

---

This is an electronic reprint of the original article.  
This reprint may differ from the original in pagination and typographic detail.

Li, Rui; John, ST; Solin, Arno

## Improving Hyperparameter Learning under Approximate Inference in Gaussian Process Models

*Published in:*  
Proceedings of the 40th International Conference on Machine Learning

Published: 01/07/2023

*Document Version*  
Publisher's PDF, also known as Version of record

*Published under the following license:*  
CC BY

*Please cite the original version:*  
Li, R., John, ST., & Solin, A. (2023). Improving Hyperparameter Learning under Approximate Inference in Gaussian Process Models. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, & J. Scarlett (Eds.), *Proceedings of the 40th International Conference on Machine Learning* (pp. 19595-19615). (Proceedings of Machine Learning Research; Vol. 202). JMLR. <https://proceedings.mlr.press/v202/li23m.html>

---

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

# Improving Hyperparameter Learning under Approximate Inference in Gaussian Process Models

Rui Li<sup>1</sup> ST John<sup>1</sup> Arno Solin<sup>1</sup>

## Abstract

Approximate *inference* in Gaussian process (GP) models with non-conjugate likelihoods gets entangled with the *learning* of the model hyperparameters. We improve hyperparameter learning in GP models and focus on the interplay between variational inference (VI) and the learning target. While VI’s lower bound to the marginal likelihood is a suitable objective for inferring the approximate posterior, we show that a direct approximation of the marginal likelihood as in Expectation Propagation (EP) is a better learning objective for hyperparameter optimization. We design a hybrid training procedure to bring the best of both worlds: it leverages conjugate-computation VI for inference and uses an EP-like marginal likelihood approximation for hyperparameter learning. We compare VI, EP, Laplace approximation, and our proposed training procedure and empirically demonstrate the effectiveness of our proposal across a wide range of data sets.

## 1. Introduction

Gaussian processes (GPs, [Rasmussen & Williams, 2006](#)) provide a plug-and-play approach for inference and learning, with principled ways of incorporating prior knowledge over functions and quantifying uncertainty. While GP regression under a conjugate (Gaussian) likelihood can be carried out elegantly in closed form, we focus on the non-conjugate case, where exact inference is intractable. Training of the GP consists of *inferring* the approximate posterior and *learning* the hyperparameters of the model. For clarity, by *training* we refer to the combination of inference and learning.

In a supervised learning setting, GPs are typically trained to optimize performance on the training samples (as in em-

<sup>1</sup>Department of Computer Science, Aalto University, Finland, and Finnish Center for Artificial Intelligence (FCAI). Correspondence to: Rui Li <rui.li@aalto.fi>.




Figure 1. **Practical benefits** on IONOSPHERE: Marginal likelihood (top) acts as a training proxy for predictive density (bottom) of unseen future data. Our training objective produces a better point for prediction that also matches the MCMC baseline.

pirical risk minimization, [Vapnick, 1998](#)). Under the GP paradigm, the go-to solution to learning is finding  $\theta^*$  that maximizes the marginal likelihood. The marginal likelihood summarizes the probability that we would generate the observations  $\mathbf{y}$  with the model parameters  $\theta$  if we would sample over the prior. It is formed by *marginalizing* over the latent functions from the GP prior, thus also known as the *evidence*. Even if this does not capture all aspects of *generalization* (see discussion in [Vehtari et al., 2016](#); [Lotfi et al., 2022](#)), it is still used as a practical proxy for performance on unseen test points (see Fig. 1 for the proxy and test performance on the IONOSPHERE benchmark data set).

Under approximate inference, the marginalization step entangles the representation of the posterior with the learning target evaluation. The common approach is to assume an approximative Gaussian form for the posterior, so that the inference problem turns into finding a ‘good’ parameterization for the Gaussian (see [Wilkinson et al., 2023](#), for a recent discussion on linearization/Gaussianization approaches). The simplest approach is the so-called Laplace’s approximation (LA, [Williams & Barber, 1998](#)), which uses a second-order Taylor approximation. It is efficient but not very accurate. Variational inference (VI, [Oppel & Archam-](#)




Figure 2. **Log marginal likelihood / predictive density surfaces** for the IONOSPHERE data set by varying kernel magnitude  $\sigma$  and lengthscale  $\ell$ . The colour scale is the same in all plots:  $-0.8$   $\rightarrow$   $0$  (normalized by  $n$ ). Optimal hyperparameters are shown by a black marker. EP and our EP-like marginal likelihood estimation match the MCMC baseline better than VI or LA, thus providing a learning proxy. For prediction, our method still leverages the same variational representation as VI.

beau, 2009) and expectation propagation (EP, Minka, 2001) are two commonly used approximate inference methods for non-conjugate GP models, which have complementary advantages: VI optimizes a lower bound of the marginal likelihood, is easy to implement, straightforward to use, and the convex optimization problem is guaranteed to converge. However, it is known to underestimate variance (Bui et al., 2017). EP on the other hand requires implementation-wise tuning per likelihood and is not guaranteed to converge (Vehtari et al., 2020). However, it does provide a good approximation for the marginal likelihood (Kuss & Rasmussen, 2005; Nickisch & Rasmussen, 2008).

For model performance on unseen test data, the learning of hyperparameters plays a crucial role. Thus we strongly advocate against the common practice of jointly optimizing variational and hyperparameters using the ELBO, as the training target is only representative for the variational parameters. We build on work by Khan & Lin (2017) and Adam et al. (2021) that separate the learning of hyperparameters from inferring the variational parameters, and capture a link between VI and EP: the approximate posterior obtained through VI has exactly the same structure as the approximate posterior of EP. We obtain an EP-like marginal likelihood estimate from the VI approximate posterior for full and sparse GPs with no added computational cost. We propose a hybrid training procedure that combines the complementary advantages of natural-gradient VI and EP.

The contributions of this paper are as follows. (i) We improve generalizability in non-conjugate GP models with no extra computational cost by augmenting VI with an EP-like

learning target for hyperparameter learning. (ii) We demonstrate our EP-like learning target is closer to an MCMC baseline and thus provides a better learning objective. We empirically compare the quality of the approximate marginal likelihood in LA, EP, and VI, and our proposed learning target. (iii) We show our method improves generalizability via experiments in binary classification for full and sparse GP models and robust regression.

## 2. Approximate Inference

In this section, we review common approximate inference methods in Gaussian process (GP) models. GP models put a GP prior over functions:

$$\text{GP prior: } f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}')), \quad (1)$$

where  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$  is an input vector,  $\mu(\mathbf{x})$  is the mean function, and  $\kappa(\mathbf{x}, \mathbf{x}')$  is the covariance (kernel) function. This GP prior is linked to the data set  $\mathcal{D} = (\mathbf{X}, \mathbf{y}) = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$  of input-output pairs through a likelihood function that maps the latent function value  $f(\mathbf{x})$  to the observations. We assume the likelihood factorizes over observations:

$$\text{Likelihood: } \mathbf{y} | \mathbf{f} \sim \prod_{i=1}^n p(y_i | f(\mathbf{x}_i)). \quad (2)$$

The posterior is given by  $p(\mathbf{f} | \mathbf{y}; \boldsymbol{\theta}) \propto p(\mathbf{y} | \mathbf{f}; \boldsymbol{\theta}) p(\mathbf{f}; \boldsymbol{\theta})$ , where  $\boldsymbol{\theta}$  denotes the model (hyper)parameters of the likelihood, mean function, and kernel, and  $\mathbf{f}$  is the vector of function values evaluated at the inputs. Prediction at a new test input  $\mathbf{x}_*$  is obtained by computing the predictive distribution  $p(f(\mathbf{x}_*) | \mathcal{D}, \mathbf{x}_*)$ .

**Probabilistic inference** For (conjugate) Gaussian likelihoods,  $p(y_i | f_i) = \mathcal{N}(y_i | f(\mathbf{x}_i), \sigma_n^2)$ , the posterior is available in closed form as a Gaussian distribution. For non-Gaussian likelihood models the inference problem needs to be approached with *approximative* inference methods. Sampling schemes (see Sec. 5.1 for our baseline solution) can tackle this, but for efficient inference one typically employs an approximative Gaussian posterior of the form

$$\text{Approximate posterior: } q(\mathbf{f}) = \mathcal{N}(\mathbf{m}, \mathbf{S}). \quad (3)$$

Its ‘optimal’ parameterization (Opper & Archambeau, 2009) is given in terms of  $2n$  parameters  $(\alpha, \beta)$  such that  $\mathbf{m} = \mathbf{K}\alpha$  and  $\mathbf{S} = (\mathbf{K}^{-1} + \text{diag}(\beta))^{-1}$ , where  $\mathbf{K}$  is an  $n \times n$  matrix with  $\kappa(\mathbf{x}_i, \mathbf{x}_j)$  as the  $ij^{\text{th}}$  entry. The inference problem thus turns into (efficiently) finding a (good) representation of the posterior in terms of Eq. (3) by minimizing some measure of error. Typical approaches for this are the Laplace approximation (local linearisation of the problem), expectation propagation (approximately minimizing  $D_{\text{KL}}[p(\mathbf{f} | \mathbf{y}) \| q(\mathbf{f})]$  from approximate to true posterior), or variational inference (minimizing  $D_{\text{KL}}[q(\mathbf{f}) \| p(\mathbf{f} | \mathbf{y})]$ ). EP is expected to be the most accurate method (see discussion in Vehtari et al., 2016) and Laplace to have the smallest computational overhead.

**Learning under the GP paradigm** In probabilistic machine learning, ‘learning’ typically amounts to finding point estimates for the hyperparameters  $\theta$  in the likelihood, mean function, and kernel by optimizing w.r.t. the log marginal likelihood:

$$\text{Learning target: } \theta^* = \arg \max_{\theta} \log p(\mathbf{y}; \theta). \quad (4)$$

For Gaussian likelihoods, the marginal likelihood is available in closed form. For non-conjugate models, we can only optimize a proxy to the marginal likelihood  $p(\mathbf{y}; \theta) = \int p(\mathbf{y} | \mathbf{f}; \theta) p(\mathbf{f}; \theta) d\mathbf{f}$ , which depends on the approximate inference scheme and how it represents the posterior.

### 2.1. Laplace Approximation (LA)

A local Taylor expansion of the log posterior gives the Laplace approximation (LA, Williams & Barber, 1998). By defining  $\Psi(\mathbf{f}) = \log(p(\mathbf{y} | \mathbf{f}) p(\mathbf{f}; \theta))$ , the approximate posterior  $q(\mathbf{f})$  is obtained through a second-order Taylor expansion of  $\Psi(\mathbf{f})$  around its maximum at  $\hat{\mathbf{f}} = \arg \max_{\mathbf{f}} \Psi(\mathbf{f})$  (the posterior mode):  $p(\mathbf{f} | \mathbf{y}; \theta) \propto \exp(\Psi(\mathbf{f})) \approx \exp(\Psi(\hat{\mathbf{f}}) + \frac{1}{2}(\mathbf{f} - \hat{\mathbf{f}})^\top \nabla^2 \Psi(\mathbf{f})|_{\mathbf{f}=\hat{\mathbf{f}}}(\mathbf{f} - \hat{\mathbf{f}}))$ . This is proportional to  $\mathcal{N}(\mathbf{f} | \hat{\mathbf{f}}, \mathbf{A}^{-1}) = q(\mathbf{f})$ , where  $\mathbf{A} = -\nabla^2 \Psi(\mathbf{f})|_{\mathbf{f}=\hat{\mathbf{f}}}$  is the Hessian of the negative log posterior at  $\hat{\mathbf{f}}$ . The log marginal likelihood is approximated as

$$\begin{aligned} \log p(\mathbf{y}; \theta) &= \log \int \exp(\Psi(\mathbf{f})) d\mathbf{f} \\ &\approx \log \int \exp(\Psi(\hat{\mathbf{f}}) - \frac{1}{2}(\mathbf{f} - \hat{\mathbf{f}})^\top \mathbf{A}(\mathbf{f} - \hat{\mathbf{f}})) d\mathbf{f}. \end{aligned} \quad (5)$$

### 2.2. Expectation Propagation (EP)

Expectation Propagation (EP, Minka, 2001) is based on an approximation  $q(\mathbf{f})$  that factorizes in the same way as the target posterior  $p(\mathbf{f} | \mathbf{y}; \theta) \propto p(\mathbf{f}; \theta) \prod_{i=1}^n p(y_i | f_i; \theta)$ : each likelihood term is approximated with a *site* function  $t_i(f_i; \zeta_i)$ , and

$$q(\mathbf{f}; \theta, \zeta) \propto p(\mathbf{f}; \theta) \prod_{i=1}^n t_i(f_i; \zeta_i). \quad (6)$$

For GP models, the sites  $t_i(f_i; \zeta_i)$  are chosen to be (unnormalized) Gaussians, and hence the global approximation  $q(\mathbf{f})$  is also Gaussian. EP aims to minimize  $D_{\text{KL}}[p(\mathbf{f} | \mathbf{y}; \theta) \| q(\mathbf{f}; \theta, \zeta)]$  w.r.t.  $\zeta$ . This KL cannot be computed directly. Instead, EP updates the sites in an iterative fashion; the parameters of one site  $\zeta_i$  are tuned by minimizing the local Kullback–Leibler divergence

$$\begin{aligned} D_{\text{KL}}[p(y_i | f_i; \theta) p(\mathbf{f}; \theta) \prod_{j \neq i} t_j(f_j; \zeta_j) \\ \| t_i(f_i; \zeta_i) p(\mathbf{f}; \theta) \prod_{j \neq i} t_j(f_j; \zeta_j)], \end{aligned} \quad (7)$$

where in the first argument the  $n - 1$  other likelihood terms have been replaced by their current site approximation. The optimal values of  $\zeta_i$  in this step can be determined by matching the first two moments. This iterative process often works well in practice, but can be numerically unstable (e.g., for Student- $t$  likelihood) and is not guaranteed to converge in the general case (see Vehtari et al., 2020).

The log marginal likelihood is directly approximated as

$$\log p(\mathbf{y}; \theta) \approx \mathcal{L}_{\text{EP}}(\zeta, \theta) = \log \int p(\mathbf{f}; \theta) \prod_{i=1}^n t_i(f_i; \zeta_i) d\mathbf{f}, \quad (8)$$

which is known to lead to a good objective for learning hyperparameters (see Jylänki et al., 2011).

