
Scalable Multi-Class Gaussian Process Classification using Expectation Propagation

Carlos Villacampa-Calvo ^{*1} Daniel Hernández-Lobato ^{*1}

Abstract

This paper describes an expectation propagation (EP) method for multi-class classification with Gaussian processes that scales well to very large datasets. In such a method the estimate of the log-marginal-likelihood involves a sum across the data instances. This enables efficient training using stochastic gradients and mini-batches. When this type of training is used, the computational cost does not depend on the number of data instances N . Furthermore, extra assumptions in the approximate inference process make the memory cost independent of N . The consequence is that the proposed EP method can be used on datasets with millions of instances. We compare empirically this method with alternative approaches that approximate the required computations using variational inference. The results show that it performs similar or even better than these techniques, which sometimes give significantly worse predictive distributions in terms of the test log-likelihood. Besides this, the training process of the proposed approach also seems to converge in a smaller number of iterations.

1. Introduction

Gaussian processes (GPs) are non-parametric models that can be used to address multi-class classification problems (Rasmussen & Williams, 2006). These models become more expressive as the number of data instances N grows. They are also very useful to introduce prior knowledge in the learning problem, as many properties of the model are specified by a covariance function. Moreover, GPs provide an estimate of the uncertainty in the predictions made which may be critical in some applications. Neverthe-

less, in spite of these advantages, GPs scale poorly to large datasets because their training cost is $\mathcal{O}(N^3)$, where N is the number of instances. An additional challenge is that exact inference in these models is generally intractable and one has to resort to approximate methods in practice.

Traditionally, GP classification has received more attention in the binary case than in the multi-class setting (Kuss & Rasmussen, 2005; Nickisch & Rasmussen, 2008). The reason is that approximate inference is more challenging in the multi-class case where there is one latent function per class. To this one has to add more complicated likelihood factors, which often have the form of softmax functions or intractable Gaussian integrals. In spite of these difficulties, there have been several works addressing multi-class GP classification (Williams & Barber, 1998; Kim & Ghahramani, 2006; Girolami & Rogers, 2006; Chai, 2012; Riihimäki et al., 2013). Nevertheless, most of the proposed methods do not scale well with the size of the training set.

In the literature there have been some efforts to scale up GPs. These techniques often introduce a set of $M \ll N$ inducing points whose location is learnt alongside with the other model hyper-parameters. The use of inducing points in the model can be understood as an approximate GP prior with a low-rank covariance structure (Quiñonero-Candela & Rasmussen, 2005). When inducing points are considered, the training cost can be reduced to $\mathcal{O}(NM^2)$. This allows to address datasets with several thousands of instances, but not millions. The reason is the difficulty of estimating the model hyper-parameters, which is often done by maximizing an estimate of the log-marginal-likelihood. Because such an estimate does not involve a sum across the data instances, one cannot rely on efficient methods for optimization based on stochastic gradients and mini-batches.

A notable exception is the work of (Hensman et al., 2015a) which uses variational inference to approximate the calculations. Such a method allows for stochastic optimization and can address datasets with millions of instances. In this work we propose an alternative based on expectation propagation (EP) (Minka, 2001) and recent advances on binary GP classification (Hernández-Lobato & Hernández-Lobato, 2016). The proposed approach also allows for efficient training using mini-batches. This leads to a training

^{*}Equal contribution ¹Universidad Autónoma de Madrid, Madrid, Spain. Correspondence to: Carlos Villacampa-Calvo <carlos.villacampa@uam.es>.

cost that is $\mathcal{O}(CM^3)$, where C is the number of classes. An experimental comparison with the variational approach and related methods from the literature shows that the proposed approach has benefits both in terms of the training speed and the accuracy of the predictive distribution.

2. Scalable Multi-class Classification

Here we describe multi-class Gaussian process classification and the proposed method. Such a method uses the expectation propagation algorithm whose original description is modified to be more efficient both in terms of memory and computational costs. For this, we consider stochastic gradients to update the hyper-parameters and an approximate likelihood that avoids one-dimensional quadratures.

2.1. Multi-class Gaussian Process Classification

We consider a dataset of N instances in the form of a matrix of attributes $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T$ with labels $\mathbf{y} = (y_1, \dots, y_N)^T$, where $y_i \in \{1, \dots, C\}$ and $C > 2$ is the total number of different classes. The task of interest is to predict the class label of a new data instance \mathbf{x}_* .

A typical approach in multi-class Gaussian process (GP) classification is to assume the following labeling rule for y_i given \mathbf{x}_i : $y_i = \arg \max_k f^k(\mathbf{x}_i)$, for $k = 1, \dots, C$, where each $f^k(\cdot)$ is a non-linear latent function (Kim & Ghahramani, 2006). Define $\mathbf{f}^k = (f^k(\mathbf{x}_1), \dots, f^k(\mathbf{x}_N))^T \in \mathbb{R}^N$ and $\mathbf{f}_i = (f^1(\mathbf{x}_i), \dots, f^C(\mathbf{x}_i))^T \in \mathbb{R}^C$. The likelihood of $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^C)^T \in \mathbb{R}^{N \times C}$, $p(\mathbf{y}|\mathbf{f}) = \prod_{i=1}^N p(y_i|\mathbf{f}_i)$, is then a product of N factors of the form:

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

where $\Theta(\cdot)$ is the Heaviside step function. This likelihood takes value one if \mathbf{f} can explain the observed data and zero otherwise. Potential classification errors can be easily introduced in (1) by considering that each f^k has been contaminated with Gaussian noise with variance σ_k^2 . That is, $f^k(\mathbf{x}_i) = \hat{f}^k(\mathbf{x}_i) + \epsilon_i^k$, where $\epsilon_i^k \sim \mathcal{N}(0, \sigma_k^2)$.

In multi-class GP classification a GP prior is assumed for each function $f^k(\cdot)$ (Rasmussen & Williams, 2006). Namely, $f^k \sim \mathcal{GP}(0, c(\cdot, \cdot; \xi^k))$, where $c(\cdot, \cdot; \xi^k)$ is some covariance function with hyper-parameters ξ^k . Often these priors are assumed to be independent. That is, $p(\mathbf{f}) = \prod_{k=1}^C p(\mathbf{f}^k)$, where each $p(\mathbf{f}^k)$ is a multivariate Gaussian distribution. The task of interest is to make inference about \mathbf{f} and for that Bayes' rule is used: $p(\mathbf{f}|\mathbf{y}) = p(\mathbf{y}|\mathbf{f})p(\mathbf{f})/p(\mathbf{y})$, where $p(\mathbf{y})$ is a normalization constant (the marginal likelihood) which can be maximized to find good hyper-parameters ξ^k , for $k = 1, \dots, C$. However, because the likelihood in (1) is non-Gaussian, evaluating $p(\mathbf{y})$

and $p(\mathbf{f}|\mathbf{y})$ is intractable. Thus, these computations must be approximated. Often, one computes a Gaussian approximation to $p(\mathbf{f}|\mathbf{y})$ (Kim & Ghahramani, 2006). This results in a non-parametric classifier with training cost $\mathcal{O}(N^3)$, where N is the number of data instances.

