, EP minimizes (approximately) the  
42 KL-divergence in the reversed way. Recently, Bui et al. [12] suggested a  
43 framework for binary GP classification that, by means of the minimization  
44 of  $\alpha$ -divergences with Power Expectation Propagation (PEP) [13], unifies  
45 previous approaches based on VB and EP. The  $\alpha$ -divergence generalizes the  
46 KL-divergence and different values of the  $\alpha$  parameter interpolate ( $\alpha \rightarrow 0$   
47 and  $\alpha = 1$ ) between the two versions of the KL-divergence described above  
48 [14]. Importantly, Bui et al. [12] show that, in the case of binary classification,  
49 one can do better than VB and EP by considering an intermediate version of  
50 the two KL-divergences.

51 Here, we extend the minimization of  $\alpha$ -divergences for approximate in-  
52 ference of Bui et al. [12] to address multi-class GP classification problems.  
53 For this, we describe a multi-class extension of the PEP algorithm for binary  
54 GP classification. This extension is not trivial due to the more complicated  
55 likelihood factors that appear in the multi-class setting. Furthermore, instead  
56 of considering a single latent function, in the multi-class case we have one  
57 latent function per class. Besides this, we address here some of the drawbacks  
58 of standard PEP, which include the difficulty of using standard optimization

59 techniques and a high memory consumption. More precisely, standard PEP  
60 combines gradient-based updates of the model hyper-parameters with closed-  
61 form updates to refine the approximate likelihood factors. These approximate  
62 factors have to be stored in memory, which is memory expensive. The variants  
63 of PEP considered are based on using ideas from approximate EP [15, 16]  
64 and from the approximate minimization of  $\alpha$ -divergences in the context of  
65 Bayesian neural networks [17]. The results obtained in our experiments show  
66 that the (approximate) minimization of an intermediate divergence between  
67 the ones considered by VB and EP, *i.e.*, setting  $\alpha = 0.5$ , may work better in  
68 terms of the prediction error and the test log-likelihood.

## 69 2. Multi-Class Gaussian Processes

70 Consider a dataset of  $N$  labelled examples in the form of a matrix  $\mathbf{X} =$   
71  $(\mathbf{x}_1, \dots, \mathbf{x}_N)^T$  and a vector of labels  $\mathbf{y} = (y_1, \dots, y_N)^T$ , where  $y_i \in \{1 \dots C\}$   
72 with  $C > 2$  the total number of classes. The goal is to predict the label of  
73 an unseen instance  $\mathbf{x}_*$ . In multi-class Gaussian process classification it is  
74 usual to use the softmax function. However, it is not the only alternative.  
75 Here, we will follow [3] assume that the label  $y_i$  of  $\mathbf{x}_i$  is generated by the rule  
76  $y_i = \arg \max_k f^k(\mathbf{x}_i)$ , where each  $f^k(\cdot)$  is a latent function associated to a  
77 class  $k \in \{1 \dots C\}$ . Based on this, the likelihood is a product of  $N$  factors  
78 such as:

$$p(\mathbf{y}|\mathbf{f}) = \prod_{i=1}^N p(y_i|\mathbf{f}_i) = \prod_{i=1}^N \prod_{k \neq y_i} \Theta(f^{y_i}(\mathbf{x}_i) - f^k(\mathbf{x}_i)), \quad (1)$$

79 where  $\Theta(\cdot)$  is the Heaviside function and we have defined

80  $\mathbf{f}^k = (f^k(\mathbf{x}_1), \dots, f^k(\mathbf{x}_N))^T \in \mathbb{R}^N$ ,  $\mathbf{f}_i = (f^1(\mathbf{x}_i), \dots, f^C(\mathbf{x}_i))^T \in \mathbb{R}^C$  and

81  $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^C) \in \mathbb{R}^{N \times C}$ . The likelihood in (1) can be made more robust  
 82 to possible labelling errors by adding a parameter  $\epsilon$  which represents the  
 83 probability of choosing at random  $y_i$  from the set of labels [18]. Then, each  
 84 factors is:

$$p(y_i|\mathbf{f}_i) = (1 - \epsilon) \prod_{k \neq y_i} \Theta(f^{y_i}(\mathbf{x}_i) - f^k(\mathbf{x}_i)) + \frac{\epsilon}{C}. \quad (2)$$

85 We assume a GP prior for each  $f^k(\cdot)$  [1]. Particularly,  $p(f^k) \sim \mathcal{GP}(0, c(\cdot, \cdot; \xi^k))$   
 86 where  $c(\cdot, \cdot; \xi^k)$  is a covariance function with hyper-parameters  $\xi^k$ . Moreover,  
 87 we assume these priors to be independent. That is,  $p(\mathbf{f}) = \prod_{k=1}^C p(\mathbf{f}^k)$ , where  
 88 every  $p(\mathbf{f}^k)$  is a multivariate Gaussian distribution. In this model, one can  
 89 easily include Gaussian additive noise around each latent function. In that  
 90 case, the labeling rule described is equivalent to the Gumbel-max trick (which  
 91 leads to a soft-max function), but adding independent Gaussian noise instead  
 92 of Gumbel noise [19]. The task of interest is to compute a posterior distribution  
 93 for  $\mathbf{f}$  using Bayes rule:  $p(\mathbf{f}|\mathbf{y}) = p(\mathbf{y}|\mathbf{f})p(\mathbf{f})/p(\mathbf{y})$ . We can then maximize  
 94 the marginal likelihood  $p(\mathbf{y})$  to find good values for the hyper-parameters  
 95  $\xi^k$ . Nevertheless, as the likelihood factors in (1) and (2) are not Gaussian,  
 96 we will be unable to compute analytically  $p(\mathbf{f}|\mathbf{y})$  and approximate inference  
 97 will be needed: the Laplace approximation [2], EP [3] or VB [20]. These  
 98 methods result in a cost of  $\mathcal{O}(CN^3)$ , where  $N$  is the number of instances and  
 99  $C$  the number of classes, assuming independent GPs per each class (this is  
 100 the hypothesis made in the rest of paper).

### 101 2.1. Sparse Gaussian Processes

102 To speed up calculations, a typical approach is to use sparse approxima-  
 103 tions. These approximations rely on introducing a different set of points of

104 size  $M \ll N$  called inducing points  $\bar{\mathbf{X}}^k = (\bar{\mathbf{x}}_1^k, \dots, \bar{\mathbf{x}}_M^k)^\top$  for each class  $k$ , with  
 105 associated latent values  $\bar{\mathbf{f}}^k = (f^k(\bar{\mathbf{x}}_1^k), \dots, f^k(\bar{\mathbf{x}}_M^k))^\top$  [21]. Now, by setting  
 106 a GP prior on the latent functions associated with the inducing points we  
 107 can obtain an approximate prior for  $\mathbf{f}^k$  as  $p(\mathbf{f}^k) = \int p(\mathbf{f}^k | \bar{\mathbf{f}}^k) p(\bar{\mathbf{f}}^k | \bar{\mathbf{X}}^k) d\bar{\mathbf{f}}^k \approx$   
 108  $\int [\prod_{i=1}^N p(f^k(\mathbf{x}_i) | \bar{\mathbf{f}}^k)] p(\bar{\mathbf{f}}^k | \bar{\mathbf{X}}^k) d\bar{\mathbf{f}}^k = p_{\text{FITC}}(\mathbf{f}^k | \bar{\mathbf{X}}^k)$ , where we have assumed that  
 109  $p(\bar{\mathbf{f}}) = \prod_{k=1}^C p(\bar{\mathbf{f}}^k | \bar{\mathbf{X}}^k)$  and that the conditional distribution  $p(\mathbf{f}^k | \bar{\mathbf{f}}^k)$  factorizes  
 110 like  $\prod_{i=1}^N p(f^k(\mathbf{x}_i) | \bar{\mathbf{f}}^k)$ . In other words, marginalizing over the latent values  
 111 associated with the inducing points  $\bar{\mathbf{f}}^k$  will effectively result in an approximate  
 112 covariance function for the prior on the latent values  $\mathbf{f}^k$  [22]. This approxima-  
 113 tion is known as the Fully Independent Training Conditional (FITC) [7] and  
 114 gives an approximate inference cost of  $\mathcal{O}(NM^2)$ .

## 115 2.2. Scalable Gaussian Processes: EP

116 A method for approximate inference in multi-class GP classification is  
 117 Expectation Propagation (EP) [23]. In EP the latent variables,  $\mathbf{f}$ , of the  
 118 process at the training points  $\mathbf{X}$  are marginalized out. The task of inter-  
 119 est is to approximate the posterior of the process values at the inducing  
 120 points  $\bar{\mathbf{f}} = (\bar{\mathbf{f}}^1, \dots, \bar{\mathbf{f}}^K)^\top$ :  $p(\bar{\mathbf{f}} | \mathbf{y}) \propto \prod_{i=1}^N \phi_i(\bar{\mathbf{f}}) p(\bar{\mathbf{f}})$ , where  $\phi_i(\bar{\mathbf{f}})$  is a likelihood  
 121 factor defined as  $\phi_i(\bar{\mathbf{f}}) = \int p(y_i | \mathbf{f}_i) p(\mathbf{f}_i | \bar{\mathbf{f}}) d\mathbf{f}_i$  and  $p(\bar{\mathbf{f}}) = \prod_{k=1}^C p(\bar{\mathbf{f}}^k)$  is the  
 122 prior distribution over the inducing values. In this last expression  $p(\mathbf{f}_i | \bar{\mathbf{f}})$   
 123 is a conditional Gaussian distribution that factorizes across classes, *i.e.*,  
 124  $p(\mathbf{f}_i | \bar{\mathbf{f}}) = \prod_{k=1}^C p(f^k(\mathbf{x}_i) | \bar{\mathbf{f}}^k)$ . EP approximates each non-Gaussian factor of  
 125 the likelihood  $\phi_i$  with a Gaussian factor  $\tilde{\phi}_i$  [10]. More precisely, it refines  
 126 at each iteration a factor  $\tilde{\phi}_i$  of the approximate posterior  $q(\bar{\mathbf{f}}) \propto \prod_{i=1}^N \tilde{\phi}_i p(\bar{\mathbf{f}})$   
 127 by computing the cavity distribution  $q^{\setminus i} \propto q / \tilde{\phi}_i$  and then minimizing locally  
 128 the Kullback-Leibler divergence between  $Z_i^{-1} \phi_i q^{\setminus i}$  and  $q$  with respect to  $q$ ,

129 *i.e.*,  $\text{KL}[Z_i^{-1}\phi_i q^{\setminus i} \parallel q]$  where  $Z_i$  is the normalization constant of  $\phi_i q^{\setminus i}$ . The  
 130 updated factor is simply  $\tilde{\phi}_i^{\text{new}} = Z_i q^{\text{new}} / q^{\setminus i}$ . The KL-divergence minimization  
 131 is done using the derivatives of  $\log Z_i$  w.r.t. the parameters of  $q^{\setminus i}$  [24] and  
 132  $Z_i$  can be computed using a one-dimensional quadrature. Note that  $q$  is  
 133 Gaussian because the prior and each  $\tilde{\phi}_i$  are Gaussian. The approximation to  
 134 the marginal likelihood  $p(\mathbf{y})$ , denoted  $Z_q$ , is simply the normalization constant  
 135 of  $\prod_{i=1}^N \tilde{\phi}_i(\bar{\mathbf{f}}) p(\bar{\mathbf{f}})$ . The gradient of  $\log Z_q$  w.r.t. a hyper-parameter  $\xi_j^k$  of the  
 136  $k$ -th covariance function can be easily obtained since the parameters of each  
 137  $\tilde{\phi}_i$  can be considered fixed after running EP [24].

138 Recent work in the literature shows that it is possible to scale to large  
 139 datasets the previous EP approach [11, 10]. One only has to jointly update  
 140 the approximate factors  $\tilde{\phi}_i$  and the model hyper-parameters  $\xi^k$ . Furthermore,  
 141 because  $\log Z_q$  contains a sum across the data points, stochastic optimization  
 142 techniques can be used to update the model hyper-parameters. This allows  
 143 to scale to very large datasets with millions of instances.

144 A limitation of EP is that the parameters of each  $\phi_i$  have to be stored  
 145 in memory. A further approximation to EP called Stochastic Expectation  
 146 Propagation (SEP) [15] assumes that all the approximate factors are tied and  
 147 only keeps in memory the product of all of them instead of their individual  
 148 parameters. This reduces the memory cost to  $\mathcal{O}(CM^2)$ .

149 Interestingly, the previous derivation of the EP algorithm for approx-  
 150 imate inference in multi-class GPC is equivalent to the one that is ob-  
 151 tained when one approximates the posterior distribution of  $\mathbf{f}$  and  $\bar{\mathbf{f}}$ , *i.e.*,  
 152  $p(\mathbf{f}, \bar{\mathbf{f}} | \mathbf{y}) \propto \prod_{i=1}^N p(\mathbf{y}_i | \mathbf{f}_i) p(\mathbf{f}_i | \bar{\mathbf{f}}) p(\bar{\mathbf{f}})$ , under the constraint that the approxi-  
 153 mate distribution is  $q(\mathbf{f}, \bar{\mathbf{f}}) \propto p(\mathbf{f} | \bar{\mathbf{f}}) \prod_{i=1} \tilde{\phi}_i p(\bar{\mathbf{f}})$ , where  $p(\mathbf{f} | \bar{\mathbf{f}}) = \prod_{k=1}^C p(\mathbf{f}^k | \bar{\mathbf{f}}^k)$ .

154 That is, each likelihood factor  $p(y_i|\mathbf{f}_i)$  has been approximated by the corre-  
 155 sponding factor  $\tilde{\phi}_i$  which depends on  $\bar{\mathbf{f}}$ . Specifically, the conditional distribu-  
 156 tion  $p(\mathbf{f}|\bar{\mathbf{f}})$  in  $q$  is fixed and we can only update the part of  $q$  that depends  
 157 on the inducing values  $\bar{\mathbf{f}}$ . In this we case, there is no need to use the FITC  
 158 approximation. See [12] for further details and the specific equivalence in the  
 159 regression case.

### 160 2.3. Scalable Gaussian Processes: VB

Another approach for approximate inference is Variational Bayes (VB) [25, 26, 27, 9]. In this section we will follow the derivation of the lower bound in [9]. VB uses the same likelihood function as EP. The approximate distribution  $q$  is the same as the one considered at the end of the previous section. Namely,  $q(\mathbf{f}, \bar{\mathbf{f}}) = p(\mathbf{f}|\bar{\mathbf{f}})q(\bar{\mathbf{f}})$ , where  $q(\bar{\mathbf{f}})$  is Gaussian and  $p(\mathbf{f}|\bar{\mathbf{f}})$  is fixed. The distribution  $q$  is found by minimizing the KL-divergence between  $q$  and the exact posterior  $p(\mathbf{f}, \bar{\mathbf{f}}|\mathbf{y})$ . It is possible to show that this minimization is equivalent to the maximization of a lower bound on the log-marginal likelihood  $\log p(\mathbf{y})$ . This lower bound is obtained by first applying Jensen’s inequality to obtain a lower bound to the log conditional  $\log p(\mathbf{y}|\bar{\mathbf{f}})$ :

$$\log p(\mathbf{y}|\bar{\mathbf{f}}) = \log \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\bar{\mathbf{f}})d\mathbf{f} \geq \mathbb{E}_{p(\mathbf{f}|\bar{\mathbf{f}})}[\log p(\mathbf{y}|\mathbf{f})]. \quad (3)$$

Then, a lower bound to the log-marginal likelihood is derived in the same way:

$$\begin{aligned} \log p(\mathbf{y}) &= \log \int q(\bar{\mathbf{f}})p(\mathbf{y}|\bar{\mathbf{f}})p(\bar{\mathbf{f}})/q(\bar{\mathbf{f}})d\bar{\mathbf{f}} \\ &\geq \mathbb{E}_{q(\bar{\mathbf{f}})}[\log p(\mathbf{y}|\bar{\mathbf{f}})] - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})], \end{aligned} \quad (4)$$

where  $q(\bar{\mathbf{f}})$  is approximate Gaussian distribution. By substituting (3) in (4) we obtain the final lower bound:

$$\begin{aligned}
\log p(\mathbf{y}) &\geq \mathbb{E}_{q(\bar{\mathbf{f}})}[\log p(\mathbf{y}|\bar{\mathbf{f}})] - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})] \\
&\geq \mathbb{E}_{q(\bar{\mathbf{f}})}[\mathbb{E}_{p(\mathbf{f}|\bar{\mathbf{f}})}[\log p(\mathbf{y}|\mathbf{f})]] - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})] \\
&\geq \mathbb{E}_{q(\mathbf{f})}[\log p(\mathbf{y}|\mathbf{f})] - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})] \\
&\geq \sum_{i=1}^N \mathbb{E}_{q(\mathbf{f}_i)}[\log p(y_i|\mathbf{f}_i)] - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})], \tag{5}
\end{aligned}$$

161 where  $q(\mathbf{f}) = \int p(\mathbf{f}|\bar{\mathbf{f}})q(\bar{\mathbf{f}})d\bar{\mathbf{f}}$  and each marginal over  $\mathbf{f}_i = (f^1(\mathbf{x}_i), \dots, f^C(\mathbf{x}_i))^T$   
162 is a product of  $C$  Gaussian conditional distributions with mean  $\hat{m}_i^k$  and  
163 variance  $\hat{s}_i^k$ , for  $k = 1, \dots, C$ . Namely,  $q(\mathbf{f}_i) = \prod_{k=1}^C \mathcal{N}(f^k(\mathbf{x}_i)|\hat{m}_i^k, \hat{s}_i^k)$ .