### 2.3. Variational Inference (VI)

Variational Inference (VI, Opper & Archambeau, 2009) approximates the GP posterior  $p(\mathbf{f} | \mathbf{y}; \theta)$  with a Gaussian distribution  $q(\mathbf{f}; \xi)$  parameterized by  $\xi$ . VI minimizes the reverse KL  $D_{\text{KL}}[q(\mathbf{f}; \xi) \| p(\mathbf{f} | \mathbf{y}; \theta)]$  by maximizing the following evidence lower bound (ELBO):

$$\begin{aligned} \log p(\mathbf{y}; \theta) \geq \mathcal{L}_{\text{VI}}(\xi, \theta) &= \sum_{i=1}^n \mathbb{E}_{q(f_i; \xi_i)} [\log p(y_i | f_i; \theta)] \\ &\quad - D_{\text{KL}}[q(\mathbf{f}; \xi) \| p(\mathbf{f}; \theta)], \end{aligned} \quad (9)$$

w.r.t. variational parameters  $\xi$ . VI optimizes a lower bound on the marginal likelihood, so is guaranteed to converge, which is a strength over EP. As known from Hensman et al. (2013) and motivated by Khan et al. (2013), it has been desirable to *not* use the optimal parameterization in terms of  $2n$  parameters, as the resulting optimization problem is

non-convex. Instead, it is common to declare a variational distribution over the full posterior,  $q(\mathbf{f}; \boldsymbol{\xi}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ , and optimize the ELBO w.r.t. this mean–covariance parameterization<sup>1</sup>  $\boldsymbol{\xi} = (\mathbf{m}, \mathbf{S})$  using a general-purpose optimizer (e.g. Adam, Kingma & Ba, 2015).

In practice, the same lower bound  $\mathcal{L}_{\text{VI}}(\boldsymbol{\xi}, \boldsymbol{\theta})$  is used to optimize variational parameters as well as hyperparameters, i.e., inference and learning are coupled into a single optimization. This approach is commonplace, even though it is well-known to result in biased hyperparameters (Kuss & Rasmussen, 2005; Nickisch & Rasmussen, 2008; Bui et al., 2017).

### 3. Learning in the Dual Parameterization

We design a hybrid training procedure that augments VI with an EP-like learning target for hyperparameter learning. Our work builds upon the dual parameterization (Khan & Lin, 2017). Because the Gaussian distribution is part of the exponential family, we can write the approximate posterior as  $q(\mathbf{f}) = \mathcal{N}(\mathbf{m}, \mathbf{S}) = \exp(\boldsymbol{\eta}^\top \mathbf{T}(\mathbf{f}) - a(\boldsymbol{\eta}))$ , where  $\boldsymbol{\eta} = (\mathbf{S}^{-1}\mathbf{m}, -\frac{1}{2}\mathbf{S}^{-1})$ ,  $\mathbf{T}(\mathbf{f}) = (\mathbf{f}, \mathbf{f}\mathbf{f}^\top)$  are the sufficient statistics, and  $\exp(-a(\boldsymbol{\eta}))$  is a normalization term. This leads to two additional parameterizations of  $q(\mathbf{f})$ : using the natural parameters  $\boldsymbol{\eta}$ , or using the expectation parameters  $\boldsymbol{\mu} = \mathbb{E}_{q(\mathbf{f})}[\mathbf{T}(\mathbf{f})] = (\mathbf{m}, \mathbf{S} + \mathbf{m}\mathbf{m}^\top)$ .

Khan & Lin (2017) showed that in the natural parameterization of the approximate posterior, *natural gradient descent* (NGD, Amari, 1998) (in the natural parameters space  $\boldsymbol{\eta}$ ) will have the same computational cost as ordinary gradient descent on  $\boldsymbol{\xi} = (\mathbf{m}, \mathbf{S})$ . The approximate posterior under this parameterization is

$$q(\mathbf{f}; \boldsymbol{\lambda}, \boldsymbol{\theta}) \propto p(\mathbf{f}; \boldsymbol{\theta}) \prod_{i=1}^n \underbrace{\exp\langle \boldsymbol{\lambda}_i, \mathbf{T}(f_i) \rangle}_{\triangleq t_i(f_i; \boldsymbol{\lambda}_i)}, \quad (10)$$

where  $\boldsymbol{\lambda}_i = \nabla_{\boldsymbol{\mu}_i} \mathbb{E}_{q(f_i; \boldsymbol{\lambda}_i, \boldsymbol{\theta})}[\log p(y_i | f_i; \boldsymbol{\theta})]$ . The natural parameters of the approximate posterior  $q(\mathbf{f})$  are  $\boldsymbol{\eta} = \boldsymbol{\lambda}_0 + \boldsymbol{\lambda}$ , where  $\boldsymbol{\lambda}_0 = (\mathbf{0}, -\frac{1}{2}\mathbf{K}^{-1})$  are the natural parameters of the prior  $p(\mathbf{f}; \boldsymbol{\theta})$  and  $\boldsymbol{\lambda}$  are the parameters of the likelihood approximation term  $t(\mathbf{f})$ . Then, we could also parameterize the approximate posterior with  $\boldsymbol{\lambda}$ , to which we refer as the ‘dual’  $\boldsymbol{\lambda}$  parameterization.

Crucially, the approximate posterior (10) has the same form as its EP counterpart (6). This links EP with VI, which is the starting point for our proposed learning objective. The similarity *per se* has been visible in, e.g., Chang et al. (2020); Adam et al. (2021), but it had not been explored further.

<sup>1</sup>In practice it may be beneficial to optimize in the *whitened* (or *non-centered*) parameterization  $\boldsymbol{\xi} = (\mathbf{m}', \mathbf{S}')$  s.t.  $\mathbf{m} = \mathbf{L}\mathbf{m}'$  and  $\mathbf{S} = \mathbf{L}\mathbf{S}'\mathbf{L}^\top$ , where  $\mathbf{L}$  = Cholesky( $\mathbf{K}$ ) (Gorinova et al., 2020).

**Algorithm 1** Training procedure for improved hyperparameter learning by a VEM-style iteration.

---

```

Initialize var. parameters  $\boldsymbol{\lambda}^{(0)}$  and hyperparameters  $\boldsymbol{\theta}^{(0)}$ 
Specify total training iteration  $K$ , the number of E-steps
and M-steps per iteration  $J_E, J_M$ 
for  $k = 1, 2, \dots, K$  do
    With  $J_E$  nat. grad. steps and learning rate  $\rho_E$ , optimize
     $\boldsymbol{\lambda}^{(k)} \leftarrow \arg \max_{\boldsymbol{\lambda}} \mathcal{L}_{\text{VI}}(\boldsymbol{\lambda}, \boldsymbol{\theta}^{(k-1)})$ 
    With  $J_M$  grad. steps and learning rate  $\rho_M$ , optimize
     $\boldsymbol{\theta}^{(k)} \leftarrow \arg \max_{\boldsymbol{\theta}} \mathcal{L}_{\text{EP}}(\boldsymbol{\lambda}^{(k)}, \boldsymbol{\theta})$ 
    if  $\mathcal{L}_{\text{EP}}(\boldsymbol{\lambda}^{(k)}, \boldsymbol{\theta}^{(k)}) < \mathcal{L}_{\text{EP}}(\boldsymbol{\lambda}^{(k-1)}, \boldsymbol{\theta}^{(k-1)})$  then
        return  $\boldsymbol{\lambda}^{(k-1)}, \boldsymbol{\theta}^{(k-1)}$ 
    end if
end for
return  $\boldsymbol{\lambda}^{(k)}, \boldsymbol{\theta}^{(k)}$ 
    
```

---

#### 3.1. Our Proposed Objective for Learning

Natural gradient descent can efficiently optimize the variational parameters  $\boldsymbol{\lambda}$ , and we can combine it with another optimizer for the hyperparameters, leading to a natural separation of the inference and learning steps (Salimbeni et al., 2018). As discussed by Adam et al. (2021), this can be seen as a Variational Expectation–Maximization (VEM) procedure. Under this setup, inference/learning is performed by alternating between optimizing the variational distribution in the  $\boldsymbol{\lambda}$  parameterization and taking gradient steps for finding  $\boldsymbol{\theta}$  by iterating the following steps at the  $k^{\text{th}}$  iteration:

$$\text{E-step (inference): } \boldsymbol{\lambda}^{(k+1)} \leftarrow \arg \max_{\boldsymbol{\lambda}} \mathcal{L}_E(\boldsymbol{\lambda}, \boldsymbol{\theta}^{(k)}),$$

$$\text{M-step (learning): } \boldsymbol{\theta}^{(k+1)} \leftarrow \arg \max_{\boldsymbol{\theta}} \mathcal{L}_M(\boldsymbol{\lambda}^{(k+1)}, \boldsymbol{\theta}),$$

where the objective for both the inference and learning steps is the ELBO in Eq. (9) under the  $\boldsymbol{\lambda}$  parameterization:  $\mathcal{L}_E \equiv \mathcal{L}_M \equiv \mathcal{L}_{\text{VI}}$ . Note: Even if the parameterization and optimization procedure are different, the inference and learning objective are the same as in Sec. 2.3 and typically expected to converge at the same optima.

VEM deals with the variational inference problem by casting inference into an optimization problem that is solved by NGD, which appears both principled and efficient. We conjecture that the ELBO in Eq. (9) is not the best objective. Conveniently, the dual parameterization in the VI posterior Eq. (10) is formed as a product of the prior and Gaussian sites  $t_i(f_i; \boldsymbol{\lambda}_i)$  just as in EP (cf. Eq. (6)). This provides a representation of the posterior that is directly EP-like and allows us to estimate the log marginal likelihood by plugging  $\boldsymbol{\lambda}_i$  from Eq. (10) into  $\zeta_i$  in Eq. (8):

$$\mathcal{L}_M \equiv \mathcal{L}_{\text{EP}}(\boldsymbol{\lambda}, \boldsymbol{\theta}) = \log \int p(\mathbf{f}; \boldsymbol{\theta}) \prod_{i=1}^n t_i(f_i; \boldsymbol{\lambda}_i) d\mathbf{f}, \quad (11)$$

giving the target for the M-step.






Figure 3. **Sparse approximation:** log marginal likelihood surfaces for the IONOSPHERE data set, changing the fraction  $m/n$  of the number of inducing points  $m$  vs.  $n = 351$  data points. The colour scale is the same in all plots:  $-0.8 \leq \text{value} \leq 0$ ; values below the predefined range are plotted as black. For moderate sparsification, our EP-like marginal likelihood estimation (top) matches the full MCMC baseline better than VI (bottom). For extreme sparsification (10%:  $m = 35$ ), neither approximation resembles the full surface.

Our hybrid training procedure uses the variational objective in the E-step to ensure a good representation of the posterior. Then in the M-step, we use an EP-like (and thus closer to marginal likelihood) objective for hyperparameter learning, at no additional computational cost. The algorithm is sketched out in Alg. 1. Although our procedure requires implementing two training objectives, this is not likelihood-specific and has minimal implementation overhead.

#### 4. Sparse Approximation for Large Data Sets

Regardless of conjugacy, inference in GP models for large-scale data sets is challenging due to an  $\mathcal{O}(n^3)$  computational bottleneck. In this section, we extend our hybrid training procedure to the sparse case.

A common approach to tackle this scalability issue is to summarize the information contained in the original data set into a smaller but more effective *pseudo*-data set, making the computational complexity tractable (see Fig. 3 and Quiñero-Candela & Rasmussen, 2005). The pseudo-inputs are referred to as inducing points and denoted as  $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^m$ , where  $m \ll n$  (Seeger, 2003; Csató, 2002; Quiñero-Candela & Rasmussen, 2005; Williams & Seeger, 2001). The pseudo-outputs are referred to as inducing variables and denoted as  $\mathbf{u} = f(\mathbf{Z})$ .

One common choice for the form of approximate posterior, as first introduced in Titsias (2009), is  $q(\mathbf{f}, \mathbf{u}; \xi_{\mathbf{u}}, \theta) = p(\mathbf{f} | \mathbf{u}; \theta) q(\mathbf{u}; \xi_{\mathbf{u}})$ , where  $p(\mathbf{f} | \mathbf{u})$  is the GP conditional and  $q(\mathbf{u}; \xi_{\mathbf{u}}) = \mathcal{N}(\mathbf{m}_{\mathbf{u}}, \mathbf{S}_{\mathbf{u}})$  the approximate posterior in  $\mathbf{u}$ . In this form,  $\mathbf{y}$  can only affect  $\mathbf{f}$  through  $\mathbf{u}$ , which means the information in the original data set is summarized in  $\mathbf{u}$ .