To reduce the computational cost of the method described a typical approach is to consider a sparse representation for each GP. With this goal, one can introduce C datasets of $M \ll N$ inducing points $\bar{\mathbf{X}}^k = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)^T$, with associated values $\bar{\mathbf{f}}^k = (f^k(\bar{\mathbf{x}}_1^k), \dots, f^k(\bar{\mathbf{x}}_M^k))^T$ for $k = 1, \dots, C$ (Snelson & Ghahramani, 2006; Naish-Guzman & Holden, 2008). Given each $\bar{\mathbf{X}}^k$ the prior for \mathbf{f}^k is approximated as $p(\mathbf{f}^k) = \int p(\mathbf{f}^k|\bar{\mathbf{f}}^k)p(\bar{\mathbf{f}}^k|\bar{\mathbf{X}}^k)d\bar{\mathbf{f}}^k \approx \int [\prod_{i=1}^N p(f_i^k(\mathbf{x}_i)|\bar{\mathbf{f}}^k)]p(\bar{\mathbf{f}}^k|\bar{\mathbf{X}}^k)d\bar{\mathbf{f}}^k = p_{\text{FITC}}(\mathbf{f}^k|\bar{\mathbf{X}}^k)$, in which the conditional Gaussian distribution $p(\mathbf{f}^k|\bar{\mathbf{f}}^k)$ has been approximated by the factorizing distribution $\prod_{i=1}^N p(f_i^k(\mathbf{x}_i)|\bar{\mathbf{f}}^k)$. This approximation is known as the full independent training conditional (FITC) (Quiñonero-Candela & Rasmussen, 2005), and it leads to a Gaussian prior $p_{\text{FITC}}(\mathbf{f}^k|\bar{\mathbf{X}}^k)$ with a low-rank covariance matrix. This allows for approximate inference with cost $\mathcal{O}(NM^2)$. The inducing points $\{\bar{\mathbf{X}}^k\}_{k=1}^C$ can be regarded as hyper-parameters and can be learnt by maximizing the estimate of the marginal likelihood $p(\mathbf{y})$.

2.2. Method Specification and Expectation Propagation

The formulation of the previous section is limited because the estimate of the log-marginal-likelihood $\log p(\mathbf{y})$ cannot be expressed as a sum across the data instances. This makes infeasible the use of efficient methods based on stochastic optimization for finding the model hyper-parameters.

A recent work focusing on the binary case has shown that it is possible to obtain an estimate of $\log p(\mathbf{y})$ that involves a sum across the data instances if the values $\bar{\mathbf{f}}^k$ associated to the inducing points are not marginalized (Hernández-Lobato & Hernández-Lobato, 2016). We follow that work and consider the posterior approximation $p(\mathbf{f}|\mathbf{y}) \approx \int p(\mathbf{f}|\bar{\mathbf{f}})q(\bar{\mathbf{f}})d\bar{\mathbf{f}}$, where $\bar{\mathbf{f}} = (\bar{\mathbf{f}}^1, \dots, \bar{\mathbf{f}}^C)^T$, $p(\mathbf{f}|\bar{\mathbf{f}}) = \prod_{k=1}^C p(\mathbf{f}^k|\bar{\mathbf{f}}^k)$, we have defined $p(\bar{\mathbf{f}}) = \prod_{k=1}^C p(\bar{\mathbf{f}}^k|\bar{\mathbf{X}}^k)$, and q is a Gaussian approximation to $p(\mathbf{f}|\mathbf{y})$. This distribution q is obtained in three steps. First, we use on the exact posterior the FITC approximation:

$$\begin{aligned} p(\bar{\mathbf{f}}|\mathbf{y}) &= \frac{\int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\bar{\mathbf{f}})d\mathbf{f}p(\bar{\mathbf{f}})}{p(\mathbf{y})} \\ &\approx \frac{\int p(\mathbf{y}|\mathbf{f})p_{\text{FITC}}(\mathbf{f}|\bar{\mathbf{f}})d\mathbf{f}p(\bar{\mathbf{f}})}{p(\mathbf{y})} \\ &= \frac{[\prod_{i=1}^N \phi_i(\bar{\mathbf{f}})]p(\bar{\mathbf{f}})}{p(\mathbf{y})}, \end{aligned} \quad (2)$$

where we have defined $p_{\text{FITC}}(\mathbf{f}|\bar{\mathbf{f}}) = \prod_{i=1}^N \prod_{k=1}^C p(f^k(\mathbf{x}_i)|\bar{\mathbf{f}}^k) \approx p(\mathbf{f}|\bar{\mathbf{f}}) = \prod_{k=1}^C p(\mathbf{f}^k|\bar{\mathbf{f}}^k)$ and

$$\begin{aligned} \phi_i(\bar{\mathbf{f}}) &= \int [\prod_{k \neq y_i} \Theta(f^{y_i}(\mathbf{x}_i) - f^k(\mathbf{x}_i))] \\ &\quad \times [\prod_{k=1}^C p(f^k(\mathbf{x}_i)|\bar{\mathbf{f}}^k)] d\mathbf{f}_i, \end{aligned} \quad (3)$$

with $p(f^k(\mathbf{x}_i)|\bar{\mathbf{f}}^k) = \mathcal{N}(f^k(\mathbf{x}_i)|m_i^k, v_i^k)$, where

$$m_i^k = (\mathbf{k}_{\mathbf{x}_i, \bar{\mathbf{x}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{x}}^k, \bar{\mathbf{x}}^k}^k)^{-1} \bar{\mathbf{f}}^k, \quad (4)$$

$$s_i^k = \kappa_{\mathbf{x}_i, \mathbf{x}_i}^k - (\mathbf{k}_{\mathbf{x}_i, \bar{\mathbf{x}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{x}}^k, \bar{\mathbf{x}}^k}^k)^{-1} \mathbf{k}_{\mathbf{x}_i, \bar{\mathbf{x}}^k}^k. \quad (5)$$

In the previous expressions $\mathcal{N}(\cdot|\mu, \sigma^2)$ is the p.d.f. of a Gaussian with mean μ and variance σ^2 . Furthermore, $\mathbf{k}_{\mathbf{x}_i, \bar{\mathbf{x}}^k}^k$ is a vector with the covariances between $f^k(\mathbf{x}_i)$ and $\bar{\mathbf{f}}^k$; $\mathbf{K}_{\bar{\mathbf{x}}^k, \bar{\mathbf{x}}^k}^k$ is a $M \times M$ matrix with the cross covariances between $\bar{\mathbf{f}}^k$; and, finally, $\kappa_{\mathbf{x}_i, \mathbf{x}_i}^k$ is the prior variance of $f^k(\mathbf{x}_i)$.