164 The lower bound contains a sum over the training examples, so stochastic  
165 optimization techniques can be used for its optimization. As in EP,  
166 one-dimensional quadratures must be used to approximate the required ex-  
167 pectations in (5). Last, this formulation minimizes the global KL-divergence  
168 between the approximate distribution  $q$  and the posterior, and can be shown  
169 to be equivalent to minimizing  $\text{KL}[q \parallel Z_i \tilde{\phi}_i q^{\lambda_i}]$  (the reversed divergence) in  
170 EP [14, 28].

### 171 3. Alpha-Divergence Minimization

We introduce the  $\alpha$ -divergence [29, 30], a divergence measure that general-  
izes the KL divergence [14], as well as the different approaches proposed for its  
minimization in the context of Gaussian processes for multi-class classification.  
The  $\alpha$ -divergence between two probability distributions  $p$  and  $q$  of a random

variable  $\boldsymbol{\theta}$  is:

$$D_\alpha[p||q] = \frac{1 - \int p(\boldsymbol{\theta})^\alpha q(\boldsymbol{\theta})^{1-\alpha} d\boldsymbol{\theta}}{\alpha(1 - \alpha)}, \quad (6)$$

172 where  $\alpha \in \mathbb{R} \setminus \{0, 1\}$ .

173 The case  $\alpha = 0.5$  is called the Hellinger distance the only member of the  
 174 family of  $\alpha$ -divergences that is symmetric in  $p$  and  $q$  [31]. More precisely,  
 175  $D_{\frac{1}{2}}[p || q] = 2 \int_{\boldsymbol{\theta}} (\sqrt{p(\boldsymbol{\theta})} - \sqrt{q(\boldsymbol{\theta})})^2 d\boldsymbol{\theta}$ . Furthermore,  $D_0[p || q] = \lim_{\alpha \rightarrow 0} D_\alpha[p ||$   
 176  $q] = \text{KL}[q || p]$  is used in VB and  $D_1[p || q] = \lim_{\alpha \rightarrow 1} D_\alpha[p || q] = \text{KL}[p || q]$  is  
 177 used in EP, so the  $\alpha$ -divergence minimization can easily interpolate between  
 178 these two methods by changing the value of  $\alpha$ .

### 179 3.1. Power Expectation Propagation (PEP)

Power Expectation Propagation (PEP) is an extension of EP that instead  
 of minimizing the KL-divergence at each step, minimizes an  $\alpha$ -divergence [13].  
 Importantly, this  $\alpha$ -divergence minimization is done by simply minimizing the  
 KL-divergence between some modified distribution and  $q$ . Specifically, when  
 computing the cavity distribution  $q^{\setminus \alpha i}$ , the approximate factor  $\tilde{\phi}_i$  to the power  
 of  $\alpha$  is removed. That is,  $q^{\setminus \alpha i} \propto q / \tilde{\phi}_i^\alpha$ . Next, the KL divergence between  
 $Z_i^{-1} \phi_i^\alpha q^{\setminus \alpha i}$  and  $q$ ,  $\text{KL}[Z_i^{-1} \phi_i^\alpha q^{\setminus \alpha i} || q]$ , is minimized with respect to  $q$ , where  
 $Z_i$  is the normalization constant of  $\phi_i^\alpha q^{\setminus \alpha i}$ . Note that the factor  $\phi_i$  is raised to  
 the power of  $\alpha$ . The updated factor is simply  $\tilde{\phi}_i^{\text{new}} = (Z_i q^{\text{new}} / q^{\setminus \alpha i})^{\frac{1}{\alpha}}$ , since the  
 exact factor  $\phi_i$  had been raised to the power of  $\alpha$ . Importantly, it is possible  
 to show that this minimization is equivalent to minimizing  $D_\alpha[Z_i \phi_i q^{\setminus i} || q]$  [14].  
 More precisely, let  $\lambda_q$  be the parameters of  $q$ . For a distribution  $p$  and  $q$  in

the exponential family:

$$\begin{aligned} \nabla_{\lambda_q} D_\alpha[p||q] &= \frac{Z_{\tilde{p}}}{\alpha} (\mathbb{E}_q[s(\boldsymbol{\theta})] - \mathbb{E}_{\tilde{p}}[s(\boldsymbol{\theta})]) \\ &\propto \nabla_{\lambda_q} \text{KL}[\tilde{p}||q], \end{aligned} \tag{7}$$

180 where  $\tilde{p} \propto p^\alpha q^{1-\alpha}$  and  $s(\boldsymbol{\theta})$  is the vector of sufficient statistics of  $q$ . At the  
 181 minimum both gradients must be equal to zero and the moments of  $q$  and  
 182  $\tilde{p}$  must match. If we let  $p \propto \phi_i q^{\setminus i}$ , as in EP, the corresponding distribution  
 183  $\tilde{p} \propto \phi_i^\alpha q^{\setminus \alpha i}$ , as in PEP. Therefore, at convergence, when the approximate  
 184 factors do not change any more and (7) is equal to zero for each approximate  
 185 factor, PEP minimizes the  $\alpha$ -divergences between the tilted distributions  
 186 defined as  $\phi_i q^{\setminus i}$ ,  $\forall i$ , and  $q$ . Importantly, this local divergence minimization  
 187 becomes a global divergence minimization (between the target posterior and  
 188  $q$ ) only when  $\alpha \rightarrow 0$  [14]. In all the other cases the global  $\alpha$ -divergence  
 189 minimization is approximate, but accurate as shown in [14].

190 The PEP algorithm consists in applying the following steps to every  
 191 approximate factor  $\tilde{\phi}_i$  and repeat them until it has converged:

**Remove** an approximate factor to the power of  $\alpha$  from the posterior  $q$  to  
 compute the cavity distribution  $q^{\setminus \alpha i}$ .

$$q^{\setminus \alpha i} = \frac{q}{(\tilde{\phi}_i)^\alpha}$$

**Include** the true factor  $\phi_i$  to the power of  $\alpha$  to compute the tilted distri-  
 bution  $\hat{p}$ .

$$\hat{p} = (\phi_i)^\alpha q^{\setminus \alpha i}$$

**Project** onto the approximating family by matching moments.

$$q^{\text{new}} = \arg \min_{q^*} \text{KL}[\hat{p} || q^*]$$

**Update** the approximate factor.

$$(\tilde{\phi}_i)^\alpha = \frac{q^{\text{new}}}{q^{\setminus \alpha i}}$$

To apply PEP to the model described in this manuscript, we consider the approximation described at the end of Section 2.2. Namely,  $q(\mathbf{f}, \bar{\mathbf{f}}) \propto p(\mathbf{f}|\bar{\mathbf{f}}) \prod_{i=1}^N \tilde{\phi}_i p(\bar{\mathbf{f}})$ , where  $\tilde{\phi}_i$  are Gaussian factors depending only on  $\bar{\mathbf{f}}$ . The marginal likelihood approximation of PEP,  $Z_q$ , is the normalization constant of the previous expression:

$$\begin{aligned} \log Z_q &= g(\boldsymbol{\theta}_{\text{post}}) - g(\boldsymbol{\theta}_{\text{prior}}) + \frac{1}{\alpha} \sum_{i=1}^N \log \tilde{Z}_i, \\ \log \tilde{Z}_i &= \log \mathbb{E}_{q(\mathbf{f}_i)} [(\phi_i / \tilde{\phi}_i)^\alpha], \end{aligned} \tag{8}$$

192 where each  $\phi_i = p(y_i|\mathbf{f}_i)$ ;  $g(\cdot)$  is the log-normalizer of a distribution in the  
 193 exponential family of  $q$ ; and  $\boldsymbol{\theta}_{\text{post}}$  and  $\boldsymbol{\theta}_{\text{prior}}$  are the natural parameters of  
 194  $q(\bar{\mathbf{f}})$  and the prior, respectively. When  $\alpha = 1$ , (8) is equivalent to the EP  
 195 approximation of the log-marginal likelihood. In the limit when  $\alpha \rightarrow 0$ , one  
 196 can show that (8) tends to the lower bound optimized in VB. See [14] for  
 197 further details. The expectation in (8) can be computed using one-dimensional  
 198 quadrature methods. In particular, it is simply related to the probability  
 199 that one Gaussian random variable is larger than several others (one per each  
 200 other class label) [18].

201 As in EP, the gradient of  $\log Z_q$  w.r.t. the model hyper-parameters  
 202 (inducing points locations and parameters of the covariance functions) involves  
 203 a sum across the data instances. Therefore, mini-batches and stochastic  
 204 optimization methods can also be used here to maximize  $\log Z_q$ . This allows  
 205 to scale-up to very large datasets. Usually, one has to wait until PEP has

206 converged to compute the gradient and update the model hyper-parameters.  
 207 Nevertheless, it is possible to follow the same approach as in [10] and jointly  
 208 optimize the approximate factors and the model hyper-parameters.

### 209 3.2. Approximate Power EP (APEP)

210 A limitation of the PEP algorithm described in Section 3.1 is that it needs  
 211 to keep in memory the parameters of all the approximate factors, which are  
 212 optimized through PEP updates by moment matching. To overcome this,  
 213 a first approximation to the PEP method considers stochastic expectation  
 214 propagation [15], which ties all the approximate factors and only keeps in  
 215 memory their product, *i.e.*,  $\tilde{\phi} = \prod_{i=1}^N \tilde{\phi}_i$ . Note that this only affects the  
 216 way of computing the cavity distribution  $q^{\setminus \alpha i}$ . Under this approximation  
 217  $q^{\setminus \alpha i}$  is computed in an approximate way. Namely,  $q^{\setminus \alpha i} \propto q / \tilde{\phi}^{\frac{\alpha}{N}}$ , where  $N$  is  
 218 the number of factors (and also data points). Besides this, we optimize the  
 219 global factor  $\tilde{\phi}$  by maximizing  $\log Z_q$ , the approximation to the log-marginal  
 220 likelihood, w.r.t. the parameters of  $\tilde{\phi}$ , instead of using the PEP updates. This  
 221 is supported by the fact that these updates also find a stationary point of this  
 222 energy function [14]. This allows the use of standard optimization techniques  
 223 to find the posterior approximation  $q$ , which is defined as  $q \propto \tilde{\phi} p(\bar{\mathbf{f}})$ .

### 224 3.3. Approximate Reparameterized PEP (ARPEP)

As another way to approximately optimize the PEP evidence or energy  
 function, we consider the approach described by Li and Gal [17] for Bayesian  
 neural networks. In that work it is described a reparameterization of the  
 PEP energy function that is compatible with an approximate distribution  $q$   
 that need not belong to the exponential family, although we will also assume

a Gaussian form here. In the large data limit, *i.e.*, when  $\alpha/N \rightarrow 0$ , the reparameterized objective is simply approximated by:

$$\begin{aligned} \log Z_q &\approx \frac{1}{\alpha} \sum_{i=1}^N \log \mathbb{E}_{q(\mathbf{f}_i)} [p(y_i | \mathbf{f}_i)^\alpha] \\ &\quad - \text{KL}[q(\bar{\mathbf{f}}) \parallel p(\bar{\mathbf{f}})]. \end{aligned} \tag{9}$$

225 This objective is a combination of the terms appearing in the PEP estimate  
 226 of the log-marginal likelihood (8) and the lower bound of VB (5). The KL-  
 227 divergence term in (9) can be understood as a regularization term enforcing  $q$   
 228 to look similar to the prior. Because this objective also involves a sum across  
 229 the data points, it can be efficiently optimized both w.r.t. the parameters of  
 230  $q$  and the model hyper-parameters using stochastic optimization techniques.

### 231 3.4. Refined Prior Approximate PEP (RPAPPEP)

232 Some of the solutions obtained by VB can not be retrieved by the methods  
 233 from Sections 3.2 and 3.1 due to the parameterization resulting in a non  
 234 positive definite covariance matrix. A last method accounts for this. It is  
 235 the same method as the one described in Section 3.2, but where we let  $q$   
 236 be an arbitrary Gaussian distribution and eliminate the assumption that  $q$   
 237 is proportional to a Gaussian times the prior. Namely,  $q \propto \tilde{\phi}p(\bar{\mathbf{f}})$ . This is  
 238 precisely the same hypothesis made by VB or the method described in Section  
 239 3.3. For this, we simply let the prior be another extra factor to be refined  
 240 by PEP. Thus, instead of considering  $N$  factors, one per each point, we will  
 241 have  $N + 1$  factors, the extra factor corresponding to the prior. Under this  
 242 setting, the PEP approximate log-marginal likelihood is:

$$\log Z_q = g(\boldsymbol{\theta}_{\text{post}}) + \frac{1}{\alpha} \sum_{i=1}^{N+1} \log \tilde{Z}_i. \tag{10}$$

243 In this method we also retrieve the approximate EP energy objective when  
244  $\alpha = 1$  and VB’s lower bound as  $\alpha \rightarrow 0$ . Stochastic optimization is also  
245 possible.

### 246 *3.5. Summary of Approximate Inference Methods*

247 In the previous sections we have described four methods that deal with  
248 the minimization of  $\alpha$ -divergences in Gaussian process models for multi-class  
249 classification. In this section we will summarize the characteristics of each of  
250 the methods with the aim of giving a better understanding of them. Table 3.5  
251 shows, for each method, a summary of what we think are the most relevant  
252 features: the ability to use standard optimization techniques, the use of  
253 stochastic EP to make the method memory efficient and whether  $q$  has a free  
254 Gaussian form.

255 The first method, PEP, was first introduced in [13] as a generalization of  
256 the EP algorithm. In [12] they successfully apply the algorithm to minimize  
257  $\alpha$ -divergences with sparse GPs and perform extensive experiments in the  
258 regression and binary classification cases. This method is precisely the one in  
259 [12], but applied to multi-class classification problems. It follows the general  
260 PEP scheme where one has to alternate between updating the approximate  
261 factors by moment matching and gradient based optimization of the model  
262 hyper-parameters. However, in the original PEP formulation [13], one has  
263 to wait until convergence before updating the hyper-parameters and here,  
264 we follow [10] and jointly optimize the approximate factors and the hyper-  
265 parameters. This method is not memory efficient due to the need of keeping  
266 in memory all the approximate factors.

267 The second method, APEP, was described in [31] as a black-box method

268 that can be applied to general probabilistic models. They propose a simplified  
269 objective by tying the approximate factors following [15], and directly optimize  
270 the posterior approximation  $q$  using the gradients of the simplified objective.  
271 This makes the method memory efficient and allows for standard optimization  
272 techniques instead of having to rely on the PEP update step to optimize the  
273 approximate factors. In this work, we use the APEP method in the specific  
274 case of GP models applied to multi-class classification problems.

275 The third method, ARPEP, was proposed in [17] for the specific case of  
276 Bayesian neural networks as an approximate way to minimize  $\alpha$ -divergences.  
277 We have apply the same idea to the case of multi-class GP classification. This  
278 method is also memory efficient, can be used with standard optimization  
279 techniques. Also, in this case the posterior approximation  $q$  takes the form  
280 of a free Gaussian, meaning that it is no longer proportional to a Gaussian  
281 times the prior.

282 The last method, RPAPPEP, has been considered because the parameter-  
283 ization used in both PEP and APEP prevents them to reach some of the  
284 solutions that can be obtained by VB. To account for this, this method is  
285 based on APEP but where we let  $q$  be a free Gaussian, instead of a Gaussian  
286 times the prior. This method is memory efficient as well and can be used  
287 with standard optimization techniques.