The marginal posterior over the function  $f(\cdot)$  is

$$q_{\mathbf{u}}(f(\cdot); \xi_{\mathbf{u}}, \theta) = \int p(f(\cdot) | \mathbf{u}; \theta) q(\mathbf{u}; \xi_{\mathbf{u}}) d\mathbf{u}, \quad (12)$$

where  $p(f(\cdot) | \mathbf{u}; \theta)$  is the distribution of the GP prior conditioned on  $f(\mathbf{Z}) = \mathbf{u}$ . Variational parameters  $\xi_{\mathbf{u}} = (\mathbf{m}_{\mathbf{u}}, \mathbf{S}_{\mathbf{u}})$  can be inferred by optimizing the following ELBO:

$$\begin{aligned} \log p(\mathbf{y}; \theta) &\geq \mathcal{L}_{\text{VI}}^{\text{sparse}}(\xi_{\mathbf{u}}, \theta) \\ &= \sum_{i=1}^n \mathbb{E}_{q_{\mathbf{u}}(f_i; \xi_{\mathbf{u}}, \theta)} [\log p(y_i | f_i; \theta)] \\ &\quad - \text{D}_{\text{KL}}[q(\mathbf{u}; \xi_{\mathbf{u}}) \| p(\mathbf{u}; \theta)], \end{aligned} \quad (13)$$

where  $q_{\mathbf{u}}(f_i; \xi_{\mathbf{u}}, \theta) = \mathcal{N}(f_i | \mathbf{a}_i^{\top} \mathbf{m}_{\mathbf{u}}, \kappa_{ii} - \mathbf{a}_i^{\top} (\mathbf{K}_{\mathbf{u}\mathbf{u}} - \mathbf{S}_{\mathbf{u}}) \mathbf{a}_i)$ ,  $\mathbf{a}_i = \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},i}$ ,  $\mathbf{K}_{\mathbf{u}\mathbf{u}} = \kappa(\mathbf{Z}, \mathbf{Z})$ , and  $\mathbf{k}_{\mathbf{u},i} = \kappa(\mathbf{Z}, \mathbf{x}_i)$  for the  $i^{\text{th}}$  data sample. Now under the dual parameterization, we denote the converged dual parameters of the posterior marginal  $q_{\mathbf{u}}^*(f_i)$  as  $\lambda_i^*$ . Adam et al. (2021) suggest designing a similar VEM procedure that exploits the structure of the  $q(\mathbf{u})$  in terms of the  $2n$  dual parameters. The natural parameters of  $q^*(\mathbf{u})$  are

$$(\mathbf{S}_{\mathbf{u}}^*)^{-1} \mathbf{m}_{\mathbf{u}}^* = \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \underbrace{\left( \sum_{i=1}^n \mathbf{k}_{\mathbf{u},i} \lambda_{1,i}^* \right)}_{=\lambda_1^*} \quad (14)$$

$$(\mathbf{S}_{\mathbf{u}}^*)^{-1} = \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} + \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \underbrace{\left( \sum_{i=1}^n \mathbf{k}_{\mathbf{u},i} \lambda_{2,i}^* \mathbf{k}_{\mathbf{u},i}^{\top} \right)}_{=\tilde{\Lambda}_2^*} \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}, \quad (15)$$

where the quantities  $\mathbf{K}_{\mathbf{u}\mathbf{u}}$  and  $\mathbf{k}_{\mathbf{u},i}$  directly depend on  $\theta$  and we can express the ELBO as the partition function of a Gaussian distribution. Adam et al. (2021) use a *tied* parameterization (motivated by site-tying in EP, see Bui

et al., 2017; Li et al., 2015) that relaxes the need of storing all the  $\{\lambda_i^*\}_{i=1}^n$  and instead stores only  $\bar{\lambda}_1^*$  (length  $m$ ) and  $\bar{\Lambda}_2^*$  (size  $m \times m$ ), which avoids the storage issue for large data sets. This extends the results of Khan & Lin to the sparse case where the resulting approximate posterior is

$$q(\mathbf{f}, \mathbf{u}; \lambda_{\mathbf{u}}, \theta) \propto p(\mathbf{f} | \mathbf{u}; \theta) p(\mathbf{u}; \theta) \times \prod_{i=1}^n \underbrace{\exp \langle \lambda_{\mathbf{u},i}, \mathbf{T}(\mathbf{a}_i^\top \mathbf{u}) \rangle}_{\triangleq t_i(\mathbf{u}; \lambda_{\mathbf{u},i})}, \quad (16)$$

where  $\lambda_{\mathbf{u},i} = \nabla_{\mu_{\mathbf{u},i}} \mathbb{E}_{q_{\mathbf{u}}(f_i; \lambda_{\mathbf{u},i}, \theta)} [\log p(y_i | f_i; \theta)]$ . This gives rise to the sparse E-step for inference under the sparse VEM scheme, where  $\mathcal{L}_{\text{EP}}^{\text{sparse}} \equiv \mathcal{L}_{\text{VI}}^{\text{sparse}}$ .

**Our proposed sparse objective for learning** In EP, the tied representation for constraining the problem to a summary  $\zeta_{\mathbf{u}}$  that scales in  $m$  rather than  $n$  gives rise to a sparse expectation propagation approach (Bui et al., 2017), where the log marginal likelihood is approximated as

$$\begin{aligned} \log p(\mathbf{y}; \theta) &\approx \mathcal{L}_{\text{EP}}^{\text{sparse}}(\zeta_{\mathbf{u}}, \theta) = \log \int q(\mathbf{f}, \mathbf{u}; \zeta_{\mathbf{u}}, \theta) d\mathbf{f} d\mathbf{u} \\ &= \log \int p(\mathbf{f} | \mathbf{u}; \theta) p(\mathbf{u}; \theta) \prod_{i=1}^n t_i(\mathbf{u}; \zeta_{\mathbf{u},i}) d\mathbf{f} d\mathbf{u}. \end{aligned} \quad (17)$$

Under dual parameterization VI, the approximate posterior Eq. (16) has the same structure as the EP approximate posterior in Eq. (17). An EP-like estimate of the log marginal likelihood can thus be calculated by injecting  $\lambda_{\mathbf{u},i}$  from Eq. (16) into  $\zeta_{\mathbf{u},i}$  in Eq. (17), thus giving  $\mathcal{L}_{\text{M}}^{\text{sparse}} \equiv \mathcal{L}_{\text{EP}}^{\text{sparse}}(\lambda_{\mathbf{u}}, \theta)$  which is a sparse EP-like learning objective under sparse variational inference. Note:  $\mathcal{L}_{\text{EP}}^{\text{sparse}}(\lambda_{\mathbf{u}}, \theta)$  has the same computational cost as VI.

## 5. Experiments

We provide a range of experiments, in which we demonstrate *effectiveness* and *practicality* of the proposed approach, and highlight similarities and differences between learning under the three most common approximative inference methods (LA, EP, and VI).

As the log marginal likelihood is a surrogate for the generalization ability of the model to unseen data, we evaluate the marginal likelihood estimations of different methods (Sec. 5.1) to see whether our EP-like marginal likelihood provides a better learning target. We then evaluate our hybrid training on non-conjugate tasks in binary classification and Student- $t$  regression on small and mid-sized data sets (Secs. 5.2 and 5.3).

How the sparsity affects the learning target is not obvious, therefore in Sec. 5.4 we first investigate the influence of sparse approximation on the learning target. We then evaluate our hybrid training procedure on binary classification

tasks with sparse approximation. For all experiments in the main paper, we use an isotropic Matérn-5/2 kernel. We also provide results under automatic relevance determination (ARD) with the same kernel (in App. D.4), where we only include results for data sets that could be confirmed to have converged.

We implement the variational methods in GPflow (Matthews et al., 2017), use reference implementations of LA and EP from GPy (GPy, since 2012), and base our MCMC implementation on the GPML toolbox (Rasmussen & Nickisch, 2010). Additionally, we use the GPstuff toolbox (Vanhatalo et al., 2013) for the custom LA and EP implementation for the Student- $t$  likelihood. We implement EP and VI convergence checks; details in App. D.4.

### 5.1. Quality of Marginal Likelihood Approximations

We compare the quality of marginal likelihood approximations of LA, VI, EP, and our EP-like VI with gold-standard MCMC. We demonstrate this on a binary classification task on the IONOSPHERE data set, with SONAR, USPS, PARKINSONS, and MONKS-2 in App. C.2. We estimate the log marginal likelihood on a  $21 \times 21$  grid of values for the log hyperparameters  $\log \theta = (\log \ell, \log \sigma)$  and plot the contour on the grid. For each hyperparameter setting, we fix the hyperparameters and evaluate the approximate log marginal likelihood based on the inferred approximate posterior.

**Markov Chain Monte Carlo baseline** MCMC is exact in the limit of long runs and thus provides a gold standard for log marginal likelihood estimation. Kuss & Rasmussen (2005) and Nickisch & Rasmussen (2008) proposed a sampling scheme based on annealed importance sampling (AIS, Neal, 2001) for obtaining a good estimate of the marginal likelihood (see App. C.1 for details). The baseline was computed by running  $21 \times 21 = 441$  jobs in parallel on a cluster.

**Experiment results** As shown in the top row of Fig. 2 on the IONOSPHERE benchmark data set, the marginal likelihood estimation of EP closely matches the MCMC baseline, whereas that of VI looks clearly different. Notably, when we estimate the marginal likelihood by  $\mathcal{L}_{\text{EP}}(\lambda, \theta)$  using the ‘site’ parameters of dual VI (Ours), the contour shapes become much closer to the MCMC result, demonstrating the improvement of using this EP-like marginal likelihood estimation. To investigate whether the improved marginal likelihood estimation also leads to better generalization, we select the optimal hyperparameter location across the grid values and compare the log predictive density on the test set (bottom row of Fig. 2). The optimal hyperparameter location of EP-like VI (Ours) is closer to MCMC than VI and generalizes well. We show the same analysis for different data sets covering different types of classification tasks (from general classification to small images) in App. C.2, with the same conclusion. In Figs. 6 to 9, VI and EP conform to their

**Table 1. Binary classification:** log predictive density (higher is better) on different data sets from the *Bayesian benchmarks* over 5-fold cross-validation with 10 different seeds. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**. We provide MCMC results for reference and exclude it from bolding. MCMC gives the best results on all data sets except Balloons. All inference methods perform well overall, while our training objective delivers the most reliable performance.

	( <i>n</i> , <i>d</i> )	LA	EP	VI	Ours	MCMC
TRAINS	(10, 30)	-0.702±0.025	<b>-0.698±0.033</b>	-0.702±0.037	<b>-0.691±0.046</b>	-0.692±0.025
BALLOONS	(16, 5)	-0.660±0.125	-0.650±0.128	-0.649±0.185	<b>-0.607±0.227</b>	-0.684±0.076
FERTILITY	(100, 10)	<b>-0.388±0.122</b>	<b>-0.384±0.149</b>	<b>-0.393±0.136</b>	<b>-0.397±0.139</b>	-0.382±0.126
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>-0.299±0.081</b>	-0.321±0.108	<b>-0.290±0.110</b>	<b>-0.293±0.116</b>	-0.306±0.115
ACUTE-NEPHRITIS	(120, 7)	-0.203±0.012	-0.046±0.007	-0.007±0.002	<b>-0.005±0.002</b>	-0.005±0.002
ACUTE-INFLAMMATION	(120, 7)	-0.184±0.018	-0.052±0.007	-0.007±0.002	<b>-0.007±0.002</b>	-0.007±0.003
ECHOCARDIOGRAM	(131, 11)	-0.424±0.093	<b>-0.418±0.095</b>	<b>-0.425±0.110</b>	-0.428±0.112	-0.437±0.127
HEPATITIS	(155, 20)	<b>-0.370±0.071</b>	-0.372±0.072	<b>-0.364±0.090</b>	-0.367±0.094	-0.369±0.091
PARKINSONS	(195, 23)	-0.260±0.031	-0.295±0.056	-0.160±0.050	<b>-0.141±0.046</b>	-0.145±0.044
BREAST-CANCER-WISC-PROG	(198, 34)	<b>-0.458±0.075</b>	-0.473±0.091	<b>-0.457±0.085</b>	<b>-0.460±0.088</b>	-0.464±0.085
SPECT	(265, 23)	<b>-0.593±0.049</b>	<b>-0.590±0.055</b>	-0.594±0.054	-0.595±0.054	-0.596±0.051
STATLOG-HEART	(270, 14)	-0.395±0.064	<b>-0.389±0.061</b>	-0.396±0.071	-0.397±0.071	-0.397±0.070
HABERMAN-SURVIVAL	(306, 4)	<b>-0.530±0.053</b>	-0.532±0.059	-0.531±0.055	-0.531±0.055	-0.520±0.063
IONOSPHERE	(351, 34)	-0.224±0.042	-0.230±0.042	<b>-0.170±0.048</b>	<b>-0.170±0.055</b>	-0.179±0.058
HORSE-COLIC	(368, 26)	-0.463±0.059	<b>-0.452±0.057</b>	-0.467±0.072	-0.473±0.082	-0.469±0.079
CONGRESSIONAL-VOTING	(435, 17)	-0.640±0.028	<b>-0.639±0.030</b>	-0.641±0.030	-0.642±0.029	-0.644±0.027
CYLINDER-BANDS	(512, 36)	-0.488±0.038	-0.500±0.041	-0.465±0.049	<b>-0.451±0.052</b>	-0.451±0.049
BREAST-CANCER-WISC-DIAG	(569, 31)	-0.085±0.026	-0.140±0.020	-0.077±0.044	<b>-0.075±0.045</b>	-0.076±0.043
ILPD-INDIAN-LIVER	(583, 10)	<b>-0.513±0.040</b>	-0.520±0.041	<b>-0.512±0.043</b>	<b>-0.512±0.043</b>	-0.512±0.042
MONKS-2	(601, 7)	-0.491±0.025	-0.512±0.028	-0.464±0.031	<b>-0.442±0.033</b>	-0.437±0.032
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>-0.630±0.026</b>	-0.639±0.036	<b>-0.630±0.026</b>	<b>-0.630±0.026</b>	-0.630±0.025
CREDIT-APPROVAL	(690, 16)	<b>-0.342±0.047</b>	<b>-0.342±0.050</b>	<b>-0.341±0.052</b>	-0.342±0.052	-0.341±0.052
BREAST-CANCER-WISC	(699, 10)	<b>-0.094±0.025</b>	<b>-0.093±0.023</b>	<b>-0.093±0.029</b>	-0.093±0.029	-0.093±0.029
BLOOD	(748, 5)	<b>-0.478±0.039</b>	<b>-0.479±0.040</b>	-0.478±0.039	-0.478±0.039	-0.478±0.039
PIMA	(768, 9)	<b>-0.474±0.033</b>	<b>-0.476±0.038</b>	<b>-0.474±0.035</b>	<b>-0.474±0.035</b>	-0.474±0.035
MAMMOGRAPHIC	(961, 6)	<b>-0.407±0.038</b>	<b>-0.407±0.040</b>	<b>-0.408±0.040</b>	-0.408±0.040	-0.408±0.040
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>-0.491±0.030</b>	<b>-0.491±0.032</b>	<b>-0.492±0.032</b>	<b>-0.492±0.032</b>	-0.492±0.032
Bold Count		14	13	13	16	/


stereotypes of being over- and under-confident, respectively, while Ours tends to have slightly better calibration.

For completeness, and motivated by the seminal work of Kuss & Rasmussen (2005), we provide back-to-back comparisons of both marginal log likelihood and predictive density surfaces also for LA. In Fig. 2, the marginal likelihood surface of LA resembles that of VI, while for the predictive density surface VI more closely resembles MCMC compared to LA.

## 5.2. Non-Conjugate Tasks in Bayesian Benchmarks

The log marginal likelihood is a surrogate for the generalization ability of the model to unseen data. To explicitly evaluate the generalization ability of our hybrid training procedure, we compare it against LA, EP, and VI on commonly used benchmark classification tasks. We use common small and mid-sized data sets ( $n \leq 1000$ ) and do full GP inference with 5-fold cross-validation. We use the *Bayesian Benchmarks* suite ([github.com/secondmind-labs/bayesian\\_benchmarks](https://github.com/secondmind-labs/bayesian_benchmarks)) for evaluating the methods.