A practical difficulty is that the integral in (3) is intractable. Although it can be evaluated using one-dimensional quadrature techniques (Hernández-Lobato et al., 2011), in this paper we follow a different approach. For that, we note that (3) is simply the probability that $f^{y_i}(\mathbf{x}_i) > f^k(\mathbf{x}_i)$ for $k \neq y_i$, given $\bar{\mathbf{f}}$. Let $f_i^{y_i} = f^{y_i}(\mathbf{x}_i)$ and $f_i^k = f^k(\mathbf{x}_i)$. The second step consists in approximating (3) as follows:

$$\begin{aligned} p(\bigcap_{k \neq y_i} f^{y_i} > f^k) &= p(f^{y_i} > f^1 | \mathcal{S}_1) \times p(f^{y_i} > f^2 | \mathcal{S}_2) \times \\ &\quad \dots \times p(f^{y_i} > f^{y_i-1} | \mathcal{S}_{y_i-1}) \times p(f^{y_i} > f^{y_i+1} | \mathcal{S}_{y_i+1}) \times \\ &\quad \dots \approx \prod_{k \neq y_i} p(f^{y_i} > f^k) = \prod_{k \neq y_i} \Phi(\alpha_i^k), \end{aligned} \quad (6)$$

where $\mathcal{S}_j = \bigcap_{k \notin \{1, \dots, j\} \cup \{y_i\}} f^{y_i} > f^k$, $\Phi(\cdot)$ is the c.d.f. of a standard Gaussian and $\alpha_i^k = (m_i^{y_i} - m_i^k) / \sqrt{s_i^{y_i} + s_i^k}$, with $m_i^{y_i}$, m_i^k , $s_i^{y_i}$ and s_i^k defined in (5). We have omitted in (6) the dependence on $\bar{\mathbf{f}}$ to improve the readability. The quality of this approximation is supported by the good experimental results obtained in Section 4. When (6) is replaced in (2) we get an approximate posterior distribution in which we can evaluate all the likelihood factors:

$$p(\bar{\mathbf{f}}|\mathbf{y}) \approx \frac{[\prod_{i=1}^N \prod_{k \neq y_i} \phi_i^k(\bar{\mathbf{f}})] p(\bar{\mathbf{f}})}{p(\mathbf{y})}, \quad (7)$$

where we have defined $\phi_i^k(\bar{\mathbf{f}}) = \Phi(\alpha_i^k)$.

The r.h.s. of (7) is intractable due to the non-Gaussian form of the likelihood factors. The third and last step uses expectation propagation (EP) (Minka, 2001) to get a Gaussian approximation q to (7). This approximation is obtained by replacing each ϕ_i^k with an approximate Gaussian factor $\tilde{\phi}_i^k$:

$$\begin{aligned} \tilde{\phi}_i^k(\bar{\mathbf{f}}) &= \tilde{s}_{i,k} \exp \left\{ -\frac{1}{2} (\bar{\mathbf{f}}^{y_i})^T \tilde{\mathbf{V}}_{i,k}^{y_i} \bar{\mathbf{f}}^{y_i} + (\bar{\mathbf{f}}^{y_i})^T \tilde{\mathbf{m}}_{i,k}^{y_i} \right\} \times \\ &\quad \exp \left\{ -\frac{1}{2} (\bar{\mathbf{f}}^k)^T \tilde{\mathbf{V}}_{i,k}^k \bar{\mathbf{f}}^k + (\bar{\mathbf{f}}^k)^T \tilde{\mathbf{m}}_{i,k}^k \right\}, \end{aligned} \quad (8)$$

where $\tilde{\mathbf{V}}_{i,k}^{y_i} = C_{i,k}^{1,y_i} \mathbf{v}_i^{y_i} (\mathbf{v}_i^{y_i})^T$, $\tilde{\mathbf{m}}_{i,k}^{y_i} = C_{i,k}^{2,y_i} \mathbf{v}_i^{y_i}$, $\tilde{\mathbf{V}}_{i,k}^k = C_{i,k}^1 \mathbf{v}_i^k (\mathbf{v}_i^k)^T$, $\tilde{\mathbf{m}}_{i,k}^k = C_{i,k}^2 \mathbf{v}_i^k$, and we have defined $\mathbf{v}_i^k = (\mathbf{k}_{\mathbf{x}_i, \bar{\mathbf{x}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{x}}^k, \bar{\mathbf{x}}^k}^k)^{-1}$. In (8) $C_{i,k}^{1,y_i}$, $C_{i,k}^{2,y_i}$, $C_{i,k}^1$, $C_{i,k}^2$ and $\tilde{s}_{i,k}$ are free parameters adjusted by EP. Because the precision matrices in (8) are one-rank (see the supplementary material for details), we only have to store in memory $\mathcal{O}(M)$ parameters for each $\tilde{\phi}_i^k$. The posterior approximation q is obtained by replacing in (7) each exact factor $\phi_{i,k}$ by the corresponding approximate factor $\tilde{\phi}_{i,k}$. That is, $q(\bar{\mathbf{f}}) = \prod_{i=1}^N \prod_{k \neq y_i} \tilde{\phi}_i^k(\bar{\mathbf{f}}) p(\bar{\mathbf{f}}) / Z_q$, where Z_q is a normalization constant that approximates the marginal likelihood $p(\mathbf{y})$. Because all the factors involved in the computation of q are Gaussian, and we assume independence among the latent functions of different classes in (8), q is a product of C multivariate Gaussians (on per class) on M dimensions.

In EP each $\tilde{\phi}_i^k$ is updated until convergence as follows: First, ϕ_i^k is removed from q by computing $q^{\setminus i,k} \propto q / \phi_i^k$. Because the Gaussian family is closed under the product and division operations, $q^{\setminus i,k}$ is also Gaussian with parameters given by the equations in (Roweis, 1999). Then, the Kullback-Leibler divergence between $Z_{i,k}^{-1} \phi_i^k q^{\setminus i,k}$ and q , i.e., $\text{KL}[Z_{i,k}^{-1} \phi_i^k q^{\setminus i,k} | q]$, is minimized with respect to q , where $Z_{i,k}$ is the normalization constant of $\phi_i^k q^{\setminus i,k}$. This is done by matching the moments of $Z_{i,k}^{-1} \phi_i^k q^{\setminus i,k}$. These moments can be obtained from the derivatives of $Z_{i,k}$ with respect to the parameters of $q^{\setminus i,k}$ (Seeger, 2006). After updating q , the new approximate factor is $\tilde{\phi}_{i,k} = Z_{i,k} q / q^{\setminus i,k}$. We update all the approximate factors at the same time, and reconstruct q afterwards by computing the product of all the $\tilde{\phi}_i^k$ and the prior, as in (Hernández-Lobato et al., 2011).

The EP approximation to the marginal likelihood is the normalization constant of q , Z_q . The log of its value is:

$$\log Z_q = g(\boldsymbol{\theta}) - g(\boldsymbol{\theta}_{\text{prior}}) + \sum_{i=1}^N \sum_{k \neq y_i} \log \tilde{s}_{i,k}, \quad (9)$$

where $\log \tilde{s}_{i,k} = \log Z_{i,k} + g(\boldsymbol{\theta}^{\setminus i,k}) - g(\boldsymbol{\theta})$; $\boldsymbol{\theta}$, $\boldsymbol{\theta}^{\setminus i,k}$, and $\boldsymbol{\theta}_{\text{prior}}$ are the natural parameters of q , $q^{\setminus i,k}$ and the prior, respectively; and $g(\boldsymbol{\theta})$ is the log-normalizer of a multi-variate Gaussian distribution with natural parameters $\boldsymbol{\theta}$.