#### 288 4. Related work

Other works in the literature have addressed the approximate minimization  
of  $\alpha$ -divergences. In particular, [31] also approximate the Power EP objective  
with a simplified energy function by tying the approximate factors. More

	PEP	APEP	ARPEP	RPAPEP
Standard optimization		✓	✓	✓
Memory efficient		✓	✓	✓
Free Gaussian			✓	✓

Table 1: Summary of the proposed methods

precisely, the objective considered by these authors is:

$$E[\boldsymbol{\theta}_{\text{prior}}, \boldsymbol{\theta}] = g(\boldsymbol{\theta}_{\text{prior}}) - g(\boldsymbol{\theta}_{\text{post}}) - \frac{1}{\alpha} \sum_{i=1}^N \log \mathbb{E}_q \left[ \left( \frac{p(\mathbf{y}_i | \mathbf{f}_i)}{\tilde{\phi}_i} \right)^\alpha \right], \quad (11)$$

289 where  $\boldsymbol{\theta}_{\text{prior}}$  and  $\boldsymbol{\theta}_{\text{post}}$  are the natural parameters of the prior and the approx-  
290 imate posterior  $q$  respectively;  $g(\boldsymbol{\theta}_{\text{prior}})$  and  $g(\boldsymbol{\theta}_{\text{post}})$  are their log-normalizers;  
291  $\boldsymbol{\theta} = (\boldsymbol{\theta}_{\text{post}} - \boldsymbol{\theta}_{\text{prior}})/N$  are the parameters of the global approximate factor  
292  $\tilde{\phi}$ ; and  $p(\mathbf{y}_i | \mathbf{f}_i)$  is the true likelihood factor. This method’s objective will be  
293 equivalent to the one obtained by the approximation in Section 3.2, but criti-  
294 cally, the expectations in (11), which may involve multiple random variables  
295 and may lack an analytic expression, are approximated via Monte Carlo. The  
296 result is black-box algorithm that can be used for approximate inference in  
297 arbitrary complicated models. In principle this method could also be used for  
298 approximate inference in the context of multi-class Gaussian process classifica-  
299 tion. Notwithstanding, in this particular case, the required expectations can  
300 be evaluated using one-dimensional quadrature methods, which is believed to  
301 be significantly more efficient than using a Monte Carlo estimate of the same  
302 quantity. Therefore, the approaches considered in the the present work are  
303 expected to be more efficient for approximately optimizing the PEP objective.  
304 Moreover, a Monte Carlo estimate of (11) will lead to a biased objective due

305 to the non-linearity of the logarithm function.

306 The minimization of  $\alpha$ -divergences for binary Gaussian process classifica-  
307 tion has also been explored by Bui et al. [12]. These authors also use power EP  
308 as a unifying framework to work with GPs and  $\alpha$ -divergences. Moreover, Bui  
309 [32] also compares in the binary classification case PEP with APEP, finding  
310 that they give similar results. However, despite the extensive experimental  
311 results in [12, 32], the multi-class classification case was not specifically con-  
312 sidered nor analyzed. Importantly, the extension from binary to multi-class  
313 problems is more challenging. Instead of having one single latent function, in  
314 the multi-class case there is one latent function per each class in the problem.  
315 Furthermore, the likelihood factors are also more complicated, and even lack  
316 an analytical expression. Our work complements that of Bui et al. [12] by  
317 providing a careful and exhaustive analysis of the multi-class case, which has  
318 been systematically overlooked by the literature on Gaussian processes. Be-  
319 sides this, we also consider alternative methods for approximately optimizing  
320 the PEP objective. These methods are memory efficient (the memory cost  
321 is independent of  $N$ ) and can make use of standard optimization techniques  
322 (*i.e.*, they do not use PEP updates of the approximate factors).

323 Other methods for approximate inference in GPs with arbitrary likelihoods  
324 can also target the VB objective in Section 2.3 [33, 34, 35]. Instead of  
325 quadratures, they rely on a Monte Carlo approximation, which is expected to  
326 lead to higher variance in the gradients estimation and to affect negatively  
327 the optimization process.

## 328 5. Experiments

329 In this section, we intend to compare the performance of the different  
330 values of  $\alpha$  when using  $\alpha$ -divergences in the multi-class setting. As we can  
331 retrieve VI solution by making  $\alpha \rightarrow 0$  and the EP solution with  $\alpha = 1$  we  
332 are comparing the proposed methods to the state-of-the-art methods for  
333 scalable approximate inference with Gaussian processes. We compare the  
334 methods described in Section 3.1 to Section 3.4 (PEP, APEP, ARPEP and  
335 RPAPPEP) in several experiments. The R code of each method is found  
336 in the supplementary material. All methods start with the same hyper-  
337 parameters (including the location of the inducing points), which are opti-  
338 mized by maximizing the estimate of the log-marginal likelihood. We use  
339 an ARD Gaussian kernel for each latent function [1], and optimize the am-  
340 plitude and additive noise parameter. An implementation in R of all the  
341 compared methods is available at [http://arantxa.ii.uam.es/%7edhernan/](http://arantxa.ii.uam.es/%7edhernan/alpha-mgpc/R-code_alpha_MGPC.zip)  
342 [alpha-mgpc/R-code\\_alpha\\_MGPC.zip](http://arantxa.ii.uam.es/%7edhernan/alpha-mgpc/R-code_alpha_MGPC.zip). In the Appendix C there is a com-  
343 parison between the PEP method and some baseline methods, including label  
344 regression, Laplace approximation and MCMC based method, showing that  
345 the predictive distribution of PEP is good.

### 346 5.1. Performance on UCI Datasets

347 We compare the four methods, for different values of  $\alpha$ , over 8 UCI-  
348 repository [36] multi-class problems. These problems are fairly small (see  
349 Appendix B for the datasets' details), but will show how each method  
350 performs on standard problems. Later on, we will consider larger datasets.  
351 Because the datasets are small we use here batch optimization. We use 90% of

352 the data for training and 10% for testing, except for *Satellite* which is bigger  
353 (we choose 20% for training and 80% for test). In *Vowel* we consider only  
354 the points belonging to the 6 first classes. Finally, in *Waveform* (synthetic)  
355 we generate 1000 instances and split them in 30% for training and 70% for  
356 testing. All methods are trained for 500 iterations using l-BFGS, except for  
357 PEP, which uses gradient ascent with an adaptive learning rate (described in  
358 Appendix A.7). We consider three values for the number of inducing points  
359  $M$ . Namely, 5%, 10% and 20% of the number of training data  $N$ . The values  
360 of  $\alpha$  considered range from  $\alpha \rightarrow 0$  to  $\alpha = 1$  with steps of size 0.1. We report  
361 averages over 20 repetitions.

362 Figure 1 shows, for each method, the average rank for each value of  $\alpha$ ,  
363 in terms of the test log-likelihood. Average ranks are computed, for a fixed  
364 method, for each value of  $\alpha$ , across dataset and splits and values of the number  
365 of inducing points  $M$ . Besides this, we analyze which method is better, given  
366 a particular value of  $\alpha$ . For that, we compute the average rank of each method  
367 across datasets, splits and values of  $M$ , this time fixing  $\alpha$  and varying the  
368 method instead of the other way around. This rank is shown using a color  
369 pattern with *red* meaning a higher average rank and *blue* a lower average  
370 rank. We observe that for PEP, the Hellinger value  $\alpha = 0.5$ , seems to give  
371 better performance in terms of the negative test log-likelihood. For APEP  
372 and ARPEP, a value of  $\alpha$  between 0.6 and 0.8 gives better results. Finally,  
373 RPAPEP gives in general worst results than the other methods for almost all  
374 values of  $\alpha$ . Results for ranks computed in terms of the test error are shown  
375 in Appendix D. They do not differ significantly from ones shown here.

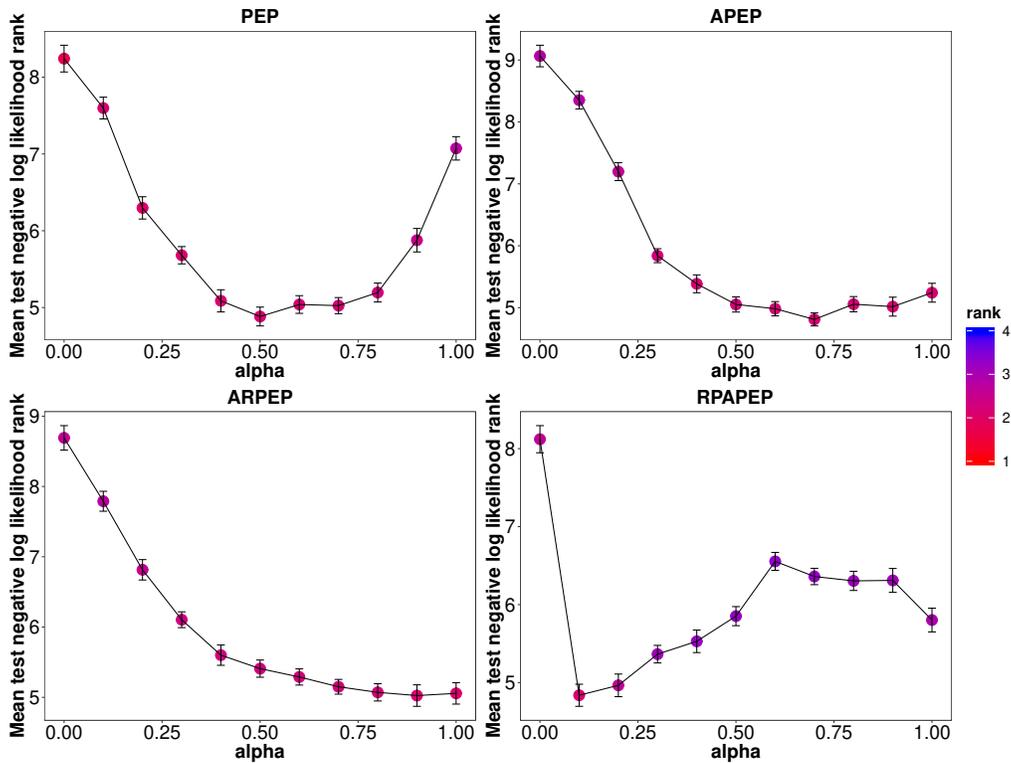


Figure 1: Avg. test neg. log-likelihood rank for each method and each value of  $\alpha$ . The color of the points indicates the average rank a method compared with the others. Best seen in color.

376 *5.2. Analysis of Inducing Points Location*

377 We also analyze, for each method, the location of the inducing points for  
 378 different values of  $\alpha$ . We use a synthetic two-dimensional problem with three  
 379 classes reproduced from [10]. We consider 1,000 training points and a fixed  
 380 number of inducing points  $M = 128$ . The initial location of the inducing  
 381 points is chosen at random and it is the same for all the methods. In these  
 382 experiments we we keep fixed the other hyper-parameters to their true value

383 [10]. PEP and APEP are trained using batch methods and ARPEP and  
 384 RPAPEP are trained using stochastic methods to avoid sub-optimal solutions.  
 385 In batch training we consider 2,000 iterations and when using mini-batches  
 386 we consider 2,000 epochs. Additionally, we use ADAM for training (default  
 387 settings) and 100 as the mini-batch size [37].

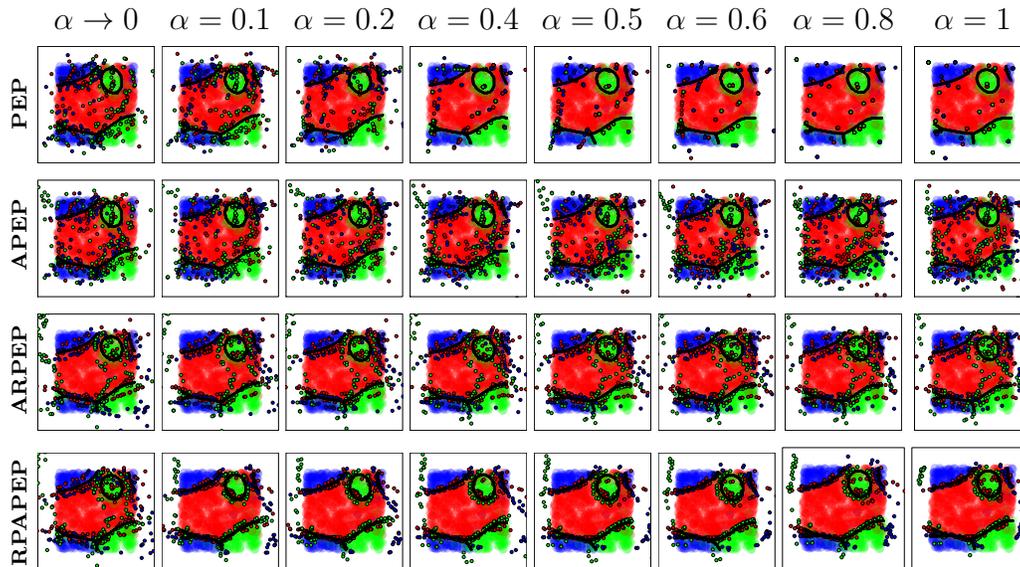


Figure 2: Decision boundaries and inducing points location for different values of  $\alpha$  ( $M = 128$ ).

388 Figure 2 shows the final location of the inducing points. Blue, red and  
 389 green points are training data, black lines are the decision boundaries and  
 390 black border points are the inducing points. We expect that for values of  $\alpha$   
 391 near zero, the inducing points would tend to place near the decision boundaries,  
 392 since this is the behavior observed in [8, 10, 12]. Indeed, this is the case  
 393 of ARPEP and RPAPEP, probably because they are the two formulations  
 394 in which  $q$  has a free Gaussian form. By contrast, in PEP and APEP this

395 behavior is not observed. As we increase  $\alpha$ , in PEP the inducing points  
 396 overlap, which can be seen as an inducing point pruning technique, previously  
 397 reported in [38, 10]. This behavior is not observed for the other methods,  
 398 probably as a consequence of using either a different parameterization for  $q$ ,  
 399 or due to the approximation employed in APEP. Interestingly, for RPAPEP  
 400 the inducing points tend to be even closer to the decision boundary as we  
 401 approach  $\alpha = 1$ . Finally, in APEP and ARPEP  $\alpha$  does not have a strong  
 402 influence in the location of the inducing points.

### 403 *5.3. Performance in Terms of Training Time*

404 We compare the performance of each method as a function of the training  
 405 time on the Satellite dataset. Training is done as in Section 5.1. We consider  
 406  $M = 4, 20$  and  $100$ . We also set  $\alpha \rightarrow 0$  and  $\alpha = 0.3, 0.5, 0.8, 1$ . We report  
 407 averages over 100 repetitions of the experiments. Figure 3 shows the average  
 408 test negative log-likelihood for each method and each value of  $\alpha$ . Similar  
 409 plots for the test error are included in Appendix D. In general, when  $\alpha \rightarrow 0$ ,  
 410 and in the case of the method VB, we obtain the worst results. For PEP the  
 411 best performance is obtained for  $\alpha = 0.3$  and  $\alpha = 0.5$ . For the rest of the  
 412 methods it seems that values between  $\alpha = 0.8$  and  $\alpha = 1$  tend to give good  
 413 over-all results. RPAPEP seems to slightly over-fit the training data, which  
 414 may explain the worse results of this method in the UCI datasets. ARPEP,  
 415 RPAPEP give almost the same results as VB for  $\alpha \rightarrow 0$ , which is the expected  
 416 behavior. In PEP and APEP this is not the case, probably because of the  
 417 different parameterization of  $q$ , which is proportional to a Gaussian times the  
 418 prior. Finally, PEP gives better results earlier, probably as a consequence of  
 419 using PEP-updates to refine  $q$ , instead of gradient-based updates.