**Evaluation on binary classification** We consider binary classification with a Bernoulli likelihood on 27 data sets from the UCI repository (Dua & Graff, 2017). For all approximate inference methods, we set the same number of maximum training iterations and use the relative changes in the parameters of the model as convergence criteria. For our



**Figure 4.** Mean relative accuracy compared to best method on each data set of Table 1 (over 5-fold CV repeated with 10 seeds). The horizontal lines indicate mean across all data sets; see legend for mean and standard deviation. Our approach yields reliable training, with the highest average relative accuracy and the least outliers.

hybrid training procedure due to the conflicting objectives in E- and M-steps discussed in Sec. 5.5, we use a decrease in the EP learning objective as an additional convergence criterion. As a gold-standard baseline, we include MCMC results. For MCMC, we use a log-uniform hyperpriors to ensure a close match to the model setup in the other models.

To reduce the variance introduced by the training–test set split, we repeat the 5-fold CV with ten different seeds. The



**Table 2. Robust regression tasks:** log predictive density (higher is better) with a Student- $t$  likelihood on different data sets (mean  $\pm$  standard deviation over 5-fold cross-validation). The best results and those not statistically significantly different from them under a paired  $t$ -test are **bolded**. Our objective performs well overall.

	$(n, d)$	LA	EP	VI	Ours
NEAL	(100, 1)	<b>.317</b> $\pm$ .440	<b>.303</b> $\pm$ .432	<b>.295</b> $\pm$ .426	<b>.301</b> $\pm$ .436
BOSTON	(506, 13)	<b>-.210</b> $\pm$ .069	<b>-.190</b> $\pm$ .053	-.206 $\pm$ .056	<b>-.195</b> $\pm$ .061
STOCK	(1000, 1)	1.910 $\pm$ .079	1.389 $\pm$ .242	1.917 $\pm$ .082	<b>1.921</b> $\pm$ .083
Bold Count		2	2	1	3

performance on the test set is given in Table 1 and Fig. 4. As shown in Table 1 for log predictive density, LA and EP have very similar performance to VI-based methods (VI and Ours) on most data sets. This empirically demonstrates that for binary classification on small and mid-sized data sets EP and LA generalize well. Our hybrid training procedure achieves the same test performance as VI on most data sets and outperforms VI on three data sets. It empirically demonstrates that when no sparse approximation is required, by using an improved estimation of the marginal likelihood for hyperparameter learning, we could potentially have better generalization ability at no additional computational cost. As the gold standard, MCMC gives the best results; notably, the gap between MCMC and approximate inference methods is relatively small on small data sets. In practice we often favour methods with stable performance over different data sets, *i.e.*, they might not always give the best performance but we can expect consistently good performances. To investigate the reliability of different methods, in Fig. 4 we plot the relative accuracy (on each data set we divide the results of each method by the highest accuracy on that data set) for individual data sets and the mean relative accuracy of each method across all data sets. Our approach achieves the most consistent performance on all data sets and thus yields reliable training. We include additional result tables with the same conclusion in App. D.3, including experiments with an ARD kernel and checks for initializing other methods with our optimal hyperparameters.

### 5.3. Robust (Student- $t$ ) Regression

We further test our hybrid training procedure on a more challenging robust regression task with a Student- $t$  likelihood, a model which is not log-concave. In the likelihood, we fix the degrees of freedom,  $\nu = 3$ , and only train the noise scale together with hyperparameters. For LA and EP we follow the methods designed by Jylänki et al. (2011). For VI and our EP-like VI, to make the training procedure numerically stable we crop the gradient w.r.t. the second element of the natural parameters to prevent the approximate posterior covariance from becoming negative. We test on three benchmark data sets previously used for robust regression: the simulated data from Neal (2022), the

**Table 3. Sparse approximation** for classification on different data sets: log predictive density (mean  $\pm$  standard deviation). Higher is better. Results that are statistically significantly different under a paired  $t$ -test are **bolded**.

	LA	EP	VI	Ours
TITANIC	-.217 $\pm$ .037	-.014 $\pm$ .004	<b>-.011</b> $\pm$ .003	-.037 $\pm$ .005
BANK	<b>-.247</b> $\pm$ .006	<b>-.246</b> $\pm$ .007	-.249 $\pm$ .006	-.247 $\pm$ .007
TWONORM	<b>-.060</b> $\pm$ .007	<b>-.061</b> $\pm$ .008	<b>-.060</b> $\pm$ .008	-.524 $\pm$ .208
MUSHROOM	-.129 $\pm$ .003	-.002 $\pm$ .000	<b>-.001</b> $\pm$ .000	-.028 $\pm$ .001
MAGIC	<b>-.285</b> $\pm$ .333	-.693 $\pm$ .000	<b>-.008</b> $\pm$ .001	-.070 $\pm$ .002
Bold Count	3	2	4	0

Boston housing regression task, and the stock data from Solin & Särkkä (2015). Jylänki et al. (2011) point out that in Student- $t$  regression EP provides a good approximation for marginal likelihood and, as shown in Table 2, by using an EP-like marginal likelihood estimation for hyperparameters learning our hybrid training procedure generalizes better than vanilla VI. The MCMC gold standard results for NEAL, BOSTON and STOCK are  $0.309 \pm 0.454$ ,  $-0.191 \pm 0.051$ , and  $1.586 \pm 0.034$ , respectively.

### 5.4. Evaluation under Sparse Approximation

It is not obvious how a sparse approximation would affect the quality of marginal likelihood approximation. To be able to compare with the MCMC baseline on the full data set, we first analyse the influence of sparsification on IONOSPHERE. We choose 75%, 50%, 25%, and 10% random subsets of training data as inducing points. We estimate the log marginal likelihood as in Sec. 5.1 on a grid of values for the log hyperparameters. Fig. 3 shows the resulting contour surfaces. Unsurprisingly, as we reduce the number of inducing points, the estimation of the log marginal likelihood becomes less accurate. For moderate sparsification (75%, 50%), our EP-like marginal likelihood estimation matches the full MCMC baseline better than VI. For more extreme sparsification, both approximations show significant biases. This is because, with very few inducing points, only larger lengthscales make sense. To match the low-complexity approximation, large lengthscales are required and the ground truth marginal likelihood provided by MCMC becomes irrelevant.

**Evaluation on large-scale binary classification** We compare LA, EP, VI, and our proposed training procedure on five data sets from 2k to 19k data points (for details, see Table 11). We use k-means to select 500 inducing points from the input data and keep them fixed. The test set performance is given in Table 3. Here we are in the regime of  $m/n$  in the range of 25% to 2.5%, where the log marginal likelihood surface approximations of both VI and our EP-like approximation are likely to be biased away from the true marginal likelihood surface, and the approximate sparse model can no longer be considered a surrogate of the true model.

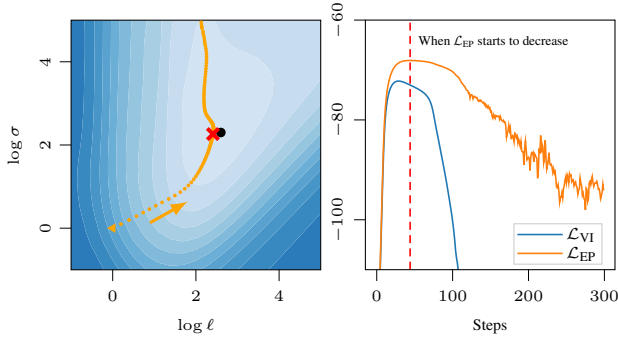


Figure 5. Interplay between the two optimization targets can result in overshooting the optimum (●). The optimizer starts at ▲ and once it passes × (dashed red; stopping point),  $\mathcal{L}_{EP}$  starts to decrease. After that, the optimization becomes increasingly unstable.

### 5.5. Practical Considerations

VEM with the same objective for E- and M-step is analogous to coordinate ascent and hence guaranteed to always improve the objective. With  $\mathcal{L}_{VI}$  for E-step and  $\mathcal{L}_{EP}$  for M-step, this guarantee no longer holds. This can introduce interplay between the two targets, and in our experiments, we observed that (with certain data-splits/model setups) the optimization can overshoot past the optimum and then becomes increasingly unstable, see Fig. 5. In our experiments this conflict between  $\mathcal{L}_E = \mathcal{L}_{VI}$  and  $\mathcal{L}_M = \mathcal{L}_{EP}$  occurred in about 54% of cases. We address this issue by ending optimization once the  $\mathcal{L}_{EP}$  objective starts to decrease (see Alg. 1). Note that this is solely based on the training data and does not require any additional validation data or tuning per data set.

## 6. Discussion and Conclusions

In GP models, the training separates into inference and learning which are typically both cast into an optimization problem. In this paper, we improved hyperparameter learning in non-conjugate GP models by augmenting VI with an EP-like learning target. Our hybrid training procedure builds upon the dual variational GP formulation, introduced by Khan & Lin (2017) and extended to sparse GPs in Adam et al. (2021), which provides a link between VI and EP. We used the representation of the posterior from VI to obtain an EP-like approximation of the marginal likelihood for hyperparameter optimization—without any added computational complexity or computation time.

In the experiments, we evaluated our hybrid training procedure on binary classification tasks and robust regression. For full (non-sparse) models, the extensive results (over 1350 runs per method) show clear benefits in stability, reliability, and performance for our method. This shows the benefits of decoupling inference and learning. When more hyper-

parameters are present, as shown in Table 5 and Table 6 in the appendix, our method has the same performance on log predictive density as VI. However, as shown in Fig. 10, our method is still more reliable than VI. For sparse problems, similar empirical benefits could not be demonstrated.

We provide a reference implementation of the methods and code to reproduce the experiments at <https://github.com/AaltoML/improved-hyperparameter-learning>.

## Acknowledgements

This work was supported by funding from the Academy of Finland (grant id 339730) and the Finnish Center for Artificial Intelligence (FCAI). We acknowledge the computational resources provided by the Aalto Science-IT project and CSC – IT Center for Science, Finland. We thank Aidan Scannell for his constructive comments on the manuscript.

## References

- Adam, V., Chang, P. E., Khan, M. E., and Solin, A. Dual parameterization of sparse variational Gaussian processes. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, volume 34, pp. 11474–11486. Curran Associates, Inc., 2021.
- Amari, S.-I. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- Bui, T. D., Yan, J., and Turner, R. E. A unifying framework for Gaussian process pseudo-point approximations using power expectation propagation. *Journal of Machine Learning Research (JMLR)*, 18:3649–3720, 2017.
- Chang, P. E., Wilkinson, W. J., Khan, M. E., and Solin, A. Fast variational learning in state-space Gaussian process models. In *International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2020.
- Csató, L. *Gaussian Processes: Iterative Sparse Approximations*. PhD thesis, Aston University, Birmingham, UK, 2002.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. <http://archive.ics.uci.edu/ml>.
- Gorinova, M., Moore, D., and Hoffman, M. Automatic reparameterisation of probabilistic programs. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3648–3657. PMLR, 2020.
- GPy. GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>, since 2012.

- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. In *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 282–290. AUAI Press, 2013.
- Jylänki, P., Vanhatalo, J., and Vehtari, A. Robust Gaussian process regression with a Student- $t$  likelihood. *Journal of Machine Learning Research (JMLR)*, 12:3227–3257, 2011.
- Khan, M. E. and Lin, W. Conjugate-computation variational inference: Converting variational inference in non-conjugate models to inferences in conjugate models. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54 of *Proceedings of Machine Learning Research*, pp. 878–887. PMRL, 2017.
- Khan, M. E., Aravkin, A., Friedlander, M., and Seeger, M. Fast dual variational inference for non-conjugate latent Gaussian models. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, volume 28 of *Proceedings of Machine Learning Research*, pp. 951–959. PMLR, 2013.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Kuss, M. and Rasmussen, C. E. Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Research (JMLR)*, 6:1679–1704, 2005.
- Li, Y., Hernández-Lobato, J. M., and Turner, R. E. Stochastic expectation propagation. In *Advances in Neural Information Processing Systems 28*, pp. 2323–2331. Curran Associates, Inc., 2015.
- Lotfi, S., Izmailov, P., Benton, G., Goldblum, M., and Wilson, A. G. Bayesian model selection, the marginal likelihood, and generalization. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14223–14247. PMLR, 2022.
- Matthews, A. G. d. G., van der Wilk, M., Nickson, T., Fujii, K., Boukouvalas, A., León-Villagrà, P., Ghahramani, Z., and Hensman, J. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research (JMLR)*, 18(40):1–6, 2017.
- Minka, T. P. Expectation propagation for approximate Bayesian inference. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI)*, *Proceedings of Machine Learning Research*, pp. 362–369, 2001.
- Murray, I., Adams, R., and MacKay, D. Elliptical slice sampling. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 541–548. JMLR Workshop and Conference Proceedings, 2010.
- Neal, R. M. Annealed importance sampling. *Statistics and Computing*, 11:125–139, 2001.
- Neal, R. M. Software for flexible Bayesian modeling and Markov Chain sampling, 2022. <http://www.cs.toronto.edu/~radford/fbm.software.html>.
- Nickisch, H. and Rasmussen, C. E. Approximations for binary Gaussian process classification. *Journal of Machine Learning Research (JMLR)*, 9:2035–2078, 2008.
- Opper, M. and Archambeau, C. The variational Gaussian approximation revisited. *Neural Computation*, 21(3):786–92, 2009.
- Quiñonero-Candela, J. and Rasmussen, C. E. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research (JMLR)*, 6:1939–1959, 2005.
- Rasmussen, C. E. and Nickisch, H. Gaussian processes for machine learning (GPML) toolbox. *Journal of Machine Learning Research (JMLR)*, 11:3011–3015, 2010.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Salimbeni, H., Eleftheriadis, S., and Hensman, J. Natural gradients in practice: Non-conjugate variational inference in Gaussian process models. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 84 of *Proceedings of Machine Learning Research*, pp. 689–697. PMLR, 2018.
- Seeger, M. *Bayesian Gaussian Process Models: PAC-Bayesian Generalisation Error Bounds and Sparse Approximations*. PhD thesis, University of Edinburgh, Edinburgh, UK, 2003.
- Solin, A. and Särkkä, S. State space methods for efficient inference in Student- $t$  process regression. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 38 of *Proceedings of Machine Learning Research*, pp. 885–893. PMLR, 2015.
- Titsias, M. Variational learning of inducing variables in sparse Gaussian processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pp. 567–574, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR.

- Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., and Vehtari, A. GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research (JMLR)*, 14(1):1175–1179, 2013.
- Vapnick, V. N. *Statistical Learning Theory*. Wiley, New York, 1998.
- Vehtari, A., Mononen, T., Tolvanen, V., Sivula, T., and Winther, O. Bayesian leave-one-out cross-validation approximations for Gaussian latent variable models. *Journal of Machine Learning Research (JMLR)*, 17(1):3581–3618, 2016.
- Vehtari, A., Gelman, A., Sivula, T., Jylänki, P., Tran, D., Sahai, S., Blomstedt, P., Cunningham, J. P., Schiminovich, D., and Robert, C. P. Expectation propagation as a way of life. *Journal of Machine Learning Research (JMLR)*, 21:1–53, 2020.
- Wilkinson, W. J., Särkkä, S., and Solin, A. Bayes–Newton methods for approximate Bayesian inference with PSD guarantees. *Journal of Machine Learning Research (JMLR)*, 24(83):1–50, 2023.
- Williams, C. K. and Barber, D. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, 1998.
- Williams, C. K. and Seeger, M. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems 13 (NIPS)*, pp. 682–688. MIT Press, 2001.



## Appendix

In the supplementary material, we include technical details of the methods that were omitted for brevity in the main paper (App. A). Additionally, we provide details on the experiments and evaluation setup (App. B–F) for reproducing the results in the main paper, and include further result tables and figures that extend the evaluation. The codes for the methods proposed in this paper are included as a separate supplement.

### A. Method Details

**Sparse energy** To extend the presentation in Sec. 4, we derive how to obtain the sparse EP marginal likelihood estimation from the VI approximate posterior. Following Eq. (77) in Bui et al. (2017), the sites  $t_i(\mathbf{u}; \zeta_{\mathbf{u},i})$  in Eq. (17) are

$$t_i(\mathbf{u}; \zeta_{\mathbf{u},i}) \propto \exp\langle \mathbf{u}^\top \mathbf{a}_i \zeta_{\mathbf{u},i,2} \zeta_{\mathbf{u},i,1} - \frac{1}{2} \mathbf{u}^\top \mathbf{a}_i \zeta_{\mathbf{u},i,2} \mathbf{a}_i^\top \mathbf{u} \rangle, \quad (18)$$

where  $\mathbf{a}_i = \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{k}_{\mathbf{u},i}$ ,  $\mathbf{K}_{\mathbf{uu}} = \kappa(\mathbf{Z}, \mathbf{Z})$ , and  $\mathbf{k}_{\mathbf{u},i} = \kappa(\mathbf{Z}, \mathbf{x}_i)$  for the  $i^{\text{th}}$  data sample. Note that  $\mathbf{u}^\top \mathbf{a}_i = \mathbf{a}_i^\top \mathbf{u}$  is a scalar.

As  $t_i(\mathbf{u}; \lambda_{\mathbf{u},i})$  in Eq. (16) is given by

$$\begin{aligned} t_i(\mathbf{u}; \lambda_{\mathbf{u},i}) &= \exp\langle \lambda_{\mathbf{u},i}, \mathbf{T}(\mathbf{a}_i^\top \mathbf{u}) \rangle \\ &= \exp\langle \lambda_{\mathbf{u},i,1} \mathbf{a}_i^\top \mathbf{u} + \lambda_{\mathbf{u},i,2} (\mathbf{a}_i^\top \mathbf{u})^2 \rangle, \end{aligned} \quad (19)$$

we have the following correspondence between  $\zeta_{\mathbf{u}}$  and  $\lambda_{\mathbf{u}}$ :

$$\zeta_{\mathbf{u},i,2} \zeta_{\mathbf{u},i,1} \Leftrightarrow \lambda_{\mathbf{u},i,1} \quad \text{and} \quad -\frac{1}{2} \zeta_{\mathbf{u},i,2} \Leftrightarrow \lambda_{\mathbf{u},i,2}. \quad (20)$$

Following Eq. (126) in Bui et al. (2017), the sparse EP energy is (we omit  $\theta$  to make notation simpler)

$$\begin{aligned} \mathcal{L}_{\text{EP}}^{\text{sparse}}(\zeta_{\mathbf{u}}, \theta) &= \frac{1}{2} \log |\mathbf{S}_{\mathbf{u}}| + \frac{1}{2} \mathbf{m}_{\mathbf{u}}^\top \mathbf{S}_{\mathbf{u}}^{-1} \mathbf{m}_{\mathbf{u}} - \frac{1}{2} \log |\mathbf{K}_{\mathbf{uu}}| + \frac{1}{\alpha} \sum_n \log \mathcal{Z}_{\text{tilted},i} \\ &+ \sum_n \left[ -\frac{1}{2\alpha} \log(1 - \mathbf{a}_i^\top \alpha \zeta_{\mathbf{u},i,2} \mathbf{S}_{\mathbf{u}} \mathbf{a}_i) + \frac{1}{2} \frac{\mathbf{m}_{\mathbf{u}}^\top \mathbf{a}_i \zeta_{\mathbf{u},i,2} \mathbf{a}_i^\top \mathbf{m}_{\mathbf{u}}}{1 - \mathbf{a}_i^\top \alpha \zeta_{\mathbf{u},i,2} \mathbf{S}_{\mathbf{u}} \mathbf{a}_i} \right. \\ &\left. + \frac{1}{2} \zeta_{\mathbf{u},i,1} \zeta_{\mathbf{u},i,2} \mathbf{a}_i^\top \mathbf{V}_{\text{cav},i} \mathbf{a}_i \alpha \zeta_{\mathbf{u},i,2} \zeta_{\mathbf{u},i,1} - \zeta_{\mathbf{u},i,1} \zeta_{\mathbf{u},i,2} \mathbf{a}_i^\top \mathbf{V}_{\text{cav},i} \mathbf{S}_{\mathbf{u}}^{-1} \mathbf{m}_{\mathbf{u}} \right], \end{aligned} \quad (21)$$

where the different terms are defined by

$$\mathbf{S}_{\text{cav},i} = \mathbf{S}_{\mathbf{u}} + \frac{\mathbf{S}_{\mathbf{u}} \mathbf{a}_i \alpha \zeta_{\mathbf{u},i,2} \mathbf{a}_i^\top \mathbf{S}_{\mathbf{u}}}{1 - \mathbf{a}_i^\top \alpha \zeta_{\mathbf{u},i,2} \mathbf{S}_{\mathbf{u}} \mathbf{a}_i}, \quad (22)$$

$$\mathbf{S}_{\text{cav},i}^{-1} \mathbf{m}_{\text{cav},i} = \mathbf{S}_{\mathbf{u}}^{-1} \mathbf{m}_{\mathbf{u}} - \alpha \mathbf{a}_i \zeta_{\mathbf{u},i,2} \zeta_{\mathbf{u},i,1}, \quad (23)$$

$$\log \mathcal{Z}_{\text{tilted},i} = \log \int q_{\text{cav},i}(f_i) p^\alpha(y_i | f_i) \mathrm{d}f_i, \quad (24)$$

$$q_{\text{cav},i}(f_i) = \int p(f_i | \mathbf{u}) \mathcal{N}(\mathbf{m}_{\text{cav},i}, \mathbf{S}_{\text{cav},i}) \mathrm{d}\mathbf{u}. \quad (25)$$

By substituting  $\lambda_{\mathbf{u},i}$  into  $\zeta_{\mathbf{u},i}$  in Eq. (21) using Eq. (20), we obtain the sparse EP marginal likelihood approximation with the VI approximate posterior. When  $\alpha = 1$ , Power-EP reduces to normal EP, which we use in our experiments.

### B. Computational Details

All experiments ran on a cluster, which allowed us to parallelize jobs. This played a central role especially for the MCMC baseline results for the marginal likelihood surfaces, where we split into 441 separate jobs (per hyperparameter value combination), each of which were allocated 1–3 CPU cores and 1 Gb memory and ran 8–40 h depending on data set size.

## C. Quality of Marginal Likelihood Approximation

Inspired by the work by [Kuss & Rasmussen \(2005\)](#) and [Nickisch & Rasmussen \(2008\)](#), we compare the quality of marginal likelihood approximations of LA, VI, EP, and our EP-like VI with a ‘ground’ truth obtained by annealed importance sampling (AIS, [Neal, 2001](#)). We demonstrate this on binary classification tasks, where we estimate the log marginal likelihood on a  $21 \times 21$  grid of values for the log hyperparameters  $\log \theta = (\log \ell, \log \sigma)$  and plot the contour on the grid. For each hyperparameter setting, we fix the hyperparameters and evaluate the approximate log marginal likelihood based on the inferred approximate posterior. We then also visualize the log predictive density on hold-out test data as a similar contour plot, showing what the performance of the model would have been if the hyperparameters would have been chosen based on the log marginal likelihood surface in question under the specific inference scheme.

### C.1. Markov Chain Monte Carlo Baseline

As in previous work, we use an MCMC approach as the gold-standard baseline. We use the annealed importance sampling ([Neal, 2001](#)) approach from [Kuss & Rasmussen \(2005\)](#) and [Nickisch & Rasmussen \(2008\)](#) that defines a sequence of  $t = 0, 1, \dots, T$  steps  $Z_t = \int p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t)} p(\mathbf{f}; \theta) d\mathbf{f}$ , where  $\tau(t) = (t/T)^4$  (such that  $\tau(0) = 0$  and  $\tau(T) = 1$ ). The marginal likelihood can be rewritten as

$$p(\mathbf{y}; \theta) = \frac{Z_T}{Z_0} = \frac{Z_T}{Z_{T-1}} \frac{Z_{T-1}}{Z_{T-2}} \dots \frac{Z_1}{Z_0}, \quad (26)$$

where  $Z_t/Z_{t-1}$  is approximated by importance sampling using samples from  $q_t(\mathbf{f}) \propto p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t-1)} p(\mathbf{f}; \theta)$ :

$$\begin{aligned} \frac{Z_t}{Z_{t-1}} &= \frac{\int p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t)} p(\mathbf{f}; \theta) d\mathbf{f}}{Z_{t-1}} \\ &= \int \frac{p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t)}}{p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t-1)}} \frac{p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t-1)} p(\mathbf{f}; \theta)}{Z_{t-1}} d\mathbf{f} \\ &\approx \frac{1}{S} \sum_{s=1}^S p(\mathbf{y} | \mathbf{f}_t^{(s)}; \theta)^{\tau(t) - \tau(t-1)}, \quad \text{where } \mathbf{f}_t^{(s)} \sim \frac{p(\mathbf{y} | \mathbf{f}; \theta)^{\tau(t-1)} p(\mathbf{f}; \theta)}{Z_{t-1}}. \end{aligned} \quad (27)$$

In practice, instead of sampling  $\mathbf{f}_t$  from  $\frac{p(\mathbf{y} | \mathbf{f})^{\tau(t-1)} N(\mathbf{f} | \mathbf{m}, \mathbf{K})}{Z_{t-1}}$  directly, we use a parameterisation in terms of  $\alpha = \mathbf{K}^{-1}(\mathbf{f}_t - \mathbf{m})$  and sample from  $\alpha \sim P(\alpha) = \frac{p(\mathbf{y} | \mathbf{K}\alpha + \mathbf{m})^{\tau(t-1)} N(\alpha | \mathbf{0}, \mathbf{K}^{-1})}{Z_{t-1}}$  to increase numerical stability since  $\log P(\alpha)$  and its gradient can be computed safely. We use elliptical slice sampling ([Murray et al., 2010](#)). Now, by using a single sample  $S = 1$  and a large number of steps  $T$ , the estimation of log marginal likelihood can be written as

$$\log p(\mathbf{y}; \theta) = \sum_{t=1}^T \log \frac{Z_t}{Z_{t-1}} \approx \sum_{t=1}^T (\tau(t) - \tau(t-1)) \log p(\mathbf{y} | \mathbf{f}_t; \theta). \quad (28)$$

Following [Kuss & Rasmussen \(2005\)](#), we set  $T = 8000$  and combine three estimates of log marginal likelihood by their geometric mean.

### C.2. Experiment Results

In addition to the figure for the IONOSPHERE data set ([Fig. 2](#)) in the main paper, we include surface plots for four additional data sets in the appendix for more comprehensive comparisons. The marginal likelihood estimation on SONAR, USPS, PARKINSONS and MONKS-2 are given in [Fig. 6](#), [Fig. 7](#), [Fig. 8](#), and [Fig. 9](#), respectively. The IONOSPHERE and SONAR were included to make it easy to compare to previous work, and the other three chosen as an interesting subset covering different type of classification tasks (from general classification to small images).

Similarly to the results on IONOSPHERE in the main paper, when using the EP-like marginal likelihood estimation from the VI approximate posterior (Ours), the contour shapes become closer to the MCMC result, and the optimal hyperparameter location of EP-like VI (Ours) is closer to MCMC than VI. These effects help explain the quantitative results in the tables in the main paper and the supplement.

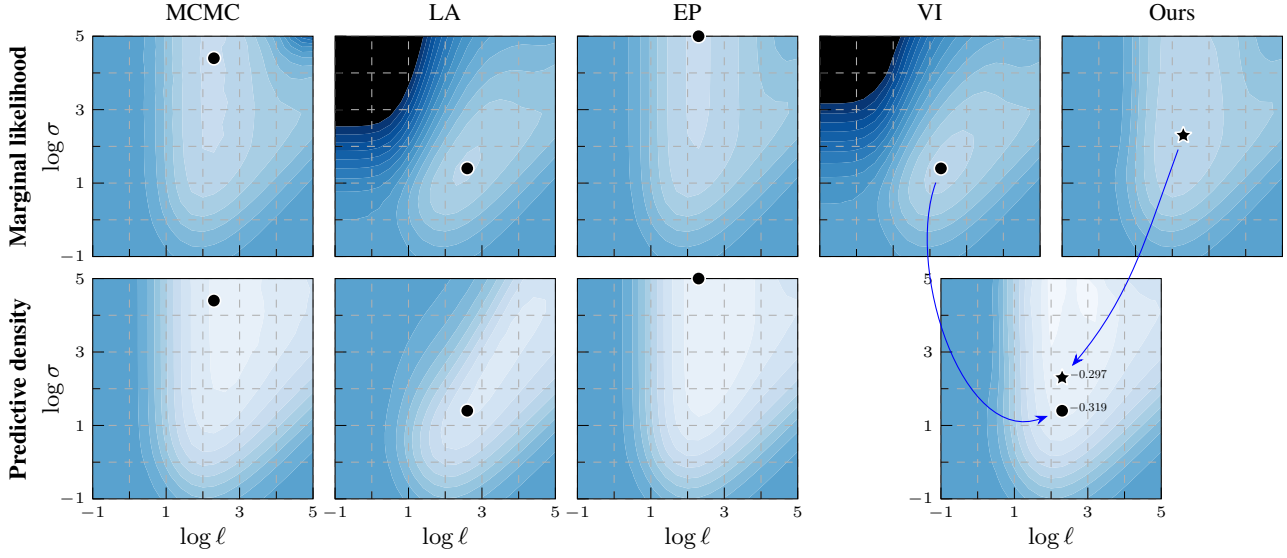
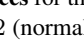


Figure 6. **Log marginal likelihood / predictive density surfaces** for the SONAR data set by varying kernel magnitude  $\sigma$  and lengthscale  $\ell$ . The colour scale is the same in all plots:  $-1.0$    $-0.2$  (normalized by  $n$ ). Optimal hyperparameters of each method shown by a black marker. EP and our EP-like marginal likelihood estimation match the MCMC baseline better than VI or LA, thus providing a learning proxy. For prediction, our method still leverages the same variational representation as VI.