It is possible to show that if EP converges, the gradient of $\log Z_q$ w.r.t the parameters of each $\tilde{\phi}_{i,k}$ is zero. Thus, the gradient of $\log Z_q$ w.r.t. a hyper-parameter ξ_j^k of the k -th covariance function (including the inducing points) is:

$$\frac{\partial \log Z_q}{\partial \xi_j^k} = (\boldsymbol{\eta}^T - \boldsymbol{\eta}_{\text{prior}}^T) \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \sum_{i=1}^N \sum_{k \neq y_i} \frac{\log Z_{i,k}}{\partial \xi_j^k}, \quad (10)$$

where $\boldsymbol{\eta}$ and $\boldsymbol{\eta}_{\text{prior}}$ are the expected sufficient statistics under q and the prior, respectively. Importantly, only the direct dependency of $\log Z_{i,k}$ on ξ_j^k has to be taken into account. See (Seeger, 2006). The dependency through $\boldsymbol{\theta}^{\setminus i,k}$, i.e., the natural parameters of $q^{\setminus i,k}$ can be ignored.

After obtaining q and finding the model hyper-parameters by maximizing $\log Z_q$, one can get an approximate predictive distribution for the label y_* of a new instance \mathbf{x}_* :

$$p(y_*|\mathbf{x}_*, \mathbf{y}) = \int p(y_*|\mathbf{f}_*, \bar{\mathbf{f}})q(\bar{\mathbf{f}})d\bar{\mathbf{f}}d\mathbf{f}_*, \quad (11)$$

where we have defined $\mathbf{f}_* = (f^1(\mathbf{x}_*), \dots, f^C(\mathbf{x}_*))^\top$, and $\int p(y_*|\mathbf{f}_*, \bar{\mathbf{f}})d\bar{\mathbf{f}}$ has the same form as the likelihood factor in (3). The resulting integral in (11) is again intractable. However, it can be approximated using a one-dimensional quadrature. See the supplementary material.

Because some simplifications occur when computing the derivatives of $\log Z_q$ w.r.t the inducing points, the total training time of EP is $\mathcal{O}(NM^2)$ while the total memory cost is $\mathcal{O}(NMC)$ (Snelson, 2007).

2.3. Scalable Expectation Propagation

Traditionally, for finding the model hyper-parameters with EP one re-runs EP until convergence (using the previous solution as the starting point), after each gradient ascent update of the hyper-parameters. The reason for this is that (10) is only true if EP has converged (*i.e.*, the approximate factors do not change any more). This approach is particularly inefficient initially, when there are strong changes to the model hyper-parameters, and EP may require several iterations to converge. Recently, a more efficient method has been proposed in (Hernández-Lobato & Hernández-Lobato, 2016). In that work the authors suggest to update both the approximate factors and the model hyper-parameters at the same time. Because we do not wait for EP to converge, one should ideally add to (10) extra terms to get the gradient. These terms account for the mismatch between the moments of $Z_{i,k}^{-1}\phi_i^k q^{\setminus i,k}$ and q . However, according to (Hernández-Lobato & Hernández-Lobato, 2016) these extra terms can be ignored and one can simply use (10) for an inner update of the hyper-parameters.




Figure 1. Estimate of $p(\mathbf{y})$ on the Vehicle dataset as a function of the training time for the proposed EP method when considering three different schemes to update the model hyper-parameters.

Figure 1 shows, for the Vehicle dataset from UCI repository (Lichman, 2013), the estimate of the marginal likelihood $\log Z_q$ with respect to the training time, for 250 updates of the hyper-parameters, and $M = N/5$. We compare three

methods: (i) re-running EP until convergence each time and using (10) to update the hyper-parameters (EP-outer); (ii) updating at the same time the approximate factors $\tilde{\phi}_i^k$ and the hyper-parameters with (10) (EP-inner-approx); and (iii) the same approach as the previous one, but using the exact gradient for the update instead of (10) (EP-inner-exact). All approaches successfully maximize $\log Z_q$. However, the inner updates are more efficient as they do not wait until EP converges. Moreover, using the approximate gradient is faster (it is cheaper to compute), and it gives almost the same results as the exact gradient.

2.3.1. STOCHASTIC EXPECTATION PROPAGATION

The memory cost of EP can be significantly reduced by a technique called stochastic EP (SEP) (Li et al., 2015). In SEP all the approximate factors $\tilde{\phi}_i^k$ are tied. This means that instead of storing their individual parameters, what is stored is their product, *i.e.*, $\tilde{\phi} = \prod_{i=1}^N \prod_{k \neq y_i} \tilde{\phi}_i^k$. A consequence of this is that we no longer have direct access to their individual parameters. This only affects the computation of the cavity distribution $q^{\setminus i,k}$ which now is obtained in an approximate way $q^{\setminus i,k} \propto q / \tilde{\phi}^{\frac{1}{n}}$, where n is the total number of factors and $\tilde{\phi}^{\frac{1}{n}}$ approximates each individual factor. Thus, SEP reduces the memory costs of EP by a factor of n . All the other steps are carried out as in the original EP algorithm, including the computation of $\log Z_q$ and its gradients. Figure 2 shows the differences between EP and SEP on a toy example. When SEP is used in the proposed method, the memory cost is reduced to $\mathcal{O}(CM^2)$.




Figure 2. (top) EP approximation of a distribution over the variable θ with complicated likelihood factors, but tractable prior. (bottom) SEP approximation of the same distribution.

2.3.2. TRAINING USING MINI-BATCHES

Both the estimate of the log-marginal-likelihood in (9) and its gradient in (10) contain a sum across the data instances. This allows to write an EP algorithm that processes mini-batches of data, as in (Hernández-Lobato & Hernández-Lobato, 2016). For this, the data are split in mini-batches \mathcal{M}_j of size $S \ll N$, where N is the number of instances. Given a mini-batch \mathcal{M}_j , we process all the approximate factors corresponding to that mini-batch, *i.e.*, $\{\{\tilde{\phi}_i^k\}_{k \neq y_i} : (\mathbf{x}_i, y_i) \in \mathcal{M}_j\}$. Then, we update the model hyper-parameters using a stochastic approximation of (10):

$$\frac{\partial \log Z_q}{\partial \xi_j^k} \approx (\boldsymbol{\eta}^\top - \boldsymbol{\eta}_{\text{prior}}^\top) \frac{\partial \boldsymbol{\theta}_{\text{prior}}}{\partial \xi_j^k} + \rho \sum_{i \in \mathcal{M}_j} \sum_{k \neq y_i} \frac{\log Z_{i,k}}{\partial \xi_j^k}, \quad (12)$$

where $\rho = N/|\mathcal{M}_j|$. We reconstruct q after each update of the approximate factors and each update of the hyper-

parameters. When using mini-batches of data, we update more frequently q and the hyper-parameters. The consequence is that the training cost is $\mathcal{O}(CM^3)$, assuming a constant number of updates until convergence. This training scheme can handle datasets with millions of instances.

