

---

# Convergent Expectation Propagation in Linear Models with Spike-and-slab Priors

---

**José Miguel Hernández-Lobato**

Department of Engineering  
University of Cambridge  
Trumpington Street, Cambridge  
CB2 1PZ, United Kingdom  
jmh233@eng.cam.ac.uk

**Daniel Hernández-Lobato**

Computer Science Department  
Universidad Autónoma de Madrid  
Francisco Tomás y Valiente, 11  
28049 Madrid, Spain  
daniel.hernandez@uam.es

## Abstract

Exact inference in the linear regression model with spike and slab priors is often intractable. Expectation propagation (EP) can be used for approximate inference. However, the regular sequential form of EP (R-EP) may fail to converge in this model when the size of the training set is very small. As an alternative, we propose a provably convergent EP algorithm (PC-EP). PC-EP is proved to minimize an energy function which, under some constraints, is bounded from below and whose stationary points coincide with the solution of R-EP. Experiments with synthetic data indicate that when R-EP does not converge, the approximation generated by PC-EP is often better. By contrast, when R-EP converges, both methods perform similarly.

## 1 Introduction

Exact Bayesian inference is often intractable in many probabilistic models of practical interest. The computational cost of marginalization operations scales exponentially in the number of variables or these operations require to compute high-dimensional integrals that do not have a closed-form analytical solution. In practice, we have to use some form of approximation. Approximate inference in large applications is frequently implemented using deterministic methods. These have often less computational cost than other alternatives such as sampling. Deterministic methods approximate the exact posterior by a tractable parametric distribution whose parameters are selected by solving optimization problems. Expectation propagation (EP) is one of the most successful techniques for deterministic approximate inference [1]. However, a disadvantage of EP is that, in its standard sequential form, it is not guaranteed to converge.

In this work we focus on the linear regression model with spike-and-slab priors (LRMSSP). The standard sequential EP method often generates in this model very accurate approximations of the posterior distribution [2]. However, EP may fail to converge in some extreme cases in which the number of training instances is very small. To avoid this, we introduce a provably convergent EP algorithm (PC-EP) for approximate inference in the LRMSSP. PC-EP is based on the double loop algorithm described in [3]. Each outer iteration of PC-EP is proved to minimize an energy function whose stationary points coincide with the solution of regular EP [4]. The main difference between PC-EP and the algorithm proposed in [3] is that PC-EP constrains the variance parameters to be positive. This ensures that the energy function minimized by PC-EP is bounded from below. The boundedness of this function is a necessary condition to guarantee convergence. Experiments with synthetic data illustrate the advantages of PC-EP over regular EP. In these experiments, when regular EP does not converge, the posterior approximation generated by PC-EP is often better. By contrast, when regular EP converges, both methods perform similarly.

## 2 The Linear Regression Model with Spike-and-slab Priors

We focus on the problem of Bayesian inference in the linear regression model with spike-and-slab priors (LRMSSP). Consider  $n$  feature vectors, with  $d$  real components each, encoded by the design matrix  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$  and associated target values  $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathbb{R}^n$ . The LRMSSP assumes that  $\mathbf{y} = \mathbf{X}\mathbf{w} + \boldsymbol{\epsilon}$ , where  $\mathbf{w}$  is an unknown  $d$ -dimensional vector of real coefficients and  $\boldsymbol{\epsilon}$  is an  $n$ -dimensional vector of independent Gaussian noise with variance  $\sigma^2$ . Given  $\mathbf{X}$  and  $\mathbf{y}$ , the likelihood of  $\mathbf{w}$  is a Gaussian function:  $\mathcal{P}(\mathbf{y}|\mathbf{w}, \mathbf{X}) = \mathcal{N}(\mathbf{y}|\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I})$ . The prior for  $\mathbf{w}$  is a product of spike-and-slab factors [5]:

$$\mathcal{P}(\mathbf{w}) = \prod_{i=1}^d [p_s \mathcal{N}(w_i|0, v_s) + (1 - p_s) \delta(w_i)], \quad (1)$$