Figure 7. **Log marginal likelihood / predictive density surfaces** for the USPS data set (MNIST-like digits image) by varying kernel magnitude  $\sigma$  and lengthscale  $\ell$ . The colour scale is the same in all plots:  $-1.0$    $0$  (normalized by  $n$ ). Optimal hyperparameters shown by a black marker. EP and our EP-like marginal likelihood estimation match the MCMC baseline better than VI or LA, thus providing a learning proxy. For prediction, our method still leverages the same variational representation as VI.





Figure 8. **Log marginal likelihood / predictive density surfaces** for the PARKINSONS data set by varying kernel magnitude  $\sigma$  and lengthscale  $\ell$ . The colour scale is the same in all plots:  $-0.8$    $0$  (normalized by  $n$ ). Optimal hyperparameters shown by a black marker. EP and our EP-like marginal likelihood estimation match the MCMC baseline better than VI or LA, thus providing a learning proxy. For prediction, our method still leverages the same variational representation as VI.





Figure 9. **Log marginal likelihood / predictive density surfaces** for the MONKS-2 data set by varying kernel magnitude  $\sigma$  and lengthscale  $\ell$ . The colour scale is the same in all plots:  $-0.9$    $-0.3$  (normalized by  $n$ ). Optimal hyperparameters shown by a black marker. EP and our EP-like marginal likelihood estimation match the MCMC baseline better than VI or LA, thus providing a learning proxy. For prediction, our method still leverages the same variational representation as VI.



## D. Non-Conjugate Tasks in Bayesian Benchmarks

We use the *Bayesian Benchmarks* suite ([github.com/secondmind-labs/bayesian\\_benchmarks](https://github.com/secondmind-labs/bayesian_benchmarks); originally by Salimbeni *et al.*) for evaluating the methods in binary classification. Bayesian benchmarks includes common evaluation data sets (typically from UCI) and makes it possible to run a large number of comparisons under a fixed evaluation setup. In the first part, we only include binary classification tasks (Bernoulli likelihood) with  $n \leq 1000$ . We follow the standard setup of input point normalization and splits in the evaluation suite.

### D.1. Evaluation Metrics

We conduct 5-fold cross-validation and use test set accuracy and log predictive density to evaluate the test performance of each method (higher is better in both). To compare different methods, we use the paired  $t$ -test (with  $p = 0.05$ ) that compares whether the best-performing method performs statistically significantly better than the others.

### D.2. Experiment Setup

We initialize the hyperparameters with unit lengthscale and magnitude for all methods. For LA and EP, the hyperparameters are optimized by the default optimizer L-BFGS-B in GPy. For VI and our hybrid training procedure, each E step and M step consists of 20 iterations. In the E-step we set the learning rate of natural gradient descent to be 0.1. In the M-step we use the Adam optimizer (Kingma & Ba, 2015) with learning rate 0.01. We use the convergence criterion described in the main text, with a maximum number of at most 10 000 steps.

For MCMC, we use log uniform prior to ensure it is the same model as the approximate inference methods. We set the burn-in step to be 200, the number of samples to be 10000, and thin the sample with 2.

### D.3. Experiment Results

The test set accuracies are given in Table 4. For test accuracies all methods achieve similar performance, which is to be expected as accuracy alone only captures where the decision boundary has been drawn, completely disregarding second-order information. The log predictive density results are included in the main paper (Table 1). Our hybrid training procedure gives the most consistent results and thus achieves the most reliable training.

### D.4. Ablation Studies

**Automatic Relevance Determination Kernel** We run experiments with automatic relevance determination (ARD) kernel to see whether our method would perform well when there are multiple hyperparameters. To ensure fair comparison we only include results on data sets where all methods have converged. The log predictive density and test set accuracy are given in Table 5 and Table 6 respectively. The mean relative accuracy is plotted in Fig. 10. Our training objective performs well overall and achieves reliable training.

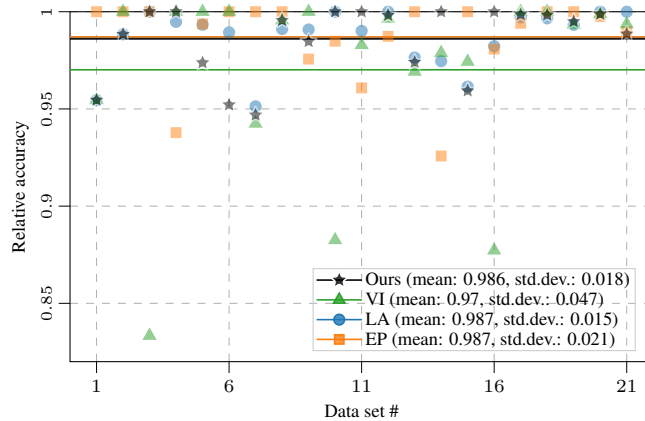


Figure 10. Mean relative accuracy compared to best method on each data set of Table 6. The horizontal lines indicate mean across all data sets; see legend for mean and standard deviation. Our approach yields reliable training.

**Table 4. Binary classification:** test set accuracy (higher is better) on different data sets from the *Bayesian benchmarks* over 5-fold cross-validation with 10 different seeds. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**. We provide MCMC results for reference (excluded from bolding). All inference methods perform well overall, while our training objective delivers the most reliable performance.

	$(n, d)$	LA	EP	VI	Ours	MCMC
TRAINS	(10, 30)	0.620±0.394	<b>0.680±0.371</b>	0.660±0.353	<b>0.710±0.348</b>	0.730±0.349
BALLOONS	(16, 5)	<b>0.618±0.275</b>	0.607±0.282	<b>0.615±0.285</b>	<b>0.650±0.291</b>	0.625±0.272
FERTILITY	(100, 10)	<b>0.879±0.062</b>	<b>0.880±0.062</b>	<b>0.879±0.062</b>	<b>0.877±0.063</b>	0.877±0.061
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>0.874±0.070</b>	<b>0.868±0.075</b>	<b>0.875±0.072</b>	<b>0.877±0.068</b>	0.868±0.069
ACUTE-NEPHRITIS	(120, 7)	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	1.000±0.000
ACUTE-INFLAMMATION	(120, 7)	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	1.000±0.000
ECHOCARDIOGRAM	(131, 11)	0.820±0.064	<b>0.840±0.056</b>	0.808±0.069	0.808±0.068	0.797±0.073
HEPATITIS	(155, 20)	<b>0.830±0.059</b>	0.819±0.061	<b>0.833±0.059</b>	<b>0.832±0.060</b>	0.830±0.057
PARKINSONS	(195, 23)	<b>0.951±0.031</b>	0.888±0.050	0.942±0.032	<b>0.949±0.034</b>	0.950±0.031
BREAST-CANCER-WISC-PROG	(198, 34)	0.808±0.058	0.793±0.070	<b>0.815±0.057</b>	<b>0.814±0.057</b>	0.808±0.060
SPECT	(265, 23)	<b>0.706±0.055</b>	<b>0.703±0.057</b>	<b>0.703±0.056</b>	<b>0.705±0.056</b>	0.707±0.058
STATLOG-HEART	(270, 14)	0.830±0.049	<b>0.838±0.047</b>	0.830±0.049	0.828±0.050	0.822±0.050
HABERMAN-SURVIVAL	(306, 4)	<b>0.741±0.048</b>	<b>0.740±0.050</b>	<b>0.741±0.049</b>	<b>0.741±0.048</b>	0.746±0.047
IONOSPHERE	(351, 34)	0.927±0.029	0.930±0.035	<b>0.942±0.030</b>	<b>0.945±0.027</b>	0.943±0.026
HORSE-COLIC	(368, 26)	0.810±0.042	<b>0.817±0.040</b>	0.805±0.044	0.806±0.043	0.805±0.045
CONGRESSIONAL-VOTING	(435, 17)	<b>0.599±0.046</b>	<b>0.597±0.046</b>	<b>0.598±0.047</b>	<b>0.597±0.049</b>	0.600±0.046
CYLINDER-BANDS	(512, 36)	0.778±0.041	0.763±0.045	0.782±0.041	<b>0.794±0.040</b>	0.795±0.040
BREAST-CANCER-WISC-DIAG	(569, 31)	<b>0.969±0.015</b>	<b>0.974±0.015</b>	<b>0.970±0.015</b>	<b>0.972±0.014</b>	0.973±0.016
ILPD-INDIAN-LIVER	(583, 10)	<b>0.718±0.035</b>	<b>0.715±0.039</b>	<b>0.719±0.035</b>	<b>0.719±0.035</b>	0.719±0.036
MONKS-2	(601, 7)	0.758±0.036	0.738±0.034	0.760±0.035	<b>0.773±0.035</b>	0.778±0.034
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>0.677±0.035</b>	0.667±0.032	<b>0.677±0.035</b>	<b>0.677±0.035</b>	0.678±0.035
CREDIT-APPROVAL	(690, 16)	<b>0.859±0.031</b>	0.856±0.029	<b>0.860±0.030</b>	<b>0.860±0.030</b>	0.859±0.030
BREAST-CANCER-WISC	(699, 10)	0.967±0.013	<b>0.969±0.012</b>	0.967±0.013	0.967±0.013	0.967±0.013
BLOOD	(748, 5)	<b>0.783±0.033</b>	<b>0.784±0.032</b>	<b>0.783±0.034</b>	<b>0.783±0.034</b>	0.779±0.031
PIMA	(768, 9)	<b>0.764±0.030</b>	<b>0.765±0.031</b>	<b>0.764±0.030</b>	<b>0.764±0.030</b>	0.763±0.030
MAMMOGRAPHIC	(961, 6)	<b>0.823±0.026</b>	<b>0.823±0.026</b>	<b>0.823±0.026</b>	<b>0.823±0.026</b>	0.823±0.026
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>0.769±0.027</b>	0.766±0.028	<b>0.768±0.027</b>	<b>0.768±0.027</b>	0.767±0.027
Bold Count		18	17	19	23	/

**Table 5. Binary classification:** log predictive density (higher is better) on different data sets with ARD kernel from the *Bayesian benchmarks* over 5-fold cross-validation. To ensure fair comparison we only include results on data sets where all methods have converged. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**.

	$(n, d)$	LA	EP	VI	Ours
FERTILITY	(100, 10)	<b>-0.417±0.096</b>	<b>-0.407±0.120</b>	<b>-0.439±0.116</b>	-0.456±0.115
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>-0.320±0.057</b>	<b>-0.325±0.085</b>	<b>-0.325±0.055</b>	<b>-0.320±0.061</b>
ACUTE-INFLAMMATION	(120, 7)	-0.112±0.005	-0.029±0.004	<b>-1.668±2.754</b>	<b>-0.005±0.001</b>
PARKINSONS	(195, 23)	-0.260±0.029	-0.214±0.086	<b>-0.062±0.062</b>	<b>-0.048±0.042</b>
BREAST-CANCER-WISC-PROG	(198, 34)	<b>-0.498±0.053</b>	<b>-0.476±0.068</b>	<b>-0.516±0.080</b>	<b>-0.644±0.092</b>
SPECT	(265, 23)	<b>-0.608±0.059</b>	<b>-0.614±0.064</b>	<b>-0.643±0.056</b>	-0.646±0.068
STATLOG-HEART	(270, 14)	-0.433±0.066	<b>-0.398±0.056</b>	-0.450±0.072	-0.472±0.082
HABERMAN-SURVIVAL	(306, 4)	<b>-0.535±0.049</b>	<b>-0.531±0.053</b>	<b>-0.539±0.050</b>	<b>-0.540±0.051</b>
IONOSPHERE	(351, 34)	<b>-0.243±0.058</b>	<b>-0.253±0.032</b>	<b>-0.208±0.086</b>	<b>-0.231±0.061</b>
CONGRESSIONAL-VOTING	(435, 17)	<b>-0.642±0.030</b>	<b>-0.639±0.033</b>	-0.687±0.043	<b>-0.644±0.081</b>
CYLINDER-BANDS	(512, 36)	<b>-0.468±0.063</b>	-0.501±0.061	<b>-0.454±0.068</b>	-0.478±0.064
BREAST-CANCER-WISC-DIAG	(569, 31)	<b>-0.098±0.024</b>	-0.122±0.031	<b>-0.075±0.044</b>	<b>-0.073±0.046</b>
ILPD-INDIAN-LIVER	(583, 10)	-0.531±0.030	<b>-0.518±0.035</b>	<b>-0.519±0.036</b>	<b>-0.518±0.037</b>
MONKS-2	(601, 7)	-0.458±0.029	-0.485±0.038	-0.424±0.037	<b>-0.399±0.036</b>
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>-0.633±0.022</b>	<b>-0.627±0.009</b>	<b>-0.633±0.022</b>	<b>-0.632±0.023</b>
CREDIT-APPROVAL	(690, 16)	-0.294±0.102	-0.148±0.015	<b>-0.905±1.126</b>	<b>-0.070±0.028</b>
BREAST-CANCER-WISC	(699, 10)	<b>-0.102±0.026</b>	<b>-0.101±0.028</b>	<b>-0.102±0.034</b>	<b>-0.104±0.035</b>
BLOOD	(748, 5)	<b>-0.477±0.047</b>	<b>-0.477±0.049</b>	<b>-0.477±0.048</b>	<b>-0.477±0.048</b>
PIMA	(768, 9)	<b>-0.473±0.034</b>	<b>-0.466±0.042</b>	<b>-0.474±0.037</b>	<b>-0.474±0.036</b>
MAMMOGRAPHIC	(961, 6)	<b>-0.399±0.039</b>	<b>-0.401±0.041</b>	<b>-0.400±0.040</b>	<b>-0.400±0.040</b>
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>-0.488±0.044</b>	<b>-0.485±0.046</b>	-0.501±0.053	<b>-0.503±0.054</b>
Bold Count		15	15	17	17

**Convergence of EP** To make sure EP fully converged, we initialized it with the trained hyperparameter values of our method and the log predictive density results and test set accuracies are given in Table 7 and Table 8. Compared with Table 9 and Table 10, regarding log predictive density EP has three more bolded results and regarding test set accuracies EP has two more bolded results. This underlines some of the issues associated with EP and speaks in favour of our method.