3. Related Work

The likelihood used in (1) was first considered for multi-class Gaussian process classification in (Kim & Ghahramani, 2006). That work considers full non-parametric GP priors, which lead to a training cost that is $\mathcal{O}(CN^3)$. The consequence is that it can only address small classification problems. It is, however, straight forward to replace the non-parametric GP priors with the FITC approximate priors $p_{\text{FITC}}(\mathbf{f}^k | \bar{\mathbf{X}}^k)$ (Quiñonero-Candela & Rasmussen, 2005). These priors are obtained by marginalizing the latent variables \mathbf{f}^k associated to the inducing points $\bar{\mathbf{X}}^k$, as indicated in Section 2.1. This allows to address datasets with a few thousand instances. This is precisely the approach followed in (Naish-Guzman & Holden, 2008) to address binary GP classification problems. We refer to such an approach as the generalized FITC approximation (GFITC). Nevertheless, such an approach cannot use stochastic optimization. The reason is that the estimate of the log-marginal-likelihood (needed for hyper-parameter estimation) does not contain a sum across the instances. Thus, GFITC cannot scale well to very large datasets. Nevertheless, unlike the proposed approach, it can run expectation propagation over the exact likelihood factors in (1). In GFITC we follow the traditional approach and run EP until convergence before updating the hyper-parameters.

Multi-class GP classification for potentially huge datasets has also been considered in (Hensman et al., 2015b) using variational inference (VI). However, such an approach cannot use the likelihood in (1) since its logarithm is not well defined (note that it takes value zero for some values of \mathbf{f}_i). As an alternative, Hensman et al. (2015b) have considered the robust likelihood of (Hernández-Lobato et al., 2011):

$$p(y_i | \mathbf{f}_i) = (1 - \epsilon) \prod_{k \neq y_i} \Theta(f^{y_i}(\mathbf{x}_i) - f^k(\mathbf{x}_i)) + \frac{\epsilon}{C}, \quad (13)$$

where ϵ is the probability of a labeling error (in that case, y_i is chosen at random from the potential class labels). In (Hensman et al., 2015b) it is suggested to set $\epsilon = 10^{-3}$.

We now describe the VI approach in detail. Using (13) and the definitions of Section 2, we know that $p(\mathbf{y} | \bar{\mathbf{f}}) = \int p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \bar{\mathbf{f}}) d\mathbf{f}$. If we take the log and use Jensen's inequality we get the bound $\log p(\mathbf{y} | \bar{\mathbf{f}}) \geq \mathbb{E}_{p(\mathbf{f} | \bar{\mathbf{f}})} [\log p(\mathbf{y} | \mathbf{f})]$. Consider now a Gaussian approximation q to $p(\bar{\mathbf{f}} | \mathbf{y})$. Then,

$$\begin{aligned} \log p(\mathbf{y}) &= \int q(\bar{\mathbf{f}}) p(\mathbf{y} | \bar{\mathbf{f}}) p(\bar{\mathbf{f}}) / q(\bar{\mathbf{f}}) d\bar{\mathbf{f}} \\ &\geq \mathbb{E}_{q(\bar{\mathbf{f}})} [\log p(\mathbf{y} | \bar{\mathbf{f}})] - \text{KL}[q(\bar{\mathbf{f}}) || p(\bar{\mathbf{f}})], \quad (14) \end{aligned}$$

where we have used Jensen's inequality and KL is the Kull-

back Leibler divergence. If we use the first bound we get

$$\begin{aligned} \log p(\mathbf{y}) &\geq \mathbb{E}_{q(\bar{\mathbf{f}})} [\mathbb{E}_{p(\mathbf{f} | \bar{\mathbf{f}})} [\log p(\mathbf{y} | \mathbf{f})]] - \text{KL}[q(\bar{\mathbf{f}}) || p(\bar{\mathbf{f}})] \\ &\geq \mathbb{E}_{q(\mathbf{f})} [\log p(\mathbf{y} | \mathbf{f})] - \text{KL}[q(\bar{\mathbf{f}}) || p(\bar{\mathbf{f}})] \\ &\geq \sum_{i=1}^N \mathbb{E}_{q(\mathbf{f}_i)} [\log p(y_i | \mathbf{f}_i)] - \text{KL}[q(\bar{\mathbf{f}}) || p(\bar{\mathbf{f}})], \quad (15) \end{aligned}$$

where $q(\mathbf{f}) = \int p(\mathbf{f} | \bar{\mathbf{f}}) q(\bar{\mathbf{f}}) d\bar{\mathbf{f}}$ and the corresponding marginal over $\mathbf{f}_i = (f^1(\mathbf{x}_i), \dots, f^C(\mathbf{x}_i))^T$ is $q(\mathbf{f}_i) = \prod_{k=1}^C \mathcal{N}(f^k(\mathbf{x}_i) | \hat{m}_i^k, \hat{s}_i^k)$. Note that $q(\mathbf{f}_i)$ is Gaussian because $q(\mathbf{f})$ involves a Gaussian convolution. As in the proposed approach, $q(\bar{\mathbf{f}})$ is assumed to be a Gaussian factorizing over the latent functions $\bar{\mathbf{f}}^1, \dots, \bar{\mathbf{f}}^C$. However, its mean and covariance parameters, *i.e.*, $\{\mathbf{m}^k\}_{k=1}^C$ and $\{\mathbf{S}^k\}_{k=1}^C$ are found by maximizing (15). The parameters of $q(\mathbf{f}_i)$ are:

$$\begin{aligned} \hat{m}_i^k &= (\mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1} \mathbf{m}^k, \quad (16) \\ \hat{s}_i^k &= \kappa_{\mathbf{x}_i \mathbf{x}_i}^k - (\mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1} \mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k \\ &\quad + (\mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k)^T (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1} \mathbf{S}^k (\mathbf{K}_{\bar{\mathbf{X}}^k \bar{\mathbf{X}}^k}^k)^{-1} \mathbf{k}_{\mathbf{x}_i \bar{\mathbf{X}}^k}^k. \quad (17) \end{aligned}$$

Hensman et al. (2015b) consider Markov chain Monte Carlo (MCMC) to sample the hyper-parameters. Here we simply maximize (15) to find the hyper-parameters and the inducing points. The reason for this is that in very large datasets MCMC is not expected to give much better results. We refer to the described approach as VI. The objective in (15) contains a sum across the data instances. Thus, VI also allows for stochastic optimization and it results in the same cost as the proposed approach. However, the expectations in (15) must be approximated using one-dimensional quadratures. This is a drawback with respect to the proposed method which is free of any quadrature. Finally, there are some methods related to the VI approach just described. Dezfouli & Bonilla (2015) assume that q can be a mixture of Gaussians, and Chai (2012) uses a soft-max likelihood (but does not consider stochastic optimization). Both works need to introduce extra approximations.

In the literature there are other research works addressing multi-class Gaussian process classification. Some examples include (Williams & Barber, 1998; Girolami & Rogers, 2006; Hernández-Lobato et al., 2011; Henao & Winther, 2012; Riihimäki et al., 2013). These works employ expectation propagation, variational inference or the Laplace approximation to approximate the computations. Nevertheless, the corresponding estimate of the log-marginal-likelihood cannot be expressed as a sum across the data instances. This avoids using efficient techniques for optimization based on stochastic gradients. Thus, one cannot address very large datasets with these methods.