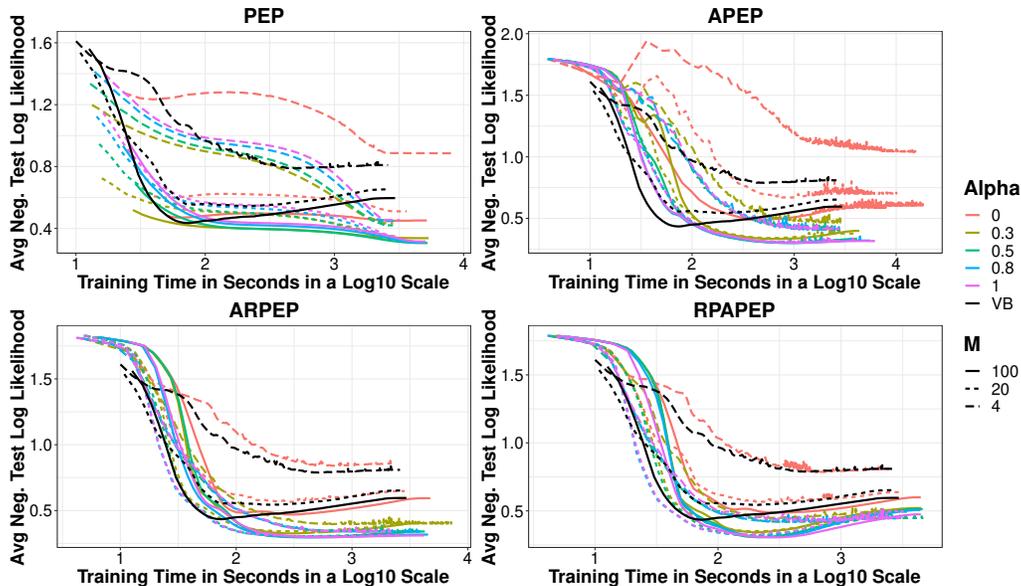


Figure 3: Neg. test log-likelihood on the Satellite dataset for different values of  $\alpha$  and  $M$ . Best seen in color.

#### 420 5.4. Performance on MNIST

421 When addressing very large datasets one can no longer rely on batch  
 422 training, and mini-batches and stochastic gradients are required. A large  
 423 problem is MNIST [39], with 60,000 instances for training and 10,000 for  
 424 testing. We train the proposed methods on this dataset setting  $M = 200$   
 425 and using a mini-batch size of 200. We consider several values for  $\alpha$  from  
 426  $\alpha \rightarrow 0$  to  $\alpha = 1$  with a step-size of 0.1. Each method is trained using ADAM  
 427 (except PEP which uses EP-updates to refine  $q$ ) with the default parameters  
 428 [37]. We include the results for Variational Bayes (VB) for reference. Figure  
 429 4 shows the test negative log-likelihood for each method and each value of  $\alpha$ .  
 430 The same plots but for the test error are included in Appendix D. All the  
 431 methods seem to give similar results, but APEP reaches the optimal solution

432 first. Importantly, in each method the lower the value of  $\alpha$  the faster the  
 433 convergence.

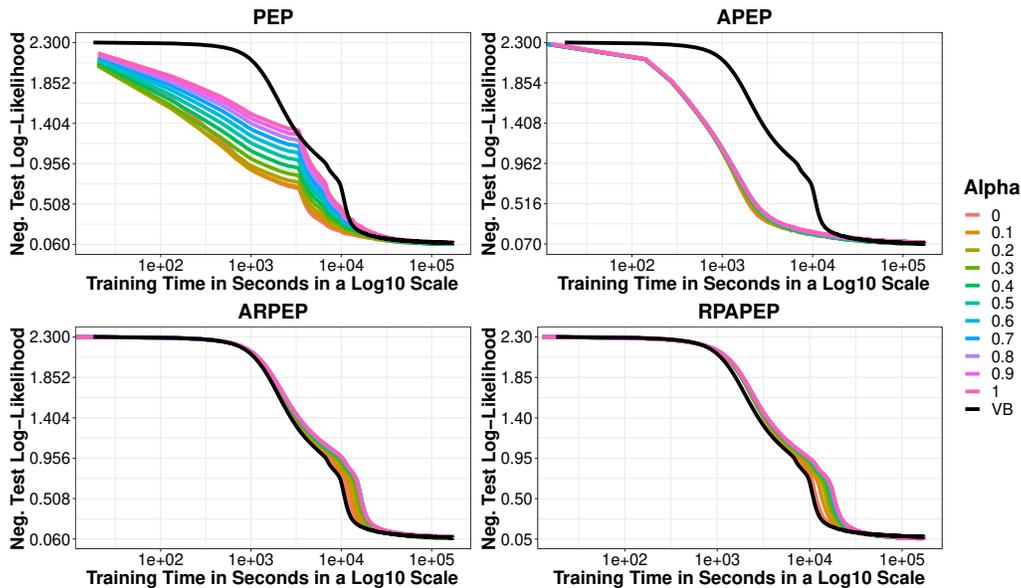


Figure 4: Negative test log-likelihood for each method on the MNIST dataset for each  $\alpha$ . Best seen in color.

#### 434 5.5. Performance on Airline Delays

435 We consider a dataset with information about the flights within the USA  
 436 between 01/2008 and 04/2008 <sup>1</sup>. It has three classes: Flight on time, with  
 437 more than 5 minutes of delay, or arrived 5 minutes before scheduled time. We  
 438 consider 8 attributes: age of the aircraft, distance covered, airtime, departure  
 439 time, arrival time, day of the week, day of the month and month. After

<sup>1</sup><http://stat-computing.org/dataexpo/2009>

440 removing the data with missing values, 2,127,068 instances remain, from  
441 which 10,000 are used for testing and the rest for training. We evaluate each  
442 method using the same setting as on the MNIST dataset. We also include  
443 the results obtained by VB for reference.

444 Figure 5 shows the negative test log-likelihood of each method as a function  
445 of time. The results are similar in terms of the test error (see Appendix D).  
446 Regarding the negative test log-likelihood, as  $\alpha$  approaches 0, worse results  
447 are obtained. This had previously been observed in [10], and is believed  
448 to be a consequence of the particular objective that is optimized by VB.  
449 As  $\alpha$  increases, the approximation to the log-marginal likelihood of PEP  
450 resembles more the EP objective, which is closer to the test-log likelihood  
451  $\log \mathbb{E}_{q(\mathbf{f}_i)}[p(y_i|\mathbf{f}_i)]$ . This explains the better results of  $\alpha = 1$ . Here,  $\alpha = 0.5$   
452 also provides good results.

### 453 *5.6. Active Learning: Waveform*

454 As a way of measuring the quality of the predictive distribution we have  
455 conducted a last experiment on the waveform dataset. We consider an active  
456 learning approach where we will iteratively add a new data point to the  
457 training set. For that, we will need an initial training set, a test set to  
458 evaluate the performance and a validation set from which to select the new  
459 data points. To choose which point to select next from the validation set,  
460 we will use the predictive distribution of the proposed methods, by selecting  
461 the point in which the entropy is highest and hence adopting an explorative  
462 approach. We compare this selection mechanism versus selecting the next  
463 point at random from the validation set. We start with 100 points for training,  
464 500 for test and 400 for validation, and we will add 100 new points to the

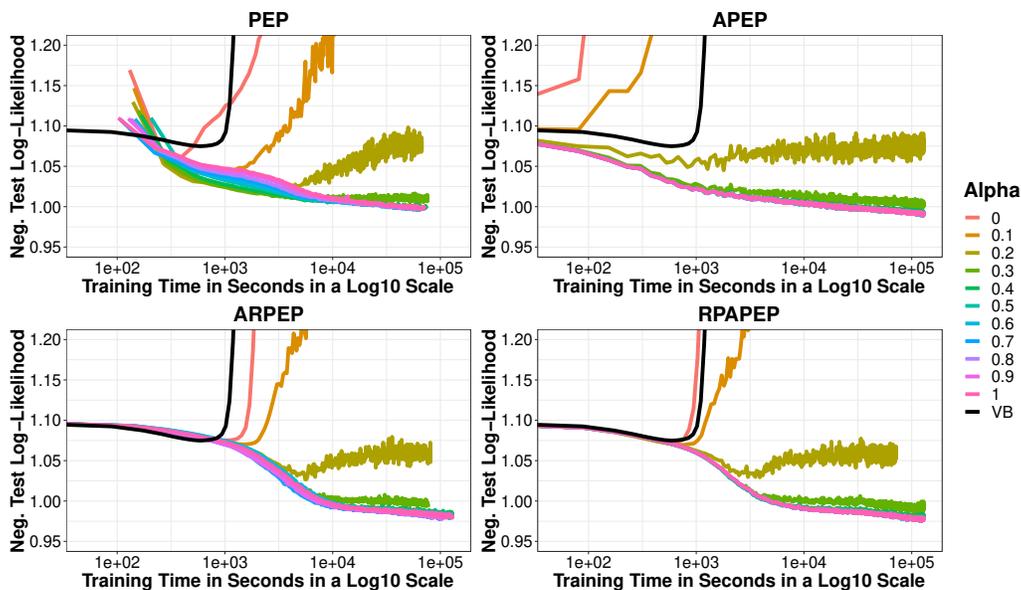


Figure 5: Neg. test log-likelihoods on the Airline Delays dataset for each  $\alpha$ . Best seen in color.

465 training set. All methods are trained using l-BFGS the first time for 250  
 466 iterations and then re-trained each time we add a new point for 25 more  
 467 iterations, reusing the solution obtained so far. For the PEP algorithm, when  
 468 adding a new point to the training set one must also add a new approximate  
 469 factor. This made the retraining process start in a bad solution leading to  
 470 bad results. In order to overcome this problem, we have combined both the  
 471 PEP updates of the approximate factors and l-BFGS, by alternating between  
 472 updating the approximate factors and optimizing the model hyper-parameters  
 473 with l-BFGS. In this case, as we are updating the factors several times at  
 474 each l-BFGS iteration in an internal loop the training process is more costly,  
 475 therefore we have reduced the initial training from 250 to 50 iterations and the

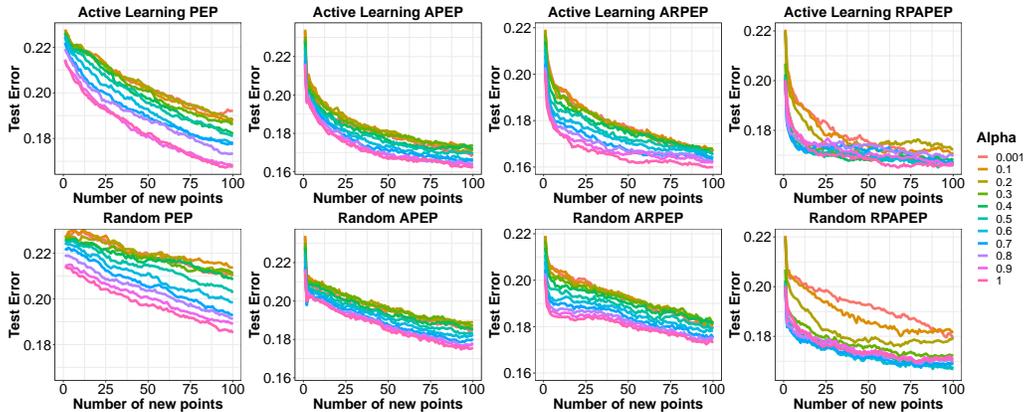


Figure 6: Test error on the Waveform dataset as a function of the number of added points to the training set, selected using an active learning approach (top) and selected at random (bottom). Best seen in color.

476 retraining from 25 to 5 iterations. We report averages over 100 repetitions.

477 Figure 6 shows, for each value of  $\alpha$ , the classification error in the test set  
 478 as a function of the number of new added points. In the top row, each new  
 479 point has been added by means of the active learning approach, selecting the  
 480 point in which the entropy is highest. In the bottom row, each new point  
 481 has been selected at random. We observe that the error is lower for values  
 482 close to  $\alpha = 1$ , both for the active learning approach and random selection.  
 483 Also, the test error is always lower for the active learning approach than for  
 484 random selection for all the methods, showing the utility of the predictive  
 485 distribution for this type of problems.

486 In Figure 7 it is shown, for each value of  $\alpha$ , the reduction in the test error  
 487 w.r.t. the initial error. For PEP, the reduction in the error is higher when we  
 488 choose higher values of  $\alpha$ , but for the other methods values of  $\alpha \rightarrow 0$  give

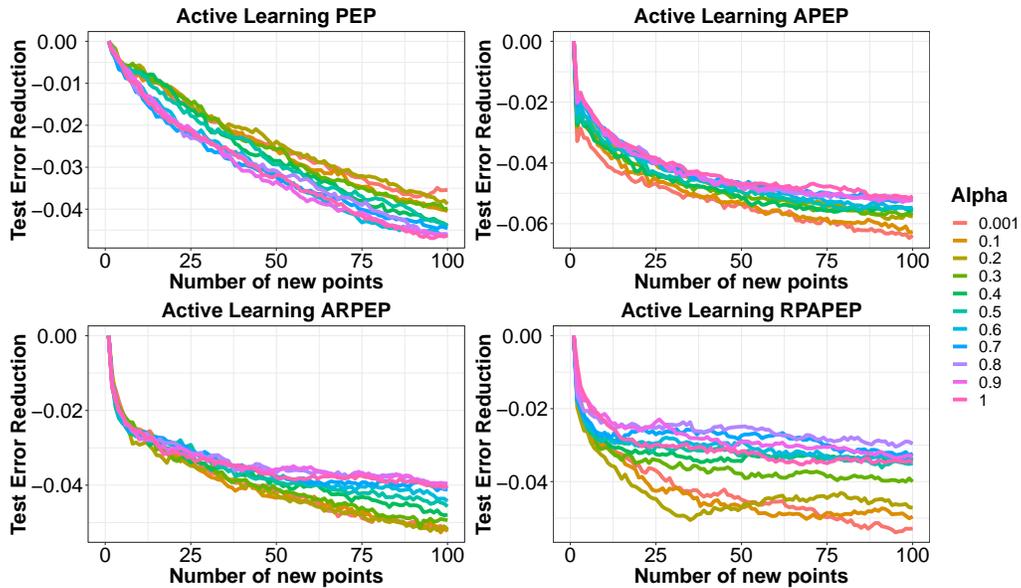


Figure 7: Test error reduction on the Waveform dataset as a function of the number of added points to the training set, selected using an active learning approach. Best seen in color.

489 better test error reduction. This is because the initial test error as  $\alpha \rightarrow 0$  is  
 490 worse for all the methods but, at the end of the training process, all the values  
 491 of  $\alpha$  give similar values for the test error, although slightly better for values  
 492 near  $\alpha = 1$ . In the case of PEP, the difference in the test error between lower  
 493 and higher values of  $\alpha$  is bigger at the end, resulting in a better reduction for  
 494 higher values of  $\alpha$ .

## 495 6. Conclusions

496 The optimization of  $\alpha$ -divergences allows to interpolate between approxi-  
 497 mate inference methods that are closer to VB when  $\alpha \rightarrow 0$  or EP as  $\alpha \rightarrow 1$ .

498 Previous work in the literature had already considered the optimization of  
499 these divergences for approximate inference [12, 31, 17]. In this work, we  
500 have analyzed its specific minimization in the case of multi-class classification  
501 using GPs. We have compared four approximate methods for this: PEP,  
502 APEP, ARPEP and RPAPEP. These approximations are memory efficient  
503 (except PEP) and can be combined with batch training methods, as well as  
504 with stochastic training methods. When using mini-batches and stochastic  
505 techniques, the training cost is  $\mathcal{O}(CM^3)$ . We have done several experiments  
506 comparing the proposed approximations for different values of  $\alpha$ .

507 While none of the proposed methods seems to be superior to the others  
508 (except RPAPEP, which performs slightly worse), there are some points that  
509 one should keep in mind when using them in practice. First of all, PEP is not  
510 memory efficient as it needs to keep in memory all the approximate factors.  
511 This clearly a drawback with respect to the other methods, especially when  
512 working with big datasets. Also, these factors are optimized through PEP  
513 updates, which makes the implementation slightly more complicated, as we  
514 cannot rely on standard optimization techniques like in the other methods.  
515 APEP and ARPEP give very similar results and can be used indistinctly. A  
516 difference between those two methods is that the ARPEP objective is more  
517 similar to the one optimized in VB, including the fact that the posterior  
518  $q$  takes the form of a free Gaussian, which is why it exhibits some of the  
519 properties previously reported for VB (*e.g.*, the inducing points tend to place  
520 near the decision boundaries).

521 The results obtained show that intermediate values of  $\alpha$ , *e.g.*,  $\alpha = 0.5$ ,  
522 can provide in general better results than standard approximate inference

523 methods based on VB or EP in some of the problems investigated. This  
524 agrees with previous results for the case of regression or binary classification  
525 problems, as indicated in [12].

## 526 **Acknowledgements**

527 We acknowledge the use of the facilities of Centro de Computación  
528 Científica (CCC) at Universidad Autónoma de Madrid and support from  
529 the Spanish *Ministerio de Economía, Industria y Competitividad*, grants  
530 TIN2016-76406-P and TEC2016-81900-REDT.