**Table 6. Binary classification:** test set accuracy (higher is better) on different data sets with ARD kernel from the *Bayesian benchmarks* over 5-fold cross-validation. To ensure fair comparison we only include results on data sets where all methods have converged. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**.

	$(n, d)$	LA	EP	VI	Ours
FERTILITY	(100, 10)	<b>0.840±0.049</b>	<b>0.880±0.060</b>	<b>0.840±0.049</b>	<b>0.840±0.049</b>
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>0.853±0.003</b>	<b>0.863±0.036</b>	<b>0.863±0.019</b>	<b>0.853±0.003</b>
ACUTE-INFLAMMATION	(120, 7)	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>0.833±0.173</b>	<b>1.000±0.000</b>
PARKINSONS	(195, 23)	<b>0.985±0.021</b>	<b>0.928±0.071</b>	<b>0.990±0.021</b>	<b>0.990±0.021</b>
BREAST-CANCER-WISC-PROG	(198, 34)	<b>0.768±0.029</b>	<b>0.768±0.048</b>	<b>0.773±0.053</b>	<b>0.753±0.049</b>
SPECT	(265, 23)	<b>0.702±0.056</b>	<b>0.709±0.054</b>	<b>0.709±0.057</b>	<b>0.675±0.066</b>
STATLOG-HEART	(270, 14)	0.796±0.039	<b>0.837±0.027</b>	0.789±0.049	<b>0.793±0.054</b>
HABERMAN-SURVIVAL	(306, 4)	<b>0.732±0.043</b>	<b>0.738±0.038</b>	<b>0.735±0.044</b>	<b>0.735±0.044</b>
IONOSPHERE	(351, 34)	<b>0.923±0.026</b>	<b>0.909±0.026</b>	<b>0.932±0.016</b>	0.917±0.016
CONGRESSIONAL-VOTING	(435, 17)	<b>0.607±0.027</b>	<b>0.598±0.030</b>	<b>0.536±0.091</b>	<b>0.607±0.034</b>
CYLINDER-BANDS	(512, 36)	<b>0.791±0.050</b>	0.768±0.043	<b>0.785±0.078</b>	<b>0.799±0.048</b>
BREAST-CANCER-WISC-DIAG	(569, 31)	<b>0.975±0.010</b>	<b>0.963±0.022</b>	<b>0.972±0.013</b>	<b>0.974±0.013</b>
ILPD-INDIAN-LIVER	(583, 10)	<b>0.707±0.033</b>	<b>0.724±0.033</b>	0.702±0.023	<b>0.705±0.030</b>
MONKS-2	(601, 7)	0.765±0.040	0.727±0.044	0.769±0.038	<b>0.785±0.031</b>
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>0.651±0.032</b>	<b>0.677±0.018</b>	<b>0.659±0.035</b>	<b>0.649±0.033</b>
CREDIT-APPROVAL	(690, 16)	<b>0.962±0.020</b>	0.961±0.013	<b>0.859±0.201</b>	<b>0.980±0.011</b>
BREAST-CANCER-WISC	(699, 10)	<b>0.961±0.011</b>	<b>0.958±0.011</b>	<b>0.964±0.010</b>	<b>0.963±0.012</b>
BLOOD	(748, 5)	<b>0.786±0.038</b>	<b>0.789±0.036</b>	<b>0.789±0.036</b>	<b>0.787±0.037</b>
PIMA	(768, 9)	<b>0.766±0.025</b>	<b>0.771±0.028</b>	<b>0.766±0.025</b>	<b>0.767±0.023</b>
MAMMOGRAPHIC	(961, 6)	<b>0.831±0.012</b>	<b>0.829±0.016</b>	<b>0.830±0.013</b>	<b>0.830±0.013</b>
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>0.787±0.028</b>	<b>0.779±0.027</b>	<b>0.782±0.028</b>	<b>0.778±0.028</b>
Bold Count		19	18	18	20

**Table 7. Binary classification:** log predictive density (higher is better) on different data sets from the *Bayesian benchmarks* over 5-fold cross-validation. We initialize EP with the trained hyperparameters values of our method. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**. The results are largely the same as Table 9, which means EP has converged.

	$(n, d)$	LA	EP (new init.)	VI	Ours
TRAINS	(10, 30)	<b>-0.695±0.011</b>	<b>-0.680±0.042</b>	<b>-0.692±0.023</b>	<b>-0.681±0.042</b>
BALLOONS	(16, 5)	<b>-0.707±0.146</b>	<b>-0.672±0.263</b>	<b>-0.711±0.239</b>	<b>-0.657±0.267</b>
FERTILITY	(100, 10)	<b>-0.379±0.099</b>	<b>-0.379±0.103</b>	<b>-0.378±0.103</b>	<b>-0.379±0.103</b>
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>-0.306±0.044</b>	<b>-0.295±0.058</b>	<b>-0.295±0.057</b>	<b>-0.296±0.059</b>
ACUTE-NEPHRITIS	(120, 7)	-0.203±0.010	-0.010±0.003	-0.006±0.002	<b>-0.005±0.001</b>
ACUTE-INFLAMMATION	(120, 7)	-0.172±0.033	-0.013±0.002	-0.008±0.001	<b>-0.007±0.001</b>
ECHOCARDIOGRAM	(131, 11)	<b>-0.420±0.073</b>	<b>-0.421±0.086</b>	<b>-0.420±0.085</b>	<b>-0.421±0.086</b>
HEPATITIS	(155, 20)	<b>-0.369±0.043</b>	<b>-0.362±0.052</b>	<b>-0.361±0.051</b>	<b>-0.361±0.052</b>
PARKINSONS	(195, 23)	-0.267±0.030	-0.153±0.050	-0.164±0.062	<b>-0.145±0.054</b>
BREAST-CANCER-WISC-PROG	(198, 34)	<b>-0.460±0.063</b>	<b>-0.458±0.070</b>	<b>-0.458±0.071</b>	<b>-0.458±0.070</b>
SPECT	(265, 23)	<b>-0.587±0.045</b>	<b>-0.590±0.051</b>	<b>-0.589±0.050</b>	<b>-0.590±0.051</b>
STATLOG-HEART	(270, 14)	<b>-0.392±0.071</b>	<b>-0.395±0.080</b>	<b>-0.394±0.080</b>	<b>-0.395±0.080</b>
HABERMAN-SURVIVAL	(306, 4)	<b>-0.535±0.050</b>	<b>-0.537±0.054</b>	<b>-0.537±0.054</b>	<b>-0.537±0.054</b>
IONOSPHERE	(351, 34)	-0.221±0.021	-0.173±0.031	<b>-0.169±0.031</b>	<b>-0.169±0.033</b>
HORSE-COLIC	(368, 26)	<b>-0.465±0.058</b>	<b>-0.469±0.068</b>	<b>-0.469±0.068</b>	<b>-0.469±0.068</b>
CONGRESSIONAL-VOTING	(435, 17)	<b>-0.636±0.026</b>	<b>-0.636±0.029</b>	<b>-0.636±0.029</b>	<b>-0.636±0.029</b>
CYLINDER-BANDS	(512, 36)	-0.482±0.030	<b>-0.442±0.038</b>	-0.457±0.038	<b>-0.441±0.040</b>
BREAST-CANCER-WISC-DIAG	(569, 31)	<b>-0.079±0.020</b>	<b>-0.069±0.040</b>	<b>-0.071±0.038</b>	<b>-0.071±0.045</b>
ILPD-INDIAN-LIVER	(583, 10)	<b>-0.516±0.033</b>	<b>-0.516±0.037</b>	<b>-0.516±0.037</b>	<b>-0.516±0.037</b>
MONKS-2	(601, 7)	-0.496±0.037	<b>-0.445±0.042</b>	-0.468±0.045	<b>-0.443±0.044</b>
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>-0.627±0.022</b>	<b>-0.627±0.023</b>	<b>-0.627±0.023</b>	<b>-0.627±0.023</b>
CREDIT-APPROVAL	(690, 16)	<b>-0.342±0.008</b>	<b>-0.341±0.009</b>	<b>-0.341±0.008</b>	<b>-0.341±0.009</b>
BREAST-CANCER-WISC	(699, 10)	<b>-0.091±0.030</b>	<b>-0.091±0.037</b>	<b>-0.091±0.037</b>	<b>-0.091±0.037</b>
BLOOD	(748, 5)	<b>-0.475±0.043</b>	<b>-0.476±0.044</b>	<b>-0.476±0.044</b>	<b>-0.476±0.044</b>
PIMA	(768, 9)	<b>-0.470±0.031</b>	<b>-0.471±0.033</b>	<b>-0.471±0.033</b>	<b>-0.471±0.033</b>
MAMMOGRAPHIC	(961, 6)	<b>-0.408±0.037</b>	<b>-0.408±0.039</b>	<b>-0.408±0.039</b>	<b>-0.408±0.039</b>
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>-0.493±0.040</b>	<b>-0.494±0.043</b>	<b>-0.493±0.043</b>	<b>-0.494±0.043</b>
Bold Count		21	23	22	27

**Convergence of VI** To ensure that the comparison to VI is as fair as possible, we ran both the vanilla VI approach where we jointly optimize  $\xi$  and  $\theta$  using L-BFGS-B and the CVI approach by Khan & Lin (2017) that uses natural gradients for the E-step for faster convergence. The speed of convergence was not the main interest, and thus we verified that the VI results obtained by the two different approaches gave the same results—thus confirming that the VI results presented in the tables are all for converged optimization runs.

**Table 8. Binary classification:** test set accuracy (higher is better) on different data sets from the *Bayesian benchmarks* over 5-fold cross-validation. We initialize EP with the trained hyperparameters values of our method. Best results and those not statistically significantly different from them under a paired *t*-test are **bolded**. The results are largely the same as Table 10, which means EP has converged.

	$(n, d)$	LA	EP (new init.)	VI	Ours
TRAINS	(10, 30)	<b>0.700±0.400</b>	<b>0.800±0.245</b>	<b>0.800±0.245</b>	<b>0.800±0.245</b>
BALLOONS	(16, 5)	<b>0.667±0.298</b>	<b>0.667±0.298</b>	<b>0.667±0.298</b>	<b>0.667±0.298</b>
FERTILITY	(100, 10)	<b>0.880±0.060</b>	<b>0.880±0.060</b>	<b>0.880±0.060</b>	<b>0.880±0.060</b>
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>0.863±0.036</b>	<b>0.863±0.036</b>	<b>0.863±0.036</b>	<b>0.863±0.036</b>
ACUTE-NEPHRITIS	(120, 7)	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>
ACUTE-INFLAMMATION	(120, 7)	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>	<b>1.000±0.000</b>
ECHOCARDIOGRAM	(131, 11)	<b>0.808±0.061</b>	<b>0.778±0.063</b>	<b>0.785±0.076</b>	<b>0.778±0.063</b>
HEPATITIS	(155, 20)	<b>0.813±0.024</b>	<b>0.832±0.024</b>	<b>0.813±0.024</b>	<b>0.832±0.024</b>
PARKINSONS	(195, 23)	<b>0.959±0.048</b>	<b>0.959±0.048</b>	<b>0.944±0.050</b>	<b>0.959±0.048</b>
BREAST-CANCER-WISC-PROG	(198, 34)	<b>0.793±0.051</b>	<b>0.793±0.051</b>	<b>0.793±0.051</b>	<b>0.793±0.051</b>
SPECT	(265, 23)	<b>0.706±0.031</b>	<b>0.706±0.023</b>	<b>0.702±0.028</b>	<b>0.706±0.023</b>
STATLOG-HEART	(270, 14)	<b>0.830±0.036</b>	<b>0.826±0.034</b>	<b>0.830±0.036</b>	<b>0.826±0.034</b>
HABERMAN-SURVIVAL	(306, 4)	<b>0.725±0.039</b>	<b>0.725±0.039</b>	<b>0.725±0.039</b>	<b>0.725±0.039</b>
IONOSPHERE	(351, 34)	<b>0.932±0.025</b>	<b>0.946±0.028</b>	<b>0.949±0.026</b>	<b>0.946±0.028</b>
HORSE-COLIC	(368, 26)	<b>0.799±0.035</b>	<b>0.802±0.039</b>	0.791±0.045	<b>0.802±0.039</b>
CONGRESSIONAL-VOTING	(435, 17)	<b>0.605±0.024</b>	<b>0.591±0.016</b>	<b>0.595±0.013</b>	<b>0.591±0.016</b>
CYLINDER-BANDS	(512, 36)	0.785±0.037	<b>0.803±0.030</b>	<b>0.791±0.034</b>	<b>0.803±0.031</b>
BREAST-CANCER-WISC-DIAG	(569, 31)	0.972±0.010	<b>0.975±0.009</b>	<b>0.979±0.009</b>	<b>0.977±0.011</b>
ILPD-INDIAN-LIVER	(583, 10)	<b>0.719±0.025</b>	<b>0.719±0.028</b>	<b>0.715±0.027</b>	<b>0.719±0.028</b>
MONKS-2	(601, 7)	<b>0.740±0.049</b>	<b>0.762±0.045</b>	<b>0.744±0.046</b>	<b>0.762±0.045</b>
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>0.678±0.028</b>	<b>0.678±0.028</b>	<b>0.678±0.028</b>	<b>0.678±0.028</b>
CREDIT-APPROVAL	(690, 16)	<b>0.859±0.026</b>	<b>0.862±0.026</b>	<b>0.864±0.024</b>	<b>0.862±0.026</b>
BREAST-CANCER-WISC	(699, 10)	<b>0.969±0.016</b>	<b>0.969±0.016</b>	<b>0.969±0.016</b>	<b>0.969±0.016</b>
BLOOD	(748, 5)	<b>0.785±0.042</b>	<b>0.786±0.044</b>	<b>0.786±0.044</b>	<b>0.786±0.044</b>
PIMA	(768, 9)	<b>0.764±0.027</b>	<b>0.767±0.026</b>	<b>0.767±0.026</b>	<b>0.767±0.026</b>
MAMMOGRAPHIC	(961, 6)	<b>0.824±0.019</b>	<b>0.823±0.019</b>	<b>0.823±0.019</b>	<b>0.823±0.019</b>
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>0.770±0.039</b>	<b>0.769±0.037</b>	<b>0.768±0.038</b>	<b>0.769±0.037</b>
Bold Count		25	27	26	27

**Table 9. Binary classification:** log predictive density (higher is better) on different data sets from the *Bayesian benchmarks* (mean  $\pm$  standard deviation over 5-fold cross-validation). The best results and those not statistically significantly different from them under a paired *t*-test are **bolded**. For the classification accuracy all methods perform comparably, which is to be expected as accuracy alone only captures where the decision boundary has been draw, completely disregarding second-order information.