4. Experiments

We evaluate the performance of the method proposed in Section 2.2. We consider two versions of it. A first

one, using expectation propagation (EP). A second, using the memory efficient stochastic EP (SEP). EP and SEP are compared with the methods described in Section 3. Namely, GFITC and VI. All methods are codified in the R language (the source code is in the supplementary material), and they consider the same initial values for the model hyper-parameters (including the inducing points, that are chosen at random from the training instances). The hyper-parameters are optimized by maximizing the estimate of the marginal likelihood. A Gaussian covariance function with automatic relevance determination, an amplitude parameter and an additive noise parameter is employed.

4.1. Performance on Datasets from the UCI Repository

We evaluate the performance of each method on 8 datasets from the UCI repository (Lichman, 2013). The characteristics of the datasets are displayed in Table 4.1. We use batch training in each method (*i.e.*, we go through all the data to compute the gradients). Batch training does not scale to large datasets. However, it is preferred on small datasets like the ones considered here. We use 90% of the data for training and 10% for testing, except for *Satellite* which is fairly big, where we use 20% for training and 80% for testing. In *Waveform*, which is synthetic, we generate 1000 instances and split them in 30% for training and 70% for testing. Finally, in *Vowel* we consider only the points that belong to the 6 first classes. All methods are trained for 250 iterations using gradient ascent (GFITC and VI use 1-BFGS, EP and SEP use an adaptive learning rate described in the supplementary material). We consider three values for M , the number of inducing points. Namely 5%, 10% and 20% of the number of training instances. We report averages over 20 repetitions of the experiments.

Table 1. Characteristics of the datasets from the UCI Repository.

Dataset	#Instances	#Attributes	#Classes
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

Table 2 shows, for each value of M , the average negative test log-likelihood of each method with the corresponding error bars (test errors are shown in the supplementary material). The average training time of each method is also displayed. The best method (the lower the better) for each dataset is highlighted in bold face. We observe that the proposed approach, EP, obtains very similar results to those of GFITC, and sometimes it obtains the best results. The memory efficient version of EP, SEP, seems to provide similar results without reducing the performance. Regarding the computational cost, SEP is the

fastest method (between 2 and 3 times faster than GFITC). VI is slower as a consequence of the quadratures required by this method. VI also gives much worse results in some datasets, *e.g.*, *Glass*, *Svmguide2* and *Waveform*. This is related to the optimization of $\mathbb{E}_{q(\mathbf{f}_i)}[\log p(y_i|\mathbf{f}_i)]$ in (15), instead of $\log \mathbb{E}_{q(\mathbf{f}_i)}[p(y_i|\mathbf{f}_i)]$, which is closer to the data log-likelihood. In the EP objective in (9), $\sum_{k \neq y_i} \log Z_{i,k}$ is probably more similar to $\log \mathbb{E}_{q(\mathbf{f}_i)}[p(y_i|\mathbf{f}_i)]$. This explains the much better results obtained by EP and SEP.

Table 2. Average negative test log likelihood for each method and average training time in seconds on UCI repository datasets.

Problem	GFITC	EP	SEP	VI	
$M = 5\%$	Glass	0.61 ± 0.05	0.78 ± 0.06	0.77 ± 0.07	2.45 ± 0.14
	New-thyroid	0.06 ± 0.01	0.11 ± 0.03	0.06 ± 0.01	0.09 ± 0.02
	Satellite	0.33 ± 0.00	0.31 ± 0.00	0.33 ± 0.00	0.61 ± 0.01
	Svmguide2	0.63 ± 0.06	0.63 ± 0.06	0.67 ± 0.06	1.03 ± 0.08
	Vehicle	0.32 ± 0.01	0.34 ± 0.02	0.34 ± 0.02	0.76 ± 0.05
	Vowel	0.16 ± 0.01	0.25 ± 0.01	0.25 ± 0.01	0.41 ± 0.05
	Waveform	0.42 ± 0.01	0.36 ± 0.00	0.39 ± 0.01	0.89 ± 0.02
	Wine	0.08 ± 0.02	0.07 ± 0.01	0.08 ± 0.01	0.08 ± 0.02
Avg. Time	131 ± 3.11	53.8 ± 0.19	48.5 ± 0.97	157 ± 0.59	
$M = 10\%$	Glass	0.58 ± 0.05	0.74 ± 0.06	0.79 ± 0.07	2.18 ± 0.14
	New-thyroid	0.07 ± 0.01	0.06 ± 0.01	0.06 ± 0.01	0.05 ± 0.01
	Satellite	0.34 ± 0.00	0.30 ± 0.00	0.34 ± 0.00	0.58 ± 0.01
	Svmguide2	0.67 ± 0.05	0.67 ± 0.05	0.74 ± 0.07	0.90 ± 0.10
	Vehicle	0.33 ± 0.01	0.33 ± 0.02	0.34 ± 0.02	0.72 ± 0.04
	Vowel	0.14 ± 0.01	0.19 ± 0.01	0.19 ± 0.01	0.30 ± 0.04
	Waveform	0.42 ± 0.01	0.36 ± 0.01	0.41 ± 0.01	0.85 ± 0.01
	Wine	0.07 ± 0.01	0.06 ± 0.01	0.07 ± 0.01	0.07 ± 0.01
Avg. Time	264 ± 6.91	102 ± 0.64	96.6 ± 1.99	179 ± 0.78	
$M = 20\%$	Glass	0.6 ± 0.07	0.75 ± 0.06	0.81 ± 0.07	2.30 ± 0.15
	New-thyroid	0.07 ± 0.01	0.06 ± 0.01	0.05 ± 0.01	0.05 ± 0.01
	Satellite	0.34 ± 0.01	0.30 ± 0.00	0.36 ± 0.00	0.53 ± 0.01
	Svmguide2	0.67 ± 0.05	0.65 ± 0.06	0.74 ± 0.07	0.94 ± 0.08
	Vehicle	0.33 ± 0.01	0.33 ± 0.02	0.34 ± 0.02	0.63 ± 0.04
	Vowel	0.12 ± 0.01	0.16 ± 0.01	0.18 ± 0.01	0.15 ± 0.03
	Waveform	0.43 ± 0.01	0.37 ± 0.01	0.45 ± 0.01	0.80 ± 0.01
	Wine	0.07 ± 0.01	0.05 ± 0.01	0.06 ± 0.01	0.06 ± 0.02
Avg. Time	683 ± 17.3	228 ± 0.78	216 ± 2.88	248 ± 0.66	

4.2. Analysis of Inducing Point Learning

We generate a synthetic two dimensional problem with three classes by sampling the latent functions from the GP prior and applying the rule $y_i = \arg \max_k f^k(\mathbf{x}_i)$. The distribution of \mathbf{x}_i is uniform in the box $[-2.5, 2.5] \times [-2.5, 2.5]$. We consider 1000 training instances and a growing number of inducing points, *i.e.*, $M = 1$ to $M = 256$. The initial location of the inducing points is chosen at random and it is the same for all the methods. We are interested in the location of the inducing points after training. Thus, we set the other hyper-parameters to their true values (specified before generating the data) and we keep them fixed. All methods but VI are trained using batch methods during 2000 iterations. VI is trained using stochastic gradients for 2000 epochs (the batch version often gets stuck in local optima). We use ADAM with the default settings (Kingma & Ba, 2015), and 100 as the mini-batch size.