## 531 **Appendix A. Details for implementing PEP**

### 532 *Appendix A.1. Introduction*

533 In this document we detail all the steps needed to implement the PEP  
534 algorithm described in the main manuscript. In particular, we describe how  
535 to reconstruct the posterior approximation from the approximate factors  
536 and how to refine these factors. We also detail the computation of the PEP  
537 approximation to the marginal likelihood and its gradients, as well as those  
538 of the proposed approximations in the main manuscript. Finally, we include  
539 some additional experimental results.

### 540 *Appendix A.2. Reconstruction of the posterior approximation*

541 In this section we show how to obtain the posterior distribution by multi-  
542 plying the approximate factors  $\tilde{\phi}_i(\bar{\mathbf{f}})$  and the prior  $p(\bar{\mathbf{f}})$ . Each factor  $\phi_i$  will  
543 be replaced by PEP for an approximate Gaussian factor  $\tilde{\phi}_i$  of the form:

$$\tilde{\phi}_i(\bar{\mathbf{f}}) = \tilde{Z}_i \prod_{k=1}^C \exp \left\{ -\frac{1}{2} (\bar{\mathbf{f}}^k)^T \tilde{\mathbf{V}}_{i,k} \bar{\mathbf{f}}^k + (\bar{\mathbf{f}}^k)^T \tilde{\mathbf{m}}_{i,k} \right\}, \quad (\text{A.1})$$

where  $\tilde{\mathbf{V}}_{i,k}$ ,  $k$  and  $\tilde{\mathbf{m}}_{i,k}$  have the following especial form (see Appendix A.4 for the detailed derivation):

$$\tilde{\mathbf{V}}_{i,k} = C_{i,k}^1 \mathbf{v}_i^k (\mathbf{v}_i^k)^T, \quad (\text{A.2})$$

$$\tilde{\mathbf{m}}_{i,k} = C_{i,k}^2 \mathbf{v}_i^k, \quad (\text{A.3})$$

where we have defined  $\mathbf{v}_i^k = (\mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1}$  and  $\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k$  is a  $M \times M$  matrix with the cross covariances between  $\bar{\mathbf{f}}^k$  and  $C_{i,k}^1$ , and  $C_{i,k}^2$  are parameters found by PEP. We also know from the main manuscript that the prior has the following form:

$$p(\bar{\mathbf{f}}) = \prod_{k=1}^C p(\bar{\mathbf{f}}^k | \bar{\mathbf{X}}^k) = \prod_{k=1}^C \mathcal{N}(\bar{\mathbf{f}}^k | \mathbf{0}, \mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k), \quad (\text{A.4})$$

So the posterior approximation will have the following form

$$q(\bar{\mathbf{f}}) = \frac{1}{Z_q} \left[ \prod_{i=1}^N \tilde{\phi}_i \right] \prod_{k=1}^C p(\bar{\mathbf{f}}^k | \bar{\mathbf{X}}^k). \quad (\text{A.5})$$

Given that all the factors are Gaussian, a distribution that is closed under product and division,  $q(\bar{\mathbf{f}})$  is also Gaussian. In particular, the posterior approximation is defined as  $q(\bar{\mathbf{f}}) = \prod_{k=1}^C \mathcal{N}(\bar{\mathbf{f}} | \mathbf{m}_k, \mathbf{V}_k)$ . The parameters of this distribution can be obtained by using the formulas given in the Appendix of [40] for the product of two Gaussians, leading to

$$\begin{aligned} \mathbf{V}_k &= \left[ (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1} + \Upsilon_k \Delta_k \Upsilon_k^T \right]^{-1}, \\ \mathbf{m}_k &= \mathbf{V}_k \Upsilon_k \tilde{\boldsymbol{\mu}}_k, \end{aligned} \quad (\text{A.6})$$

544 where  $\mathbf{\Upsilon}_k = (\mathbf{v}_1^k, \dots, \mathbf{v}_N^k)$  is a  $M \times N$  matrix,  $\mathbf{\Delta}_k$  is a diagonal  $N \times N$  matrix  
545 with diagonal entries equal to  $C_{i,k}^1$  and  $\tilde{\boldsymbol{\mu}}_k$  is a vector where each component  
546 is equal to  $C_{i,k}^2$ .

547 *Appendix A.3. Computation of the cavity distribution*

Here we will obtain the expressions for the parameters of the cavity distribution  $q^{\setminus \alpha i}$ . This distribution is computed by dividing the posterior approximation by the corresponding approximate factor to the power of  $\alpha$ :

$$q(\bar{\mathbf{f}})^{\setminus \alpha i} \propto \frac{q(\bar{\mathbf{f}})}{\tilde{\phi}_i(\bar{\mathbf{f}})^\alpha}. \quad (\text{A.7})$$

Given that all factors are Gaussian, the resulting distribution will also be Gaussian. The parameters can be obtained by using again the formulas in the Appendix of [40]. However, because  $\tilde{\phi}_i$  only depends on  $\bar{\mathbf{f}}^k$ , only these components of  $q(\bar{\mathbf{f}})^{\setminus \alpha i}$  will change. The corresponding parameters of  $q(\bar{\mathbf{f}})^{\setminus \alpha i}$  are:

$$\begin{aligned} \mathbf{V}_k^{\setminus \alpha i} &= (\mathbf{V}_k^{-1} - \alpha \tilde{\mathbf{V}}_{i,k})^{-1} \\ &= (\mathbf{V}_k^{-1} - \alpha C_{i,k}^{1,k} \mathbf{v}_i^k (\mathbf{v}_i^k)^\text{T})^{-1} \\ &= \mathbf{V}_k + \mathbf{V}_k \mathbf{v}_i^k [(\alpha C_{i,k}^{1,k})^{-1} - \mathbf{v}_i^k \mathbf{V}_k (\mathbf{v}_i^k)^\text{T}]^{-1} (\mathbf{v}_i^k)^\text{T} \mathbf{V}_k, \\ \mathbf{m}_k^{\setminus \alpha i} &= \mathbf{V}_k^{\setminus \alpha i} (\mathbf{V}_k^{-1} \mathbf{m}_k - \alpha \tilde{\mathbf{m}}_{i,k}^k) \\ &= \mathbf{V}_k^{\setminus \alpha i} (\mathbf{V}_k^{-1} \mathbf{m}_k - C_{i,k}^{2,k} \mathbf{v}_i^k) \\ &= \mathbf{V}_k^{\setminus \alpha i} \mathbf{V}_k^{-1} \mathbf{m}_k - \alpha C_{i,k}^{2,k} \mathbf{v}_i^k \mathbf{V}_k^{\setminus \alpha i} \\ &= \mathbf{V}_k \mathbf{V}_k^{-1} \mathbf{m}_k + \mathbf{V}_k \mathbf{v}_i^k [(\alpha C_{i,k}^{1,k})^{-1} - \mathbf{v}_i^k \mathbf{V}_k (\mathbf{v}_i^k)^\text{T}]^{-1} (\mathbf{v}_i^k)^\text{T} \mathbf{V}_k \mathbf{V}_k^{-1} \mathbf{m}_k \\ &\quad - \alpha C_{i,k}^{2,k} \mathbf{v}_i^k \mathbf{V}_k^{\setminus \alpha i} \\ &= \mathbf{m}_k + \mathbf{V}_k \mathbf{v}_i^k [(\alpha C_{i,k}^{1,k})^{-1} - \mathbf{v}_i^k \mathbf{V}_k (\mathbf{v}_i^k)^\text{T}]^{-1} (\mathbf{v}_i^k)^\text{T} \mathbf{m}_k - \mathbf{V}_k \mathbf{v}^\text{T} \alpha C_{i,k}^{2,k} \\ &\quad - \mathbf{V}_k \mathbf{v}_i^k [(\alpha C_{i,k}^{1,k})^{-1} - \mathbf{v}_i^k \mathbf{V}_k (\mathbf{v}_i^k)^\text{T}]^{-1} (\mathbf{v}_i^k)^\text{T} \mathbf{V}_k \mathbf{v}_i^k \alpha C_{i,k}^{2,k}, \end{aligned} \quad (\text{A.9})$$

548 where we have used the Woodbury matrix identity and  $\mathbf{v}_i^k$ ,  $C_{i,k}^1$  and  $C_{i,k}^2$  are  
549 the parameters specified in Appendix A.2.

In this section we show how to find the approximate factors  $\tilde{\phi}_i$  once the cavity distribution  $q^{\setminus\alpha i}$  has already been computed. We can compute the moments of  $\phi_i q^{\setminus\alpha i}$  by getting the derivatives of  $\log Z_i$  with respect to the parameters of  $q^{\setminus\alpha i}$ , as indicated in the Appendix of [40]. For that, note that:

$$\mathbf{m}_k = \mathbf{v}_i^k \mathbf{m}^{\setminus\alpha i} \quad (\text{A.10})$$

$$\mathbf{V}_k = \kappa_{\mathbf{x}_i \mathbf{x}_i}^k - (\mathbf{v}_i^k)^\top (\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{v}_i^k + (\mathbf{v}_i^k)^\top \mathbf{V}^{\setminus\alpha i} \mathbf{v}_i^k \quad (\text{A.11})$$

The derivatives are:

$$\frac{\partial \log Z_i}{\partial \mathbf{m}_k^{\setminus\alpha i}} = \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \frac{\partial \mathbf{m}_k}{\partial \mathbf{m}_k^{\setminus\alpha i}} = \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^k, \quad (\text{A.12})$$

$$\frac{\partial \log Z_i}{\partial \mathbf{V}_k^{\setminus\alpha i}} = \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \frac{\partial \mathbf{V}_k}{\partial \mathbf{V}_k^{\setminus\alpha i}} = \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \mathbf{v}_i^k (\mathbf{v}_i^k)^\top, \quad (\text{A.13})$$

where  $\mathbf{v}_i^k$  is the parameter specified in Appendix A.2. By following the Appendix of [40] we can obtain the moments of  $\phi_i q^{\setminus\alpha i}$  (mean  $\hat{\mathbf{m}}_c$  and covariance  $\hat{\mathbf{V}}_c$ ) from the derivatives of  $\log Z_i$  with respect to the parameters of  $q^{\setminus\alpha i}$ . Namely:

$$\hat{\mathbf{m}}_{i,k} = \mathbf{m}_k^{\setminus\alpha i} + \mathbf{V}_k^{\setminus\alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k^{\setminus\alpha i}} = \mathbf{m}_k^{\setminus\alpha i} + \mathbf{V}_k^{\setminus\alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^k \quad (\text{A.14})$$

$$\begin{aligned} \hat{\mathbf{V}}_{i,k} &= \mathbf{V}_k^{\setminus\alpha i} - \mathbf{V}_k^{\setminus\alpha i} \left( \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k^{\setminus\alpha i}} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k^{\setminus\alpha i}} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k^{\setminus\alpha i}} \right) \mathbf{V}_k^{\setminus\alpha i} \\ &= \mathbf{V}_k^{\setminus\alpha i} - \mathbf{V}_k^{\setminus\alpha i} \left( \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top \mathbf{v}_i^k (\mathbf{v}_i^k)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \mathbf{v}_i^k (\mathbf{v}_i^k)^\top \right) \mathbf{V}_k^{\setminus\alpha i} \\ &= \mathbf{V}_k^{\setminus\alpha i} - \mathbf{V}_k^{\setminus\alpha i} \left( \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \right) \mathbf{v}_i^k (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus\alpha i}. \end{aligned} \quad (\text{A.15})$$

Now we can find the parameters of the approximate factor  $\tilde{\phi}_i$ , which is obtained as  $\tilde{\phi}_i = Z_i q^{\text{new}} / q^{\setminus\alpha i}$ , where  $q^{\text{new}}$  is a Gaussian distribution with the

parameters of  $\phi_i q^{\setminus \alpha i}$  just computed. By following the equations given in the Appendix of [40] we obtain the precision matrices of the approximate factor:

$$\begin{aligned}
\tilde{\mathbf{V}}_{i,k} &= (\hat{\mathbf{V}}_{i,k})^{-1} - (\mathbf{V}_k^{\setminus \alpha i})^{-1} \\
&= \left( \mathbf{V}_k^{\setminus \alpha i} - \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \left[ \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \right] (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} \right)^{-1} - (\mathbf{V}_k^{\setminus \alpha i})^{-1} \\
&= (\mathbf{V}_k^{\setminus \alpha i})^{-1} + (\mathbf{V}_k^{\setminus \alpha i})^{-1} \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \left( \left[ \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \right]^{-1} \right. \\
&\quad \left. - (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} (\mathbf{V}_k^{\setminus \alpha i})^{-1} \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \right)^{-1} (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} (\mathbf{V}_k^{\setminus \alpha i})^{-1} - (\mathbf{V}_k^{\setminus \alpha i})^{-1} \\
&= \mathbf{v}_i^k \left( \left[ \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \right]^{-1} - (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \right)^{-1} (\mathbf{v}_i^k)^\top, \tag{A.16}
\end{aligned}$$

where we have used the Woodbury matrix identity to compute  $(\hat{\mathbf{V}}_{i,k})^{-1}$ . Let us define  $C_{i,k}^1$  as:

$$C_{i,k}^1 = \frac{1}{\alpha} \left( \left[ \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right) \left( \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \right)^\top - 2 \frac{\partial \log Z_i}{\partial \mathbf{V}_k} \right]^{-1} - (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \right)^{-1}, \tag{A.17}$$

551 where we divide by  $\alpha$  because we retrieve  $\alpha$  times the approximate factor  
552 from PEP. The precision matrix of the approximate factors will be then:

$$\tilde{\mathbf{V}}_{i,k} = C_{i,k}^1 \mathbf{v}_i^k (\mathbf{v}_i^k)^\top. \tag{A.18}$$

For the first natural parameter we proceed in a similar way

$$\begin{aligned}
\tilde{\mathbf{m}}_{i,k}^{y_i} &= (\hat{\mathbf{V}}_{i,k}^{y_i})^{-1} \hat{\mathbf{m}}_{i,k}^{y_i} - (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\
&= ((\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} + \tilde{\mathbf{V}}_{i,k}^{y_i}) \hat{\mathbf{m}}_{i,k}^{y_i} - (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\
&= (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \hat{\mathbf{m}}_{i,k}^{y_i} + \tilde{\mathbf{V}}_{i,k}^{y_i} \hat{\mathbf{m}}_{i,k}^{y_i} - (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\
&= (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \left[ \mathbf{m}_k^{\setminus \alpha i} + \mathbf{V}_k^{\setminus \alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^k \right] \\
&\quad + \tilde{\mathbf{V}}_{i,k}^{y_i} \left[ \mathbf{m}_k^{\setminus \alpha i} + \mathbf{V}_k^{\setminus \alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^k \right] - (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\
&= (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} + (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{V}_{y_i}^{\setminus \alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^{y_i} \\
&\quad + \tilde{\mathbf{V}}_{i,k}^{y_i} \left[ \mathbf{m}_{y_i}^{\setminus \alpha i} + \mathbf{V}_{y_i}^{\setminus \alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^{y_i} \right] - (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\
&= \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^{y_i} + \tilde{\mathbf{V}}_{i,k}^{y_i} \mathbf{m}_{y_i}^{\setminus \alpha i} + \tilde{\mathbf{V}}_{i,k}^{y_i} \mathbf{V}_{y_i}^{\setminus \alpha i} \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^{y_i} \\
&= \frac{\partial \log Z_i}{\partial \mathbf{m}_k} \mathbf{v}_i^{y_i} + C_{i,k}^{1,y_i} \mathbf{v}_i^{y_i} (\mathbf{v}_i^{y_i})^\top \mathbf{m}_{y_i}^{\setminus \alpha i} + \frac{\partial \log Z_i}{\partial \mathbf{m}_k} C_{i,k}^{1,y_i} \mathbf{v}_i^{y_i} (\mathbf{v}_i^{y_i})^\top \mathbf{V}_{y_i}^{\setminus \alpha i} \mathbf{v}_i^{y_i} \\
&= \left[ \frac{\partial \log Z_i}{\partial \mathbf{m}_k} + C_{i,k}^{1,y_i} (\mathbf{v}_i^{y_i})^\top \mathbf{m}_{y_i}^{\setminus \alpha i} + \frac{\partial \log Z_i}{\partial \mathbf{m}_k} C_{i,k}^{1,y_i} (\mathbf{v}_i^{y_i})^\top \mathbf{V}_{y_i}^{\setminus \alpha i} \mathbf{v}_i^{y_i} \right] \mathbf{v}_i^{y_i},
\end{aligned} \tag{A.19}$$

where we have used that  $(\mathbf{V}_{y_i}^{\text{new}})^{-1} = \mathbf{V}_{y_i}^{-1} + \tilde{\mathbf{V}}_{i,k}^{y_i}$ . If we define  $C_{i,k}^2$  as:

$$C_{i,k}^2 = \frac{1}{\alpha} \left[ \frac{\partial \log Z_i}{\partial \mathbf{m}_k} + C_{i,k}^1 (\mathbf{v}_i^k)^\top \mathbf{m}_k^{\setminus \alpha i} + \frac{\partial \log Z_i}{\partial \mathbf{m}_k} C_{i,k}^1 (\mathbf{v}_i^k)^\top \mathbf{V}_k^{\setminus \alpha i} \mathbf{v}_i^k \right], \tag{A.20}$$

we obtain the following expressions for the first natural parameter:

$$\tilde{\mathbf{m}}_{i,k} = C_{i,k}^2 \mathbf{v}_i^k. \tag{A.21}$$

Once we have these parameters we can compute the value of the normal-

ization constant  $\tilde{Z}_i$ , which guarantees that the approximate factor integrates

the same as the exact factor with respect to  $q^{\setminus \alpha i}$ . The log of this constant is:

$$\log \tilde{Z}_i = \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{\phi_i}{\tilde{\phi}_i} \right)^\alpha \right] = \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(y_i | \mathbf{f}_i)}{\tilde{\phi}_i} \right)^\alpha \right] + g(\boldsymbol{\theta}^{\setminus \alpha i}) - g(\boldsymbol{\theta}). \tag{A.22}$$

As we are using the robust likelihood  $p(y_i | \mathbf{f}_i) = (1-\epsilon) \prod_{k \neq y_i} \Theta(f^{y_i}(\mathbf{x}_i) - f^k(\mathbf{x}_i))$

+  $\frac{\epsilon}{C}$  we will need to use one-dimensional quadrature techniques to compute

559 the expectation  $\mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( p(y_i|\mathbf{f}_i)/\tilde{\phi}_i \right)^\alpha \right]$ . This expectation is simply related to  
 560 the probability that a Gaussian random variable is larger than several others  
 561 (one per each other class label) [18].