#		$(n, d)$	LA	EP	VI	Ours
1	TRAINS	(10, 30)	<b>-0.695±0.011</b>	<b>-0.687±0.023</b>	<b>-0.692±0.023</b>	<b>-0.681±0.042</b>
2	BALLOONS	(16, 5)	<b>-0.707±0.146</b>	<b>-0.684±0.161</b>	<b>-0.711±0.239</b>	<b>-0.657±0.267</b>
3	FERTILITY	(100, 10)	<b>-0.379±0.099</b>	<b>-0.384±0.138</b>	<b>-0.378±0.103</b>	<b>-0.379±0.103</b>
4	PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>-0.306±0.044</b>	<b>-0.316±0.060</b>	<b>-0.295±0.057</b>	<b>-0.296±0.059</b>
5	ACUTE-NEPHRITIS	(120, 7)	-0.203±0.010	-0.047±0.009	-0.006±0.002	<b>-0.005±0.001</b>
6	ACUTE-INFLAMMATION	(120, 7)	-0.172±0.033	-0.055±0.005	-0.008±0.001	<b>-0.007±0.001</b>
7	ECHOCARDIOGRAM	(131, 11)	<b>-0.420±0.073</b>	<b>-0.412±0.084</b>	<b>-0.420±0.085</b>	<b>-0.421±0.086</b>
8	HEPATITIS	(155, 20)	<b>-0.369±0.043</b>	<b>-0.376±0.046</b>	<b>-0.361±0.051</b>	<b>-0.361±0.052</b>
9	PARKINSONS	(195, 23)	-0.267±0.030	-0.302±0.090	-0.164±0.062	<b>-0.145±0.054</b>
10	BREAST-CANCER-WISC-PROG	(198, 34)	<b>-0.460±0.063</b>	<b>-0.478±0.083</b>	<b>-0.458±0.071</b>	<b>-0.458±0.070</b>
11	SPECT	(265, 23)	<b>-0.587±0.045</b>	<b>-0.587±0.052</b>	<b>-0.589±0.050</b>	<b>-0.590±0.051</b>
12	STATLOG-HEART	(270, 14)	<b>-0.392±0.071</b>	<b>-0.380±0.058</b>	<b>-0.394±0.080</b>	<b>-0.395±0.080</b>
13	HABERMAN-SURVIVAL	(306, 4)	<b>-0.535±0.050</b>	<b>-0.539±0.059</b>	<b>-0.537±0.054</b>	<b>-0.537±0.054</b>
14	IONOSPHERE	(351, 34)	-0.221±0.021	-0.227±0.016	<b>-0.169±0.031</b>	<b>-0.169±0.033</b>
15	HORSE-COLIC	(368, 26)	-0.465±0.058	<b>-0.455±0.060</b>	<b>-0.466±0.068</b>	<b>-0.469±0.068</b>
16	CONGRESSIONAL-VOTING	(435, 17)	<b>-0.636±0.026</b>	<b>-0.633±0.027</b>	<b>-0.636±0.029</b>	<b>-0.636±0.029</b>
17	CYLINDER-BANDS	(512, 36)	-0.482±0.030	-0.495±0.030	-0.457±0.038	<b>-0.441±0.040</b>
18	BREAST-CANCER-WISC-DIAG	(569, 31)	<b>-0.079±0.020</b>	-0.140±0.015	<b>-0.071±0.038</b>	<b>-0.071±0.045</b>
19	ILPD-INDIAN-LIVER	(583, 10)	<b>-0.516±0.033</b>	<b>-0.521±0.033</b>	<b>-0.516±0.037</b>	<b>-0.516±0.037</b>
20	MONKS-2	(601, 7)	-0.496±0.037	-0.518±0.042	-0.468±0.045	<b>-0.443±0.044</b>
21	STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>-0.627±0.022</b>	<b>-0.636±0.030</b>	<b>-0.627±0.023</b>	<b>-0.627±0.023</b>
22	CREDIT-APPROVAL	(690, 16)	<b>-0.342±0.008</b>	<b>-0.343±0.012</b>	<b>-0.341±0.008</b>	<b>-0.341±0.009</b>
23	BREAST-CANCER-WISC	(699, 10)	<b>-0.091±0.030</b>	<b>-0.092±0.028</b>	<b>-0.091±0.037</b>	<b>-0.091±0.037</b>
24	BLOOD	(748, 5)	<b>-0.475±0.043</b>	<b>-0.476±0.045</b>	<b>-0.476±0.044</b>	<b>-0.476±0.044</b>
25	PIMA	(768, 9)	<b>-0.470±0.031</b>	<b>-0.472±0.035</b>	<b>-0.471±0.033</b>	<b>-0.471±0.033</b>
26	MAMMOGRAPHIC	(961, 6)	<b>-0.408±0.037</b>	<b>-0.408±0.041</b>	<b>-0.408±0.039</b>	<b>-0.408±0.039</b>
27	STATLOG-GERMAN-CREDIT	(1000, 25)	<b>-0.493±0.040</b>	<b>-0.493±0.041</b>	<b>-0.493±0.043</b>	<b>-0.494±0.043</b>
Bold Count			20	20	22	27



Table 10. **Binary classification:** test set accuracy (higher is better) on different data sets from the *Bayesian benchmarks* (mean  $\pm$  standard deviation over 5-fold cross-validation). The best results and those not statistically significantly different from them under a paired  $t$ -test are **bolded**. For the classification accuracy all methods perform comparably, which is to be expected as accuracy alone only captures where the decision boundary has been draw, completely disregarding second-order information.

	$(n, d)$	LA	EP	VI	Ours
TRAINS	(10, 30)	<b>0.700<math>\pm</math>0.400</b>	<b>0.800<math>\pm</math>0.245</b>	<b>0.800<math>\pm</math>0.245</b>	<b>0.800<math>\pm</math>0.245</b>
BALLOONS	(16, 5)	<b>0.667<math>\pm</math>0.298</b>	<b>0.667<math>\pm</math>0.298</b>	<b>0.667<math>\pm</math>0.298</b>	<b>0.667<math>\pm</math>0.298</b>
FERTILITY	(100, 10)	<b>0.880<math>\pm</math>0.060</b>	<b>0.880<math>\pm</math>0.060</b>	<b>0.880<math>\pm</math>0.060</b>	<b>0.880<math>\pm</math>0.060</b>
PITTSBURG-BRIDGES-T-OR-D	(102, 8)	<b>0.863<math>\pm</math>0.036</b>	<b>0.883<math>\pm</math>0.066</b>	<b>0.863<math>\pm</math>0.036</b>	<b>0.863<math>\pm</math>0.036</b>
ACUTE-NEPHRITIS	(120, 7)	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>
ACUTE-INFLAMMATION	(120, 7)	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>	<b>1.000<math>\pm</math>0.000</b>
ECHOCARDIOGRAM	(131, 11)	0.808 $\pm$ 0.061	<b>0.847<math>\pm</math>0.043</b>	<b>0.785<math>\pm</math>0.076</b>	0.778 $\pm$ 0.063
HEPATITIS	(155, 20)	<b>0.813<math>\pm</math>0.024</b>	<b>0.813<math>\pm</math>0.032</b>	<b>0.813<math>\pm</math>0.024</b>	<b>0.832<math>\pm</math>0.024</b>
PARKINSONS	(195, 23)	<b>0.959<math>\pm</math>0.048</b>	0.892 $\pm$ 0.075	<b>0.944<math>\pm</math>0.050</b>	<b>0.959<math>\pm</math>0.048</b>
BREAST-CANCER-WISC-PROG	(198, 34)	<b>0.793<math>\pm</math>0.051</b>	<b>0.793<math>\pm</math>0.069</b>	<b>0.793<math>\pm</math>0.051</b>	<b>0.793<math>\pm</math>0.051</b>
SPECT	(265, 23)	<b>0.706<math>\pm</math>0.031</b>	<b>0.698<math>\pm</math>0.027</b>	<b>0.702<math>\pm</math>0.028</b>	<b>0.706<math>\pm</math>0.023</b>
STATLOG-HEART	(270, 14)	<b>0.830<math>\pm</math>0.036</b>	<b>0.833<math>\pm</math>0.039</b>	<b>0.830<math>\pm</math>0.036</b>	<b>0.826<math>\pm</math>0.034</b>
HABERMAN-SURVIVAL	(306, 4)	<b>0.725<math>\pm</math>0.039</b>	<b>0.719<math>\pm</math>0.036</b>	<b>0.725<math>\pm</math>0.039</b>	<b>0.725<math>\pm</math>0.039</b>
IONOSPHERE	(351, 34)	<b>0.932<math>\pm</math>0.025</b>	<b>0.932<math>\pm</math>0.024</b>	<b>0.946<math>\pm</math>0.024</b>	<b>0.946<math>\pm</math>0.028</b>
HORSE-COLIC	(368, 26)	<b>0.799<math>\pm</math>0.035</b>	<b>0.807<math>\pm</math>0.032</b>	<b>0.791<math>\pm</math>0.045</b>	<b>0.802<math>\pm</math>0.039</b>
CONGRESSIONAL-VOTING	(435, 17)	<b>0.605<math>\pm</math>0.024</b>	<b>0.593<math>\pm</math>0.014</b>	<b>0.595<math>\pm</math>0.013</b>	<b>0.591<math>\pm</math>0.016</b>
CYLINDER-BANDS	(512, 36)	0.785 $\pm$ 0.037	<b>0.777<math>\pm</math>0.044</b>	<b>0.791<math>\pm</math>0.034</b>	<b>0.803<math>\pm</math>0.031</b>
BREAST-CANCER-WISC-DIAG	(569, 31)	<b>0.972<math>\pm</math>0.010</b>	<b>0.974<math>\pm</math>0.019</b>	<b>0.977<math>\pm</math>0.012</b>	<b>0.977<math>\pm</math>0.011</b>
ILPD-INDIAN-LIVER	(583, 10)	<b>0.719<math>\pm</math>0.025</b>	<b>0.712<math>\pm</math>0.020</b>	<b>0.717<math>\pm</math>0.029</b>	<b>0.719<math>\pm</math>0.028</b>
MONKS-2	(601, 7)	<b>0.740<math>\pm</math>0.049</b>	0.719 $\pm$ 0.039	<b>0.744<math>\pm</math>0.046</b>	<b>0.762<math>\pm</math>0.045</b>
STATLOG-AUSTRALIAN-CREDIT	(690, 15)	<b>0.678<math>\pm</math>0.028</b>	<b>0.677<math>\pm</math>0.021</b>	<b>0.678<math>\pm</math>0.028</b>	<b>0.678<math>\pm</math>0.028</b>
CREDIT-APPROVAL	(690, 16)	<b>0.859<math>\pm</math>0.026</b>	<b>0.864<math>\pm</math>0.020</b>	<b>0.864<math>\pm</math>0.024</b>	<b>0.862<math>\pm</math>0.026</b>
BREAST-CANCER-WISC	(699, 10)	<b>0.969<math>\pm</math>0.016</b>	<b>0.969<math>\pm</math>0.016</b>	<b>0.969<math>\pm</math>0.016</b>	<b>0.969<math>\pm</math>0.016</b>
BLOOD	(748, 5)	<b>0.785<math>\pm</math>0.042</b>	<b>0.785<math>\pm</math>0.041</b>	<b>0.786<math>\pm</math>0.044</b>	<b>0.786<math>\pm</math>0.044</b>
PIMA	(768, 9)	<b>0.764<math>\pm</math>0.027</b>	<b>0.766<math>\pm</math>0.025</b>	<b>0.767<math>\pm</math>0.026</b>	<b>0.767<math>\pm</math>0.026</b>
MAMMOGRAPHIC	(961, 6)	<b>0.824<math>\pm</math>0.019</b>	<b>0.823<math>\pm</math>0.017</b>	<b>0.823<math>\pm</math>0.019</b>	<b>0.823<math>\pm</math>0.019</b>
STATLOG-GERMAN-CREDIT	(1000, 25)	<b>0.770<math>\pm</math>0.039</b>	<b>0.763<math>\pm</math>0.038</b>	<b>0.768<math>\pm</math>0.038</b>	<b>0.769<math>\pm</math>0.037</b>
Bold Count		25	25	27	26

## E. Robust (Student- $t$ ) Regression

For additional insight, we include results for a Student- $t$  likelihood model that allows heavy-tailed noise in the observations. This likelihood is not log-concave, which makes LA and EP more challenging. Previous works in the field do not list many standard data sets for benchmarking, but we have gathered three previously used benchmark problems for comparison.

### E.1. Experiment Setup

We preprocessed the data and selected an interval of 1000 data points from the original data. We initialize all methods with unit lengthscale, kernel magnitude, and likelihood variance. We fix the degrees of freedom in the likelihood to  $\nu = 3$ . For LA and EP, we follow the methods designed in [Jylänki et al. \(2011\)](#). For VI and our hybrid training procedure, E and M steps each have 20 iterations. In the E-step we set the learning rate of natural gradient descent to be 0.001. In the M-step we use the Adam optimizer ([Kingma & Ba, 2015](#)) and set the learning rate to 0.001. For NEAL we set the maximal steps to be 2000 and for BOSTON and STOCK we set the maximal steps to be 5000.

## F. Evaluation on Sparse Approximation

To further study the proposed approach, we include results for a couple of sparse GP classification problems as proof-of-concept.

Table 11. Data sets information for sparse binary classification.

Data set	TITANIC	BANK	TWONORM	MUSHROOM	MAGIC
$(n, d)$	(2201, 4)	(4521, 17)	(7400, 21)	(8124, 22)	(19020, 11)

## F.1. Experiment Setup

We initialize all methods with unit lengthscale and magnitude. We use  $k$ -means to select a set of 500 inducing points and then fix the location of inducing points. Learning the inducing points would add an additional layer of optimization, and this is not in the scope of this paper. We use 20 iterations for both the E and M steps. In the E-step we set the learning rate of natural gradient descent to 0.1. In the M-step we use the Adam optimizer (Kingma & Ba, 2015) and set the learning rate to 0.01. The data sets information are given in Table 11.