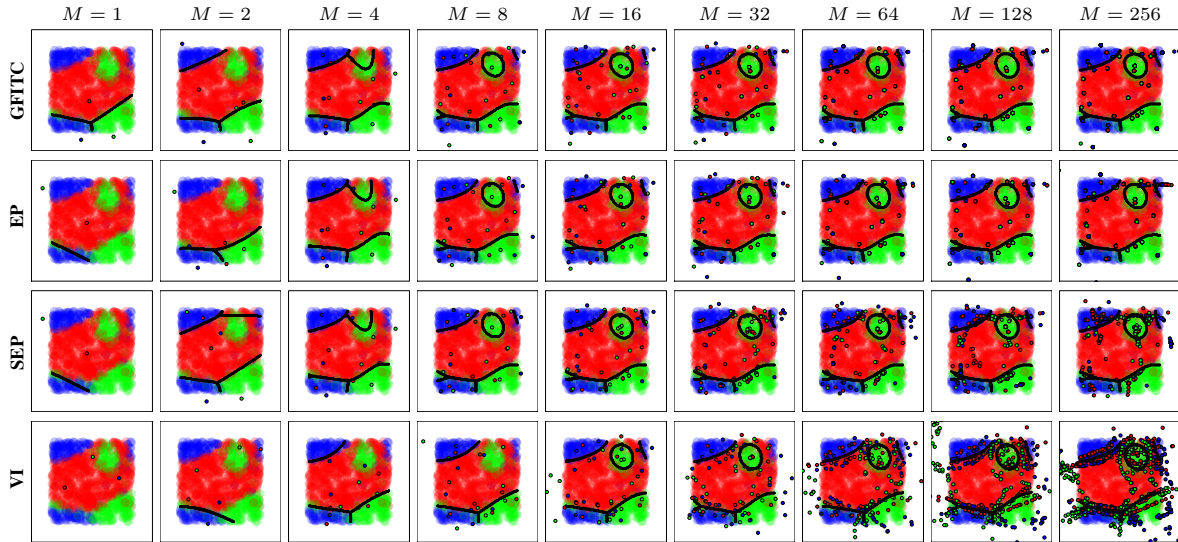


Figure 3. Decision boundaries and location of the inducing points after training for each method. GFITC, EP and SEP seem to place the inducing points one on top of each other. By contrast, VI prefers to place them near the decision boundaries. Best seen in color.

Figure 3 shows the location learnt by each method for the inducing points. Blue, red and green points represent the training data, black lines are decision boundaries and black border points are the inducing points. As we increase the number of inducing points the methods become more accurate. However, GFITC, EP and SEP identify decision boundaries that are better with a smaller number of inducing points. VI fails in this task. This is probably because VI updates the inducing-points with a bad estimate of q during the initial iterations. VI uses gradient steps to update q , which is less efficient than the EP updates (free of any learning rate). GFITC, EP and SEP overlap the inducing points, which can be seen as a pruning mechanism (if two inducing points are equal, it is like having only one). This has already been observed in regression problems (Bauer et al., 2016). By contrast, VI places the inducing points near the decision boundaries. This agrees with previous results on binary classification (Hensman et al., 2015a).

4.3. Performance as a Function of the Training Time

Figure 4 shows the negative test log-likelihood of each method as a function of the training time on the Satellite dataset (EP results are not shown since it performs equal to SEP). Training is done as in Section 4.1. We consider a growing number of inducing points $M = 4, 20, 100$ and report averages over 100 repetitions of the experiments. In all methods we use batch training. We observe that SEP is the method with the best performance at the lowest cost. Again, it is faster than GFITC because it optimizes q and the hyper-parameters at the same time, while GFITC waits until EP has converged to update the hyper-parameters. VI is not very efficient for small values of M , due to the quadratures. It also takes more time to get a good estimate of q , which is updated by gradient descent and is less ef-

ficient than the EP updates. Similar results are obtained in terms of the test error. See the supplementary material. However, in that case VI does not overfit the training data.




Figure 4. Negative test log-likelihood for GFITC, SEP and VI on *Satellite* as a function of the training time. Best seen in color.

4.4. Performance When Using Stochastic Gradients

In very large datasets batch training is infeasible, and one must use mini-batches to update q and to approximate the required gradients. We evaluate the performance of each method on the MNIST dataset (LeCun et al., 1998) with $M = 200$ inducing points and mini-batches with 200 instances. This dataset has 60,000 instances for training and 10,000 for testing. The learning rate of each method is set using ADAM with the default parameters (Kingma & Ba, 2015). GFITC does not allow for stochastic optimization. Thus, it is ignored in the comparison. The test error and the negative test log-likelihood of each method is displayed in Figure 5 (top) as a function of the training time. In this larger dataset all methods perform similarly. However, EP and SEP take less time to converge than VI. SEP obtains a test error that is 2.08% and average negative test log-likelihood that is 0.0725. The results of VI are 2.02% and 0.0686, respectively. These results are similar to the

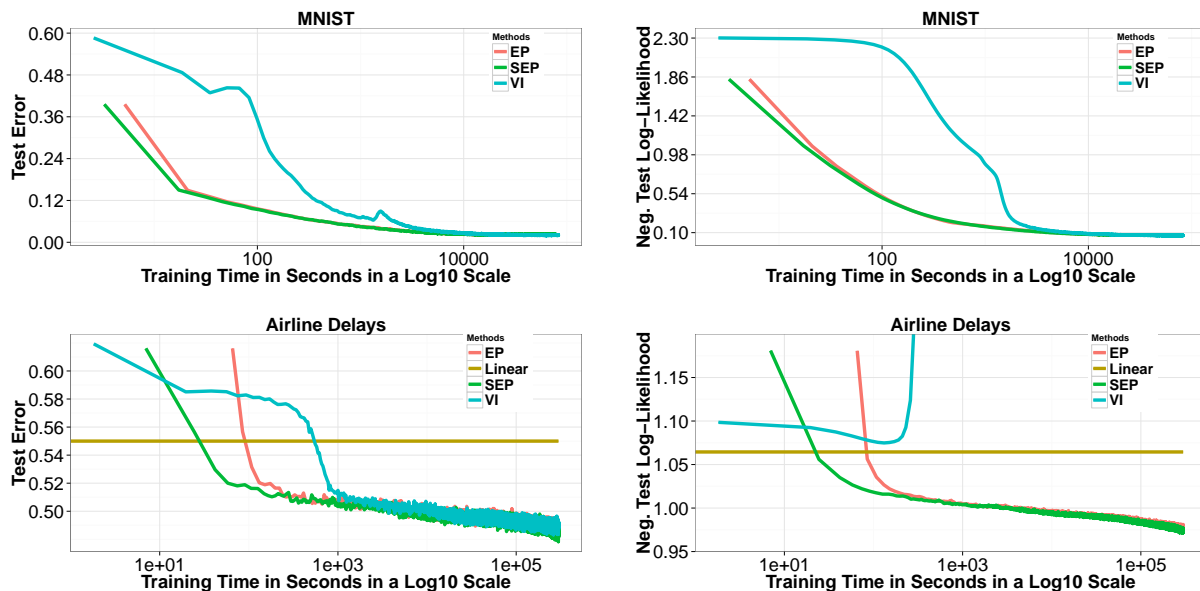


Figure 5. Average test error and average negative test log-likelihood for each method on the MNIST (top) and the Airline (bottom) dataset. In the Airline dataset a linear model based on logistic regression is included in the comparison. Best seen in color.

ones reported in (Hensman et al., 2015a) using $M = 500$.