562 *Appendix A.5. Parallel EP updates and damping*

563 We update all approximate factors in parallel. This means that we compute  
 564 all the quantities required for updating each of the approximate factors at  
 565 once (in particular the quantities derived for the cavity distribution  $q^{\setminus \alpha i}$ ).  
 566 Parallel updates are faster than sequential EP updates because there is no  
 567 need to introduce a loop over the data. All computations can be carried out  
 568 in terms of matrix vector multiplications that are often more efficient. A  
 569 disadvantage of parallel updates is, however, that they may lead to unstable  
 570 PEP updates. To prevent unstable PEP updates we used damped PEP  
 571 updates. These simply replace the PEP updates of each approximate factor  
 572 with a linear combination of old and new parameters. For example, we set  
 573  $\tilde{C}_{i,k}^1 = (\tilde{C}_{i,k}^1)^{\text{new}} \rho + (\tilde{C}_{i,k}^1)^{\text{old}} (1 - \rho)$  in the case of the  $\tilde{C}_{i,k}$  parameter of the  
 574 approximate factor (we do this with all the parameters). In the previous  
 575 expression  $\rho \in [0, 1]$  a value that specifies the amount of damping. If  $\rho = 0$   
 576 no update happens. If  $\rho = 1$  we obtain the original EP update. Importantly,  
 577 damping does not change the EP convergence points so it does not affect to  
 578 the quality of the solution.

579 *Appendix A.6. Estimate of the marginal likelihood*

As we have seen in the main manuscript, the estimate of the log marginal likelihood is:

$$\log Z_q = g(\boldsymbol{\theta}) - g(\boldsymbol{\theta}_{\text{prior}}) + \frac{1}{\alpha} \sum_{i=1}^N \log \tilde{Z}_i \quad (\text{A.23})$$

$$\log \tilde{Z}_i = \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(y_i | \mathbf{f}_i)}{\tilde{\phi}_i} \right)^\alpha \right] + g(\boldsymbol{\theta}^{\setminus \alpha i}) - g(\boldsymbol{\theta}) \quad (\text{A.24})$$

580 where  $\boldsymbol{\theta}$ ,  $\boldsymbol{\theta}^{\setminus \alpha i}$  and  $\boldsymbol{\theta}_{\text{prior}}$  are the natural parameters of  $q$ ,  $q^{\setminus \alpha i}$  and  $p(\bar{\mathbf{f}})$   
 581 respectively and  $g(\boldsymbol{\theta}')$  is the log-normalizer of a multivariate Gaussian with  
 582 natural parameters  $\boldsymbol{\theta}'$ . If  $\boldsymbol{\mu}$  and  $\boldsymbol{\Sigma}$  are the mean and covariance matrix of  
 583 that Gaussian distribution over  $D$  dimensions, then

$$g(\boldsymbol{\theta}') = \frac{D}{2} \log 2\pi + \frac{1}{2} \log |\boldsymbol{\Sigma}| + \frac{1}{2} \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}, \quad (\text{A.25})$$

584 which leads to

$$\log Z_q = \sum_{k=1}^C \frac{1}{2} \log |\mathbf{V}_k| + \frac{1}{2} \mathbf{m}_k^\top \mathbf{V}_k^{-1} \mathbf{m}_k - \frac{1}{2} |\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k| + \frac{1}{\alpha} \sum_{i=1}^N \log \tilde{Z}_i, \quad (\text{A.26})$$

585 with

$$\begin{aligned} \log \tilde{Z}_i &= \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(y_i | \mathbf{f}_i)}{\tilde{\phi}_i} \right)^\alpha \right] + \frac{1}{2} \log |\mathbf{V}_{y_i}^{\setminus \alpha i}| \\ &\quad + \frac{1}{2} (\mathbf{m}_{y_i}^{\setminus \alpha i})^\top (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i} \\ &\quad - \frac{1}{2} \log |\mathbf{V}_{y_i}^{\setminus \alpha i}| - \frac{1}{2} (\mathbf{m}_{y_i}^{\setminus \alpha i})^\top (\mathbf{V}_{y_i}^{\setminus \alpha i})^{-1} \mathbf{m}_{y_i}^{\setminus \alpha i}. \end{aligned} \quad (\text{A.27})$$

586 This expression can be evaluated very efficiently using the Woodbury  
 587 matrix identity; the matrix determinant lemma; that  $(\mathbf{V}_k^{\setminus \alpha i})^{-1} = \mathbf{V}_k^{-1} - \tilde{\mathbf{V}}_{i,k}$ ;  
 588 that  $\mathbf{m}_k^{\setminus \alpha i} = \mathbf{V}_k^{\setminus \alpha i} (\mathbf{V}_k^{-1} \mathbf{m}_k - \tilde{\mathbf{m}}_{i,k})$ ; and the special form of the parameters of  
 589 the approximate factors  $\tilde{\mathbf{V}}_{i,k}$  and  $\tilde{\mathbf{m}}_{i,k}$ .

590 *Appendix A.7. Gradient of log Z<sub>q</sub> after convergence and learning rate*

591 We derive the expression for the gradient of log Z<sub>q</sub> after PEP has converged.  
 592 Let  $\xi_j^k$  be one hyper-parameter of the model (*i.e.*, a parameter of one of the  
 593 covariance functions or a component of the inducing points) and  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}_{\text{prior}}$  to  
 594 the natural parameters of  $q$  and  $p(\bar{\mathbf{f}})$  respectively. When PEP has converged,  
 595 the approximate factors can be considered to be fixed (it does not change  
 596 with the model hyper-parameters) [24]. In this case, it is only necessary to  
 597 consider the direct dependency of log Z<sub>i</sub> on  $\xi_j^k$  [24]. But in our case, we update  
 598 the hyper-parameters at each PEP iteration, so we will need to consider the  
 599 indirect dependency too. The gradient is given by:

$$\begin{aligned}
 \frac{\partial \log Z_q}{\partial \xi_j^k} &= \left( \frac{\partial g(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^\top \frac{\partial \boldsymbol{\theta}}{\partial \xi_j^k} - \left( \frac{\partial g(\boldsymbol{\theta}_{\text{prior}})}{\partial \boldsymbol{\theta}_{\text{prior}}} \right)^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \frac{1}{\alpha} \sum_{i=1}^N \frac{\partial \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(\mathbf{y}_i | \mathbf{f}_i)}{\phi_i} \right)^\alpha \right]}{\partial \xi_j^k} \\
 &\quad + \left( \frac{\partial g(\boldsymbol{\theta}^{\setminus \alpha i})}{\partial \boldsymbol{\theta}^{\setminus \alpha i}} \right)^\top \frac{\partial \boldsymbol{\theta}^{\setminus \alpha i}}{\partial \xi_j^k} - \left( \frac{\partial g(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^\top \frac{\partial \boldsymbol{\theta}}{\partial \xi_j^k} \\
 &= \boldsymbol{\eta}^\top \frac{\partial \boldsymbol{\theta}}{\partial \xi_j^k} - (\boldsymbol{\eta}_{\text{prior}})^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \frac{1}{\alpha} \sum_{i=1}^N \frac{\partial \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(\mathbf{y}_i | \mathbf{f}_i)}{\phi_i} \right)^\alpha \right]}{\partial \xi_j^k} \\
 &\quad + (\boldsymbol{\eta}^{\setminus \alpha i})^\top \frac{\partial \boldsymbol{\theta}^{\setminus \alpha i}}{\partial \xi_j^k} - \boldsymbol{\eta}^\top \frac{\partial \boldsymbol{\theta}}{\partial \xi_j^k} \\
 &= \boldsymbol{\eta}^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} - (\boldsymbol{\eta}_{\text{prior}})^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \frac{1}{\alpha} \sum_{i=1}^N \frac{\partial \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(\mathbf{y}_i | \mathbf{f}_i)}{\phi_i} \right)^\alpha \right]}{\partial \xi_j^k} \\
 &\quad + (\boldsymbol{\eta}^{\setminus \alpha i})^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} - \boldsymbol{\eta}^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} \\
 &= (\boldsymbol{\eta}^\top - \boldsymbol{\eta}_{\text{prior}}^\top) \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \frac{1}{\alpha} \sum_{i=1}^N \frac{\partial \log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( \frac{p(\mathbf{y}_i | \mathbf{f}_i)}{\phi_i} \right)^\alpha \right]}{\partial \xi_j^k} \\
 &\quad + (\boldsymbol{\eta}^{\setminus \alpha i} - \boldsymbol{\eta})^\top \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k},
 \end{aligned} \tag{A.28}$$

600 where we have used the chain rule of matrix derivatives [41], the especial form  
 601 of the derivatives when using inducing points [42] and that  $\boldsymbol{\theta} = \boldsymbol{\theta}_{\text{prior}} + \sum_{i=1}^N \boldsymbol{\theta}_i$ ,  
 602 with  $\boldsymbol{\theta}_i$  the natural parameters of the approximate factor  $\tilde{\phi}_i$ . Furthermore,  $\boldsymbol{\eta}$

603 and  $\boldsymbol{\eta}_{\text{prior}}$  are expected sufficient statistics under the posterior approximation  
 604  $q$  and the prior, respectively. This gradient coincides with the one in the main  
 605 manuscript.

It is important to note that one has to use the chain rule of matrix deriva-  
 tives when trying to use the previous expression to compute the gradient. In  
 particular, natural parameters and expected sufficient statistics are expressed  
 in the form of matrices. Thus, one has to use in practice the chain rule of  
 matrix derivatives, as indicated in [41]. For example:

$$(\boldsymbol{\eta} - \boldsymbol{\eta}_{\text{prior}})^{\text{T}} \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} = -\frac{1}{2} \text{trace} \left( \mathbf{M}_k^{\text{T}} \frac{\partial \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k}{\partial \xi_j^k} \right), \quad (\text{A.29})$$

where

$$\mathbf{M}_k = \left( \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k \right)^{-1} - \left( \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k \right)^{-1} \mathbf{v}_k \left( \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k \right)^{-1} - \left( \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k \right)^{-1} \mathbf{m}_k \mathbf{m}_k^{\text{T}} \left( \mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k \right)^{-1}, \quad (\text{A.30})$$

606 where  $\mathbf{V}_k$  and  $\mathbf{m}_k$  are the covariance matrix and mean vector of the  $k$ -  
 607 th component of  $q$ . Furthermore, several standard properties of the trace  
 608 can be employed to simplify the computations. In particular, the trace is  
 609 invariant to cyclic rotations. Namely,  $\text{trace}(\mathbf{ABCD}) = \text{trace}(\mathbf{DABC})$ . The  
 610 derivatives with respect to each  $\log \mathbb{E}_{q(\mathbf{f}_i)} \left[ \left( p(y_i | \mathbf{f}_i) / \tilde{\phi}_i \right)^\alpha \right]$  can be computed  
 611 using quadrature techniques.

612 In our experiments we use an adaptive learning rate for the batch PEP  
 613 methods. This learning rate is different for each hyper-parameter. The  
 614 rule that we use is to increase the learning rate by 2% if the sign of the  
 615 estimate of the gradient for that hyper-parameter does not change between  
 616 two consecutive iterations. If a change is observed, we multiply the learning  
 617 rate by 1/2. When applying stochastic optimization methods, we use the  
 618 ADAM method with the default settings to estimate the learning rate [37].

619 *Appendix A.8. Predictive distribution*

Once the training has completed, we can use the posterior approximation to make predictions for new instances. For that, we first compute an approximate posterior evaluated at the location of the new instance  $\mathbf{x}^*$ , denoted by  $\mathbf{f}^* = (f^1(\mathbf{x}^*), \dots, f^C(\mathbf{x}^*))^\top$ :

$$p(\mathbf{f}^*|\mathbf{y}) = \int p(\mathbf{f}^*|\bar{\mathbf{f}})p(\bar{\mathbf{f}}|\mathbf{y})d\bar{\mathbf{f}} = \int p(\mathbf{f}^*|\bar{\mathbf{f}})q(\bar{\mathbf{f}})d\bar{\mathbf{f}} \approx \prod_{k=1}^C \mathcal{N}(f^k(\mathbf{x}^*)|m_k^*, v_k^*), \quad (\text{A.31})$$

where:

$$m_k^* = (\mathbf{k}_{\mathbf{x}^*, \bar{\mathbf{x}}^k}^k)^\top (\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{m}_k \quad (\text{A.32})$$

$$v_k^* = \kappa_{\mathbf{x}^*, \mathbf{x}^*}^k - (\mathbf{k}_{\mathbf{x}^*, \bar{\mathbf{x}}^k}^k)^\top (\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{k}_{\mathbf{x}^*, \bar{\mathbf{x}}^k}^k + (\mathbf{k}_{\mathbf{x}^*, \bar{\mathbf{x}}^k}^k)^\top (\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{V}_k (\mathbf{K}_{\bar{\mathbf{x}}^k \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{k}_{\mathbf{x}^*, \bar{\mathbf{x}}^k}^k. \quad (\text{A.33})$$

620 This approximate posterior can be used to obtain an approximate predictive  
621 distribution for the class label  $y^*$ :

$$\begin{aligned} p(y^*|\mathbf{x}^*, \mathbf{y}) &= \int p(y^*|\mathbf{x}^*, \mathbf{f}^*)p(\mathbf{f}^*|\mathbf{y})d\mathbf{f}^* \\ &= \int p(y^*|\mathbf{x}^*, \mathbf{f}^*) \prod_{k=1}^C \mathcal{N}(f^k(\mathbf{x}^*)|\mathbf{m}_k^*, \mathbf{V}_k^*)d\mathbf{f}^* \\ &= \int \left[ (1 - \epsilon) \prod_{k \neq y^*} \Theta(f^{y^*}(\mathbf{x}^*) - f^k(\mathbf{x}^*)) + \frac{\epsilon}{C} \right] \\ &\quad \prod_{k=1}^C \mathcal{N}(f^k(\mathbf{x}^*)|m_k^*, v_k^*)d\mathbf{f}^* \\ &= \int \left[ (1 - \epsilon) \prod_{k \neq y^*} \Theta(f^{y^*}(\mathbf{x}^*) - f^k(\mathbf{x}^*)) + \frac{\epsilon}{C} \right] \\ &\quad \prod_{k \neq y^*} \mathcal{N}(f^k(\mathbf{x}^*)|\mathbf{m}_k^*, \mathbf{V}_k^*)d\mathbf{f}^* \mathcal{N}(f^{y^*}(\mathbf{x}^*)|\mathbf{m}_{y^*}^*, \mathbf{V}_{y^*}^*) \\ &= \int \left[ (1 - \epsilon) \prod_{k \neq y^*} \Phi\left(\frac{f^{y^*}(\mathbf{x}^*) - \mathbf{m}_k^*}{\sqrt{v_k^*}}\right) + \frac{\epsilon}{C} \right] \\ &\quad \mathcal{N}(f^{y^*}(\mathbf{x}^*)|m_{y^*}^*, v_{y^*}^*)df^{y^*}(\mathbf{x}^*), \end{aligned} \quad (\text{A.34})$$

622 where  $\Phi(\cdot)$  is the cumulative distribution function of a Gaussian distribu-  
 623 tion. This is an integral in one dimension and can easily be approximated by  
 624 quadrature techniques.