where  $\mathcal{N}(\cdot|0, v_s)$  denotes a Gaussian density with zero mean and variance  $v_s$  (the slab),  $\delta(\cdot)$  is a delta function centered at zero (the spike), and  $p_s$  is the prior probability that any of the components of  $\mathbf{w}$  is different from zero. The LRMSSP has special practical interest in regression problems with more features than training instances (that is,  $d \gg n$ ). In this scenario, the assumption of sparsity is used to reduce over-fitting by limiting the number of influential features [6]. The spike-and-slab prior (1) incorporates this assumption by inducing a bi-separation in the model coefficients. A small number of coefficients are considered to be different from zero. These are the coefficients sampled from the slab. At the same time, a reduced number of coefficients are considered to be exactly zero. These are the coefficients sampled from the spike.

Given  $\mathbf{X}$  and  $\mathbf{y}$ , the posterior distribution for  $\mathbf{w}$  can be computed using Bayes' rule:

$$\mathcal{P}(\mathbf{w}|\mathbf{X}, \mathbf{y}) = \frac{\mathcal{P}(\mathbf{y}|\mathbf{w}, \mathbf{X})\mathcal{P}(\mathbf{w})}{\mathcal{P}(\mathbf{y}|\mathbf{X})}. \quad (2)$$

The central operation in the application of Bayesian methods is the computation of marginalizations or expectations with respect to this distribution. However, these operations are usually intractable. As a solution, we use EP [4] to produce a deterministic approximation to (2) which is tractable. The numerator in the right part of (2) can be written as

$$\mathcal{P}(\mathbf{y}|\mathbf{w}, \mathbf{X})\mathcal{P}(\mathbf{w}) = \mathcal{N}(\mathbf{y}|\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I}) \prod_{i=1}^d f_i(w_i), \quad (3)$$

where  $f_i(w_i) = p_s \mathcal{N}(w_i|0, v_s) + (1 - p_s) \delta(w_i)$ . EP replaces each  $f_i$  by a Gaussian factor  $\tilde{f}_i(w_i) = \exp\{\tilde{v}_{1i}w_i - \frac{1}{2}\tilde{v}_{2i}w_i^2 + \tilde{z}_i\}$ , where  $\tilde{v}_{1i}$ ,  $\tilde{v}_{2i}$  and  $\tilde{z}_i$  are free parameters. Let  $\mathcal{Q}(\mathbf{w})$  be the product of the likelihood  $\mathcal{N}(\mathbf{y}|\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I})$  and the Gaussian factors  $\tilde{f}_1, \dots, \tilde{f}_d$ . Then  $\mathcal{Q}$  is an unnormalized multivariate Gaussian distribution:

$$\mathcal{Q}(\mathbf{w}) = \mathcal{N}(\mathbf{y}|\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I}) \prod_{i=1}^d \tilde{f}_i(w_i) \propto \mathcal{N}(\mathbf{w}|\mathbf{m}, \mathbf{A}^{-1}), \quad (4)$$

where  $\mathbf{A} = \sigma^{-2}\mathbf{X}^T\mathbf{X} + \text{diag}(\tilde{\mathbf{v}}_2)$ ,  $\mathbf{m} = \mathbf{A}^{-1}(\tilde{\mathbf{v}}_1 + \sigma^{-2}\mathbf{X}^T\mathbf{y})$ ,  $\tilde{\mathbf{v}}_1 = (v_{11}, \dots, v_{1d})^T$  and  $\tilde{\mathbf{v}}_2 = (v_{21}, \dots, v_{2d})^T$ . Note that the posterior (2) is obtained when (3) is normalized so that it integrates one. Similarly, the EP approximation to the posterior is given by the normalized version of  $\mathcal{Q}$ .