A last experiment considers all flights within the USA between 01/2008 and 04/2008 (<http://stat-computing.org/dataexpo/2009>). The task is to classify the flights according to their delay using three classes: On time, more than 5 minutes of delay, or more than 5 minutes before time. We consider 8 attributes: age of the aircraft, distance covered, airtime, departure time, arrival time, day of the week, day of the month and month. After removing all instances with missing data 2, 127, 068 instances remain, from which 10, 000 are used for testing and the rest for training. We use the same settings as on the MNIST dataset and evaluate each method. The results obtained are shown in Figure 5 (bottom). We also report the performance of a logistic regression classifier. Again, all methods perform similarly in terms of test error. However, EP and SEP converge faster and quickly outperform the linear model. Importantly, the negative test log-likelihood of VI starts increasing at some point, which is again probably due to the optimization of $\mathbb{E}_{q(\mathbf{f}_i)}[\log p(y_i|\mathbf{f}_i)]$ in (15). The supplementary material has further evidence supporting this.

5. Conclusions

We have proposed the first method for multi-class classification with Gaussian processes, based on expectation propagation (EP), that scales well to very large datasets. Such a method uses the FITC approximation to reduce the number of latent variables in the model from $\mathcal{O}(N)$ to $\mathcal{O}(M)$, where $M \ll N$, and N is the number of data instances. For this, M inducing points are introduced for each latent function in the model. Importantly, the proposed method allows for stochastic optimization as the estimate of the log-

marginal-likelihood involves a sum across the data. Moreover, we have also considered a stochastic version of EP (SEP) to reduce the memory usage. When mini-batches and stochastic gradients are used for training, the computational cost of the proposed approach is $\mathcal{O}(CM^3)$, with C the number of classes. The memory cost is $\mathcal{O}(CM^2)$.

We have compared the proposed method with other approaches from the literature based on variational inference (VI) (Hensman et al., 2015b), and with the model considered by Kim & Ghahramani (2006), which has been combined with FITC approximate priors (GFITC) (Quiñonero-Candela & Rasmussen, 2005). The proposed approach outperforms GFITC in large datasets as this method does not allow for stochastic optimization, and in small datasets it produces similar results. The proposed method, SEP, is slightly faster than VI which also allows for stochastic optimization. In particular, VI requires one-dimensional quadratures which in small datasets are expensive. We have also observed that SEP converges faster than VI. This is probably because the EP updates, free of any learning rate, are more efficient for finding a good posterior approximation than the gradient ascent updates employed by VI.

An important conclusion of this work is that VI sometimes gives bad predictive distributions in terms of the test log-likelihood. The EP and SEP methods do not seem to have this problem. Thus, if one cares about accurate predictive distributions, VI should be avoided in favor of the proposed methods. In our experiments we have also observed that the proposed approaches tend to place the inducing points one on top of each other, which can be seen as an inducing point pruning technique (Bauer et al., 2016). By contrast, VI tends to place them near the decision boundaries.

Acknowledgements

The authors gratefully acknowledge the use of the facilities of Centro de Computación Científica (CCC) at Universidad Autónoma de Madrid. The authors also acknowledge financial support from Spanish Plan Nacional I+D+i, Grants TIN2013-42351-P, TIN2016-76406-P, TIN2015-70308-REDT and TEC2016-81900-REDT (MINECO/FEDER EU), and from Comunidad de Madrid, Grant S2013/ICE-2845.

References

- Bauer, M., van der Wilk, M., and Rasmussen, C. E. Understanding probabilistic sparse Gaussian process approximations. In *Advances in Neural Information Processing Systems 29*, pp. 1533–1541. 2016.
- Chai, K. M. A. Variational multinomial logit Gaussian process. *Journal of Machine Learning Research*, 13:1745–1808, 2012.
- Dezfouli, A. and Bonilla, E. V. Scalable inference for Gaussian process models with black-box likelihoods. In *Advances in Neural Information Processing Systems 28*, pp. 1414–1422. 2015.
- Girolami, M. and Rogers, S. Variational Bayesian multinomial probit regression with Gaussian process priors. *Neural Computation*, 18:1790–1817, 2006.
- Henao, R. and Winther, O. Predictive active set selection methods for Gaussian processes. *Neurocomputing*, 80: 10–18, 2012.
- Hensman, J., Matthews, A., and Ghahramani, Z. Scalable variational Gaussian process classification. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pp. 351–360, 2015a.
- Hensman, J., Matthews, A. G., Filippone, M., and Ghahramani, Z. MCMC for variationally sparse Gaussian processes. In *Advances in Neural Information Processing Systems 28*, pp. 1648–1656. 2015b.
- Hernández-Lobato, D. and Hernández-Lobato, J. M. Scalable Gaussian process classification via expectation propagation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pp. 168–176, 2016.
- Hernández-Lobato, D., ández Lobato, J.M., and Dupont, P. Robust multi-class Gaussian process classification. In *Advances in Neural Information Processing Systems 24*, pp. 280–288, 2011.
- Kim, H.-C. and Ghahramani, Z. Bayesian Gaussian process classification with the EM-EP algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28: 1948–1959, 2006.
- Kingma, D. P. and Ba, J. ADAM: a method for stochastic optimization. In *Inernational Conference on Learning Representations*, pp. 1–15, 2015.
- Kuss, M. and Rasmussen, C. E. Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Research*, 6:1679–1704, 2005.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86:2278–2324, 1998.
- Li, Y., Hernández-Lobato, J. M., and Turner, R. E. Stochastic expectation propagation. In *Advances in Neural Information Processing Systems 28*, pp. 2323–2331. 2015.
- Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Minka, T. Expectation propagation for approximate Bayesian inference. In *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence*, pp. 362–36, 2001.
- Naish-Guzman, A. and Holden, S. The generalized FITC approximation. In *Advances in Neural Information Processing Systems 20*, pp. 1057–1064. 2008.
- Nickisch, H. and Rasmussen, C. E. Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9:2035–2078, 2008.
- Quiñonero-Candela, J. and Rasmussen, C. E. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2006.
- Riihimäki, J., Jylänki, P., and Vehtari, A. Nested expectation propagation for Gaussian process classification with a multinomial probit likelihood. *Journal of Machine Learning Research*, 14:75–109, 2013.
- Roweis, S. Gaussian identities. Technical report, New York University, 1999.
- Seeger, M. Expectation propagation for exponential families. Technical report, Department of EECS, University of California, Berkeley, 2006.

Snelson, E. *Flexible and efficient Gaussian process models for machine learning*. PhD thesis, 2007.

Snelson, E. and Ghahramani, Z. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems 18*, pp. 1257–1264, 2006.

Williams, C. K. I. and Barber, D. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:1342–1351, 1998.