## 625 **Appendix B. Details of the UCI Datasets**

626 Table B.2 shows the characteristics of the datasets considered from the  
 627 UCI repository in the main document. This table shows, for each problem,  
 628 the number of samples, the number of attributes and the number of class  
 629 labels.

Table B.2: Characteristics of the datasets from the UCI Repository.

<b>Dataset</b>	<b>#Instances</b>	<b>#Attributes</b>	<b>#Classes</b>
Glass	214	9	6
New-thyroid	215	5	3
Satellite	6435	36	6
Svmguide2	391	20	3
Vehicle	846	18	4
Vowel	540	10	6
Waveform	1000	21	3
Wine	178	13	3

## 630 **Appendix C. Comparison to Baseline Methods**

631 In this section we compare the PEP algorithm ( $\alpha \rightarrow 0$ ,  $\alpha = 0.5$  and  
 632  $\alpha = 1$ ) with three baseline methods: label regression, Laplace approximation  
 633 and a MCMC method that uses Gibbs sampling. We have performed these

634 experiments on the 8 UCI repository datasets that are summarized in Table  
635 B.2.

636 Label regression implementation uses the inducing point approximation  
637 and EP algorithm. Note that in the regression case, EP results in exact  
638 inference, as the likelihood factors are Gaussian [1]. The approximate factors  
639 are updated via regular EP updates, and the model hyper-parameters are  
640 optimized by gradient ascent with an adaptive learning rate (described in  
641 Appendix A.7). We consider three values for the number of inducing points  
642  $M$ . Namely, 5%, 10% and 20% of the number of training data  $N$ . We report  
643 averages over 100 repetitions.

644 In the case of Laplace and MCMC, we first obtain the model hyper-  
645 parameters using EP (PEP with  $\alpha = 1$ ) and then we train the methods  
646 to optimize the approximation. This is done because learning the hyper-  
647 parameters with these two methods is not scalable.

648 For Laplace, the gradients of the approximation to the marginal likelihood  
649 w.r.t. to the hyper-parameters cannot be computed efficiently using sparse ap-  
650 proximations, since they have an explicit dependence on the hyper-parameters  
651 and an indirect dependence through the mode [43]. This is precisely why  
652 there are no works in the literature that use Laplace approximation with  
653 sparse GPs. The Laplace approximation uses the softmax likelihood function.

654 The MCMC method that we have considered used Gibbs sampling. Gibbs  
655 sampling generates samples from the joint target distribution by replacing  
656 the value of one of the variables by a value drawn from the distribution of  
657 that variable conditioned on the values of the remaining variables [43]. This  
658 method is asymptotically unbiased.

659 In Table C.3 it is shown the test error and in Table C.4 we report the  
660 test log-likelihood. By looking at the results, we observe that performance of  
661 PEP is similar to the one of MCMC, so the predictive distribution of PEP is  
662 fairly good.

663 In conclusion, PEP gives good predictive distributions with the chosen  
664 likelihood function, and better than using the softmax. The softmax likelihood  
665 can be more robust than considering Gaussian noise like this work, but it  
666 makes inference more complicated, and the lack of robustness can be partially  
667 compensated by using the robust-max likelihood, where we introduce some  
668 noise in the labels by considering possible labelling errors with probability  $\epsilon$ ,  
669 even if  $\epsilon$  is small.

## 670 **Appendix D. Additional Experimental Results**

671 In this section we add some extra experimental results that did not fit in  
672 the main manuscript. In Figure D.8 we show the mean test error rank for  
673 each of the proposed methods and several values of  $\alpha$ . We report averages  
674 over 8 datasets from the UCI repository and 20 splits. Results for PEP and  
675 APEP are similar to the ones in the main manuscript in terms of the negative  
676 test log likelihood. However, ARPEP seems to give better results in terms of  
677 the test error with  $\alpha = 0.8$  or  $\alpha = 0.9$ .

	Problem	MCMC	Laplace	Label	PEP	PEP	PEP
				Regression	( $\alpha \rightarrow 0$ )	( $\alpha = 0.5$ )	( $\alpha = 1$ )
M = 5%	glass	0.32 ± 0.01	<b>0.26 ± 0.01</b>	0.44 ± 0.01	0.36 ± 0.01	0.33 ± 0.01	0.36 ± 0.01
	new-thyroid	0.05 ± 0	0.05 ± 0	0.15 ± 0.01	<b>0.03 ± 0</b>	0.04 ± 0	0.12 ± 0.01
	satellite	0.11 ± 0	0.12 ± 0	0.26 ± 0	0.11 ± 0	<b>0.11 ± 0</b>	0.11 ± 0
	svmguid2	0.18 ± 0.01	0.18 ± 0.01	0.21 ± 0.01	0.19 ± 0.01	<b>0.17 ± 0.01</b>	0.24 ± 0.01
	vehicle	0.19 ± 0	0.22 ± 0.01	0.25 ± 0	0.18 ± 0.01	<b>0.18 ± 0</b>	0.19 ± 0
	vowel	0.08 ± 0	0.15 ± 0.01	0.29 ± 0.01	0.06 ± 0	<b>0.05 ± 0</b>	0.09 ± 0.01
	waveform	<b>0.16 ± 0</b>	0.16 ± 0	0.26 ± 0	0.18 ± 0	0.16 ± 0	0.22 ± 0
	wine	0.03 ± 0	0.04 ± 0	0.03 ± 0	0.03 ± 0	<b>0.03 ± 0</b>	0.05 ± 0.01
	<b>Avg. Time</b>	5.02 ± 0.18	62.64 ± 3.77	131.61 ± 9.78	1684.83 ± 56.14	1625.81 ± 59.68	1495.37 ± 71.49
M = 10%	glass	0.32 ± 0.01	<b>0.25 ± 0.01</b>	0.43 ± 0.01	0.36 ± 0.01	0.31 ± 0.01	0.33 ± 0.01
	new-thyroid	0.04 ± 0	0.04 ± 0	0.12 ± 0.01	<b>0.03 ± 0</b>	0.03 ± 0	0.08 ± 0.01
	satellite	<b>0.11 ± 0</b>	0.12 ± 0	0.24 ± 0	0.11 ± 0	0.11 ± 0	0.11 ± 0
	svmguid2	0.18 ± 0.01	<b>0.18 ± 0.01</b>	0.22 ± 0.01	0.2 ± 0.01	0.18 ± 0.01	0.19 ± 0.01
	vehicle	0.19 ± 0	0.22 ± 0.01	0.23 ± 0	<b>0.17 ± 0</b>	0.17 ± 0	0.18 ± 0
	vowel	0.05 ± 0	0.19 ± 0.01	0.21 ± 0.01	0.04 ± 0	<b>0.03 ± 0</b>	0.05 ± 0
	waveform	<b>0.16 ± 0</b>	0.16 ± 0	0.26 ± 0	0.18 ± 0	0.17 ± 0	0.19 ± 0
	wine	<b>0.02 ± 0</b>	0.03 ± 0	0.03 ± 0	0.03 ± 0	0.02 ± 0	0.03 ± 0
	<b>Avg. Time</b>	9.87 ± 0.36	202.17 ± 14.77	153.83 ± 9.79	1865.4 ± 68.1	1814.39 ± 78.13	1724.34 ± 82.62
M = 20%	glass	0.32 ± 0.01	<b>0.24 ± 0.01</b>	0.39 ± 0.01	0.36 ± 0.01	0.31 ± 0.01	0.32 ± 0.01
	new-thyroid	0.04 ± 0	0.04 ± 0	0.1 ± 0.01	0.04 ± 0.01	<b>0.03 ± 0</b>	0.06 ± 0.01
	satellite	<b>0.11 ± 0</b>	0.12 ± 0	0.23 ± 0	0.11 ± 0	0.11 ± 0	0.11 ± 0
	svmguid2	0.18 ± 0.01	0.18 ± 0.01	0.23 ± 0.01	0.19 ± 0.01	0.18 ± 0.01	<b>0.18 ± 0.01</b>
	vehicle	0.17 ± 0	0.22 ± 0	0.22 ± 0	0.17 ± 0.01	<b>0.16 ± 0</b>	0.17 ± 0
	vowel	0.03 ± 0	0.23 ± 0.01	0.11 ± 0	0.03 ± 0	<b>0.02 ± 0</b>	0.03 ± 0
	waveform	0.16 ± 0	0.16 ± 0	0.26 ± 0	0.18 ± 0	0.17 ± 0	<b>0.16 ± 0</b>
	wine	0.03 ± 0	0.03 ± 0	0.04 ± 0	0.02 ± 0	<b>0.02 ± 0</b>	0.02 ± 0
	<b>Avg. Time</b>	22.89 ± 0.71	766.85 ± 58.39	224.41 ± 11.01	2169.13 ± 93.34	2138.39 ± 82.94	2073.18 ± 108.69

Table C.3: Average test error for each method and average training time in seconds.

	Problem	MCMC	Laplace	Label	PEP	PEP	PEP
				Regression	( $\alpha \rightarrow 0$ )	( $\alpha = 0.5$ )	( $\alpha = 1$ )
M = 5%	glass	0.8 ± 0.02	<b>0.8 ± 0.01</b>	1.1 ± 0.03	2.05 ± 0.07	0.81 ± 0.02	0.9 ± 0.02
	new-thyroid	0.12 ± 0.01	0.26 ± 0.01	0.38 ± 0.02	0.11 ± 0.01	<b>0.09 ± 0.01</b>	0.35 ± 0.01
	satellite	<b>0.3 ± 0</b>	0.5 ± 0.01	0.74 ± 0	0.51 ± 0	0.31 ± 0	0.3 ± 0
	svmguide2	<b>0.53 ± 0.02</b>	0.54 ± 0.01	0.7 ± 0.02	0.9 ± 0.06	0.57 ± 0.02	0.65 ± 0.02
	vehicle	0.37 ± 0.01	0.6 ± 0.01	0.58 ± 0.01	0.57 ± 0.04	<b>0.36 ± 0.01</b>	0.37 ± 0.01
	vowel	0.27 ± 0.01	0.62 ± 0.01	0.7 ± 0.01	0.38 ± 0.02	<b>0.17 ± 0.01</b>	0.29 ± 0.01
	waveform	<b>0.37 ± 0</b>	0.39 ± 0	0.59 ± 0	0.67 ± 0.01	0.4 ± 0	0.69 ± 0.01
	wine	<b>0.08 ± 0.01</b>	0.14 ± 0.01	0.09 ± 0	0.1 ± 0.01	0.08 ± 0.01	0.49 ± 0.01
	<b>Avg. Time</b>	5.02 ± 0.19	62.64 ± 3.87	131.61 ± 10.43	1684.83 ± 49.43	1625.81 ± 58.99	1495.37 ± 67.77
M = 10%	glass	0.79 ± 0.02	<b>0.78 ± 0.01</b>	1.05 ± 0.03	1.97 ± 0.07	0.8 ± 0.02	0.8 ± 0.02
	new-thyroid	0.09 ± 0.01	0.15 ± 0	0.33 ± 0.02	0.1 ± 0.01	<b>0.08 ± 0.01</b>	0.31 ± 0.01
	satellite	0.3 ± 0	0.36 ± 0	0.71 ± 0	0.5 ± 0	0.32 ± 0	<b>0.29 ± 0</b>
	svmguide2	0.54 ± 0.02	<b>0.52 ± 0.01</b>	0.73 ± 0.03	0.88 ± 0.05	0.6 ± 0.03	0.56 ± 0.02
	vehicle	<b>0.36 ± 0</b>	0.59 ± 0.01	0.55 ± 0.01	0.54 ± 0.02	0.36 ± 0.01	0.36 ± 0.01
	vowel	0.21 ± 0	0.8 ± 0.02	0.55 ± 0.01	0.26 ± 0.02	<b>0.14 ± 0</b>	0.22 ± 0.01
	waveform	<b>0.37 ± 0</b>	0.38 ± 0	0.59 ± 0	0.69 ± 0.01	0.43 ± 0	0.62 ± 0.01
	wine	0.07 ± 0.01	0.13 ± 0.01	0.1 ± 0.01	0.08 ± 0.01	<b>0.07 ± 0.01</b>	0.39 ± 0.01
	<b>Avg. Time</b>	9.87 ± 0.36	202.17 ± 12.79	153.83 ± 11.07	1865.4 ± 73.57	1814.39 ± 70.72	1724.34 ± 89.06
M = 20%	glass	0.78 ± 0.02	<b>0.77 ± 0.01</b>	1 ± 0.02	1.94 ± 0.07	0.8 ± 0.02	0.79 ± 0.02
	new-thyroid	<b>0.09 ± 0.01</b>	0.15 ± 0.01	0.28 ± 0.03	0.16 ± 0.04	0.1 ± 0.01	0.27 ± 0.01
	satellite	0.29 ± 0	0.4 ± 0.01	0.69 ± 0	0.48 ± 0	0.32 ± 0	<b>0.29 ± 0</b>
	svmguide2	0.55 ± 0.02	<b>0.53 ± 0.01</b>	0.77 ± 0.03	0.78 ± 0.04	0.59 ± 0.03	0.55 ± 0.02
	vehicle	<b>0.35 ± 0.01</b>	0.6 ± 0.01	0.53 ± 0.01	0.53 ± 0.02	0.36 ± 0.01	0.35 ± 0.01
	vowel	0.2 ± 0	1.09 ± 0.02	0.37 ± 0.01	0.16 ± 0.02	<b>0.13 ± 0</b>	0.19 ± 0
	waveform	0.38 ± 0	<b>0.38 ± 0</b>	0.61 ± 0	0.7 ± 0.01	0.46 ± 0.01	0.53 ± 0.01
	wine	<b>0.07 ± 0.01</b>	0.18 ± 0.01	0.1 ± 0.01	0.08 ± 0.01	0.07 ± 0.01	0.32 ± 0.01
	<b>Avg. Time</b>	22.89 ± 0.86	766.85 ± 49.9	224.41 ± 12.44	2169.13 ± 82.92	2138.39 ± 90.18	2073.18 ± 106.53

Table C.4: Average negative test log likelihood for each method and average training time in seconds.

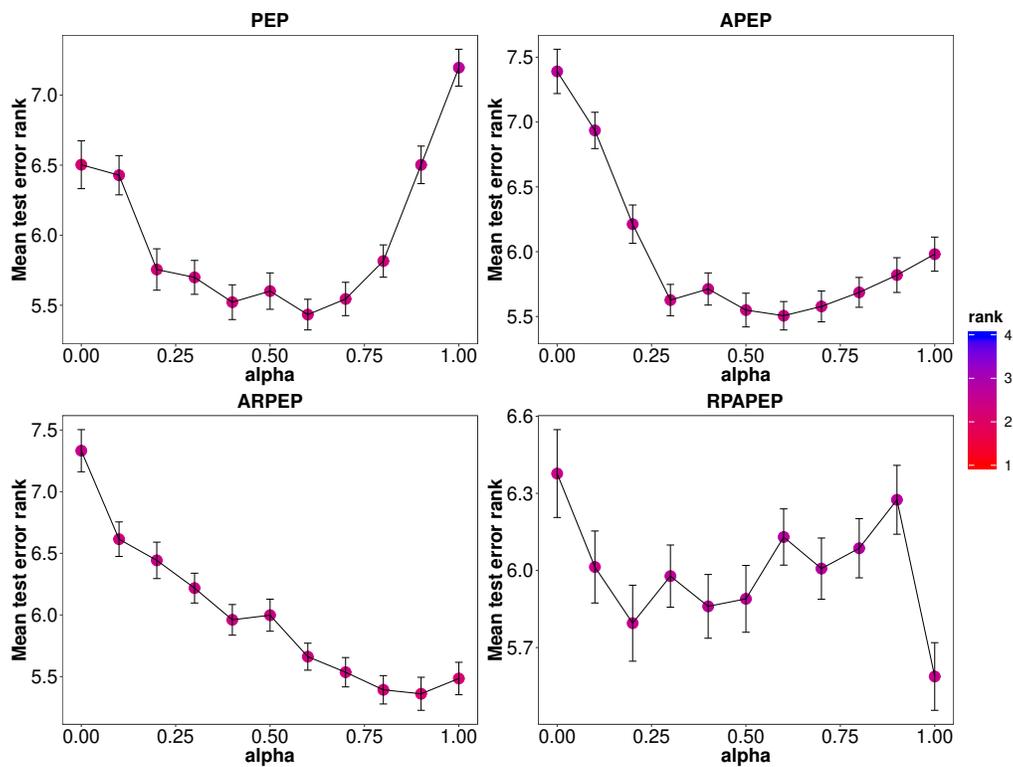


Figure D.8: Mean test error rank for different methods and different values of  $\alpha$ . The color of the points indicates the average rank of the method compared with the others. Best seen in color.

678 The next result is from the Satellite dataset of the UCI repository. We  
 679 show the performance as a function of the time in terms of the test error. It  
 680 gives similar results as for the negative test log likelihood (shown in the main  
 681 manuscript).

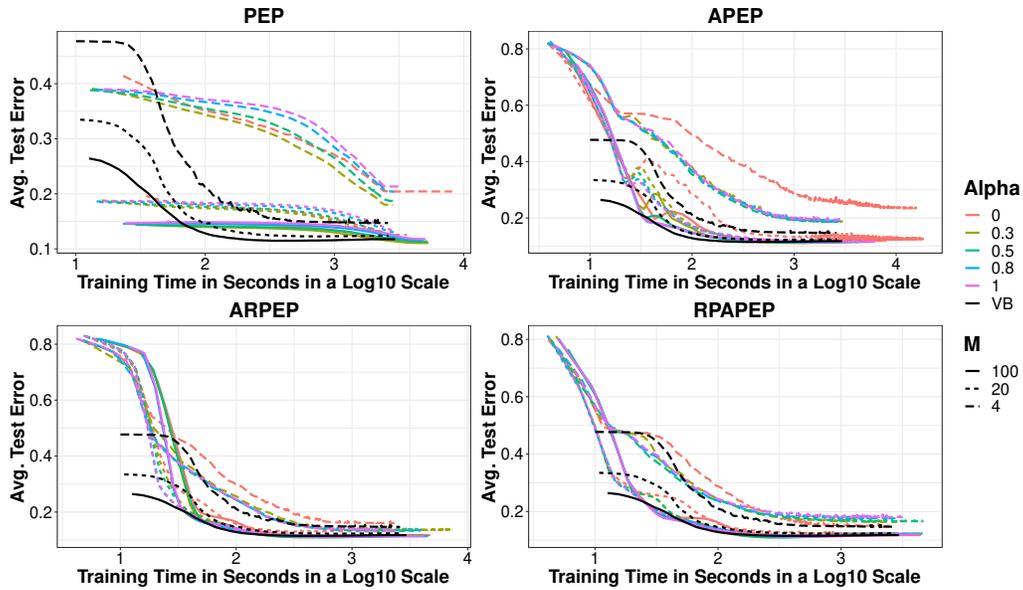


Figure D.9: Mean test error rank for different methods and different values of  $\alpha$ . Best seen in color.

682 In Figure D.10 we show the results for the MNIST dataset in terms of  
 683 the test error. Here we see that for the test error, values near  $\alpha \rightarrow 0$  do not  
 684 converge first, but instead intermediate values such as  $\alpha = 0.5$  tend to arrive  
 685 faster to the good solution.

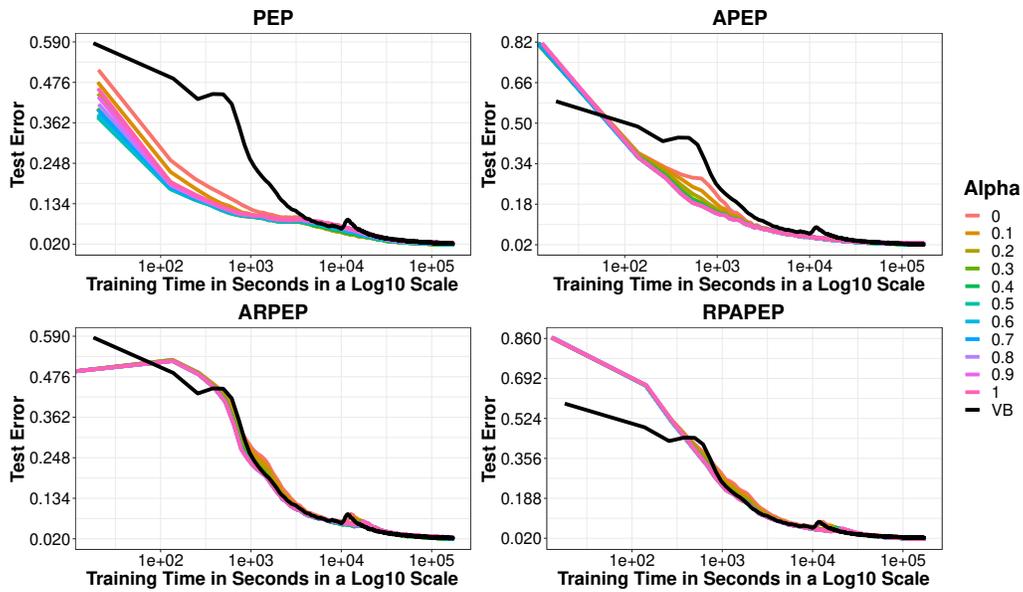


Figure D.10: Mean test error rank for different methods and different values of  $\alpha$  for MNIST dataset. Best seen in color.

686 Regarding the Airline Delays dataset, we observe similar results when  
 687 talking of the test error as in MNIST. The results are shown in Figure D.11

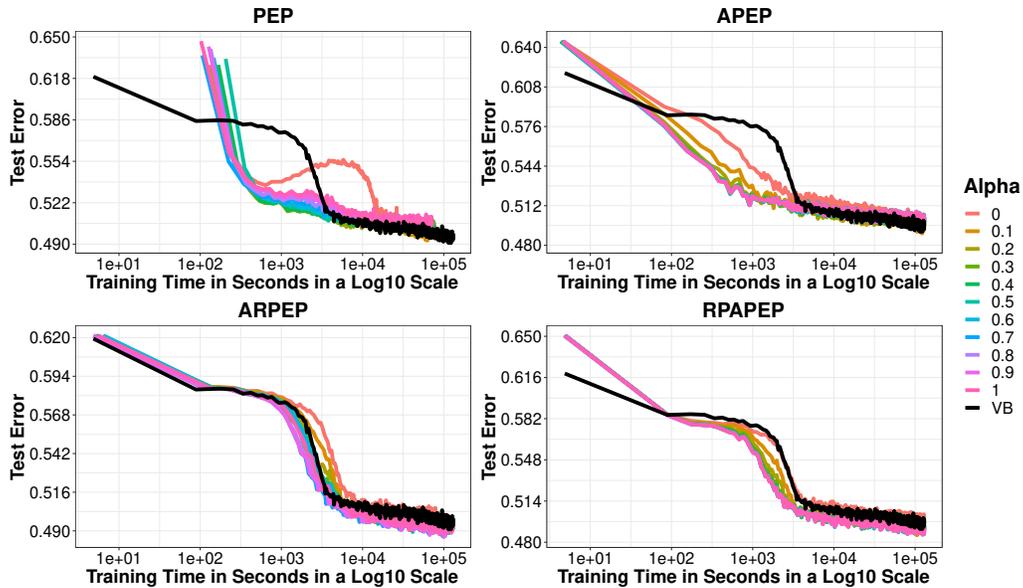


Figure D.11: Mean test error rank for different methods and different values of  $\alpha$  for Airline Delays dataset. Best seen in color.

688 **References**

689 **References**

690 [1] C. E. Rasmussen, C. K. I. Williams, Gaussian Processes for Machine  
 691 Learning (Adaptive Computation and Machine Learning), The MIT  
 692 Press, 2006.

693 [2] C. K. I. Williams, D. Barber, Bayesian classification with Gaussian pro-  
 694 cesses, IEEE Transactions on Pattern Analysis and Machine Intelligence  
 695 20 (1998) 1342–1351.

696 [3] H.-C. Kim, Z. Ghahramani, Bayesian Gaussian process classification

- 697 with the EM-EP algorithm, *IEEE Transactions on Pattern Analysis and*  
698 *Machine Intelligence* 28 (2006) 1948–1959.
- 699 [4] M. Girolami, S. Rogers, Variational Bayesian multinomial probit re-  
700 gression with Gaussian process priors, *Neural Computation* 18 (2006)  
701 1790–1817.
- 702 [5] K. M. A. Chai, Variational multinomial logit gaussian process, *Journal*  
703 *of Machine Learning Research* 13 (2012) 1745–1808.
- 704 [6] J. Riihimäki, P. Jylänki, A. Vehtari, Nested expectation propagation  
705 for Gaussian process classification with a multinomial probit likelihood,  
706 *Journal of Machine Learning Research* 14 (2013) 75–109.
- 707 [7] J. Quiñonero-Candela, C. E. Rasmussen, A unifying view of sparse  
708 approximate Gaussian process regression, *Journal of Machine Learning*  
709 *Research* 6 (2005) 1939–1959.
- 710 [8] J. Hensman, A. Matthews, Z. Ghahramani, Scalable variational Gaussian  
711 process classification, in: *Proceedings of the Eighteenth International*  
712 *Conference on Artificial Intelligence and Statistics*, 2015.
- 713 [9] J. Hensman, A. G. Matthews, M. Filippone, Z. Ghahramani, MCMC  
714 for variationally sparse Gaussian processes, in: *Advances in Neural*  
715 *Information Processing Systems* 28, 2015, pp. 1648–1656.
- 716 [10] C. Villacampa-Calvo, D. Hernández-Lobato, Scalable Multi-Class Gaus-  
717 sian Process Classification using Expectation Propagation, in: *International*  
718 *Conference on Machine Learning*, 2017, pp. 3550–3559.

- 719 [11] D. Hernández-Lobato, J. M. Hernández-Lobato, Scalable Gaussian  
720 process classification via expectation propagation, in: Proceedings of the  
721 19th International Conference on Artificial Intelligence and Statistics,  
722 2016, p. 168–176.
- 723 [12] T. D. Bui, J. Yan, R. E. Turner, A Unifying Framework for Gaussian Pro-  
724 cess Pseudo-Point Approximations using Power Expectation Propagation,  
725 Journal of Machine Learning Research 18 (2017) 1–72.
- 726 [13] T. Minka, Power EP, Technical Report, Microsoft Research, 2004.
- 727 [14] T. Minka, Divergence Measures and Message Passing, Technical Report,  
728 Microsoft Research, 2005.
- 729 [15] Y. Li, J. M. Hernández-Lobato, R. E. Turner, Stochastic expectation  
730 propagation, in: Advances in Neural Information Processing Systems 28,  
731 2015, pp. 2323–2331.
- 732 [16] T. Bui, D. Hernández-Lobato, J. M. Hernández-Lobato, Y. Li, R. Turner,  
733 Deep Gaussian processes for regression using approximate expectation  
734 propagation, in: International Conference on Machine Learning, 2016,  
735 pp. 1472–1481.
- 736 [17] Y. Li, Y. Gal, Dropout Inference in Bayesian Neural Networks with  
737 Alpha-divergences, in: International Conference on Machine Learning,  
738 2017, pp. 2052–2061.
- 739 [18] D. Hernández-Lobato, J. Hernández-Lobato, P. Dupont, Robust multi-  
740 class Gaussian process classification, in: Advances in Neural Information  
741 Processing Systems 24, 2011, pp. 280–288.

- 742 [19] C. J. Maddison, D. Tarlow, T. Minka, A\* sampling, in: Advances in  
743 Neural Information Processing Systems, 2014, pp. 3086–3094.
- 744 [20] M. N. Gibbs, D. J. C. Mackay, Variational Gaussian process classifiers,  
745 IEEE Transactions on Neural Networks 11 (2000) 1458–1464.
- 746 [21] E. Snelson, Z. Ghahramani, Sparse Gaussian processes using pseudo-  
747 inputs, in: Advances in Neural Information Processing Systems, 2006,  
748 pp. 1257–1264.
- 749 [22] A. Naish-Guzman, S. Holden, The generalized FITC approximation,  
750 in: Advances in Neural Information Processing Systems, 2008, pp. 1057–  
751 1064.
- 752 [23] T. Minka, Expectation propagation for approximate Bayesian infer-  
753 ence, in: Proceedings of the 17th Annual Conference on Uncertainty in  
754 Artificial Intelligence, 2001, pp. 362–36.
- 755 [24] M. Seeger, Expectation Propagation for Exponential Families, Technical  
756 Report, Department of EECS, University of California, Berkeley, 2006.
- 757 [25] M. Opper, C. Archambeau, The variational gaussian approximation  
758 revisited, Neural Computation 21 (2009) 786–792.
- 759 [26] M. Titsias, Variational learning of inducing variables in sparse Gaussian  
760 processes, in: International Conference on Artificial Intelligence and  
761 Statistics, 2009, pp. 567–574.
- 762 [27] E. Khan, S. Mohamed, K. P. Murphy, Fast bayesian inference for

- 763 non-conjugate gaussian process regression, in: Advances in Neural  
764 Information Processing Systems 25, 2012, pp. 3140–3148.
- 765 [28] A. G. d. G. Matthews, J. Hensman, R. Turner, Z. Ghahramani, On  
766 sparse variational methods and the kullback-leibler divergence between  
767 stochastic processes, in: Artificial Intelligence and Statistics, 2016, pp.  
768 231–239.
- 769 [29] S. Amari, Differential-geometrical methods in statistics, volume 28,  
770 Springer Science & Business Media, 1985.
- 771 [30] H. Zhu, R. Rohwer, Information Geometric Measurements of Generalisa-  
772 tion, Technical Report, Aston University, 1995.
- 773 [31] J. Hernández-Lobato, Y. Li, M. Rowland, T. Bui, D. Hernández-Lobato,  
774 R. Turner, Black-Box Alpha Divergence Minimization, in: International  
775 Conference on Machine Learning, 2016, pp. 1511–1520.
- 776 [32] T. Bui, Efficient Deterministic Approximate Bayesian Inference for Gaus-  
777 sian Process Models, Ph.D. thesis, 2017.
- 778 [33] A. Dezfouli, E. V. Bonilla, Scalable inference for Gaussian process  
779 models with black-box likelihoods, in: Advances in Neural Information  
780 Processing Systems, 2015, pp. 1414–1422.
- 781 [34] K. Krauth, E. V. Bonilla, K. Cutajar, M. Filippone, Autogp: Explor-  
782 ing the capabilities and limitations of Gaussian process models, in:  
783 Uncertainty in Artificial Intelligence, 2017.

- 784 [35] R. Sheth, Y. Wang, R. Khardon, Sparse variational inference for gener-  
785 alized GP models, in: International Conference on Machine Learning,  
786 2015, pp. 1302–1311.
- 787 [36] M. Lichman, UCI machine learning repository, 2013. URL: [http://](http://archive.ics.uci.edu/ml)  
788 [archive.ics.uci.edu/ml](http://archive.ics.uci.edu/ml).
- 789 [37] D. P. Kingma, J. Ba, ADAM: a method for stochastic optimization, in:  
790 Intrernational Conference on Learning Representations, 2015, pp. 1–15.
- 791 [38] M. Bauer, M. van der Wilk, C. E. Rasmussen, Understanding proba-  
792 bilistic sparse Gaussian process approximations, in: Advances in Neural  
793 Information Processing Systems 29, 2016, pp. 1533–1541.
- 794 [39] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning  
795 applied to document recognition, Proceedings of the IEEE 86 (1998)  
796 2278–2324.
- 797 [40] D. Hernández-Lobato, Prediction Based on Averages over Automatically  
798 Induced Learners: Ensemble Methods and Bayesian Techniques, Ph.D.  
799 thesis, 2010.
- 800 [41] K. B. Petersen, M. S. Pedersen, The Matrix Cookbook, Technical Uni-  
801 versity of Denmark, 2012.
- 802 [42] E. Snelson, Flexible and efficient Gaussian process models for machine  
803 learning, Ph.D. thesis, 2007.
- 804 [43] C. M. Bishop, Pattern Recognition and Machine Learning (Information  
805 Science and Statistics), Springer-Verlag New York, Inc., 2006.