The sequential implementation of EP iteratively updates the  $\tilde{f}_i$  until reaching a stationary point of a specific energy function [4]. However, the EP update operations do not guarantee the minimization of this energy function and the method may sometimes fail to converge. Non-convergence can be prevented by *damping* the EP update operations [7]. After performing one update operation, damping sets the updated factor  $\tilde{f}_i$  to a log-convex combination of the factor before (old) and after (new) the update, that is,  $\tilde{f}_i = (\tilde{f}_i^{\text{new}})^\tau (\tilde{f}_i^{\text{old}})^{1-\tau}$ , with  $\tau \in [0, 1]$ . However, even when damping is used, EP can fail to converge in some extreme cases in which the number of training instances is very small. When this happens, we can limit the maximum number of iterations of the algorithm and hope that the resulting approximation is accurate enough. As an alternative, we present a provably convergent EP algorithm (PC-EP).

### 3 Provably Convergent EP for the LRMSSP

The EP solutions in the LRMSSP are in a one-to-one correspondence with stationary points of the following objective [8, 9]:

$$\min_v \max_{\hat{v}, \tilde{v}} E(v, \hat{v}, \tilde{v}), \quad (5)$$

where  $v = \{\mathbf{v}_1, \mathbf{v}_2\}$ ,  $\tilde{v} = \{\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2\}$ ,  $\hat{v} = \{\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2\}$ , the elements in these tuples are  $d$ -dimensional real vectors satisfying

$$\mathbf{v}_1 = \tilde{\mathbf{v}}_1 + \hat{\mathbf{v}}_1, \quad \mathbf{v}_2 = \tilde{\mathbf{v}}_2 + \hat{\mathbf{v}}_2, \quad (6)$$

$E(v, \hat{v}, \tilde{v})$  is an energy function given by

$$E(v, \hat{v}, \tilde{v}) = -\log Z(\tilde{v}) - \log \hat{Z}(\hat{v}) + \log \tilde{Z}(v) \quad (7)$$

and  $Z(\tilde{v})$ ,  $\hat{Z}(\hat{v})$  and  $\tilde{Z}(v)$  are the following normalization constants:

$$Z(\tilde{v}) = \int \mathcal{N}(\mathbf{y}|\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I}) \prod_{i=1}^d \exp \left\{ \tilde{v}_{1i}w_i - \frac{1}{2}\tilde{v}_{2i}w_i^2 \right\} d\mathbf{w}, \quad (8)$$

$$\hat{Z}(\hat{v}) = \prod_{i=1}^d \int \exp \left\{ \hat{v}_{1i}w_i - \frac{1}{2}\hat{v}_{2i}w_i^2 \right\} f_i(w_i)dw_i, \quad (9)$$

$$\tilde{Z}(v) = \prod_{i=1}^d \int \exp \left\{ v_{1i}w_i - \frac{1}{2}v_{2i}w_i^2 \right\} dw_i, \quad (10)$$

where  $\hat{\mathbf{v}}_1 = (\hat{v}_{11}, \dots, \hat{v}_{1d})^T$ ,  $\hat{\mathbf{v}}_2 = (\hat{v}_{21}, \dots, \hat{v}_{2d})^T$ ,  $\mathbf{v}_1 = (v_{11}, \dots, v_{1d})^T$ ,  $\mathbf{v}_2 = (v_{21}, \dots, v_{2d})^T$  and the vectors  $\mathbf{v}_1$  and  $\mathbf{v}_2$  contain the natural parameters of the marginals of  $\mathcal{Q}$ . Once a stationary point of (5) is found, the optimal value of  $\tilde{z}_1, \dots, \tilde{z}_d$  can be computed very easily as a function of the solution for  $\tilde{\mathbf{v}}_1$  and  $\tilde{\mathbf{v}}_2$ . See [10] for further details. To find a stationary point of the objective energy, we follow [3] and attempt to find a minimum of this function. However, further constraints have to be imposed on  $\tilde{\mathbf{v}}_2$ ,  $\hat{\mathbf{v}}_2$  and  $\mathbf{v}_2$  to guarantee that the energy  $E(v, \hat{v}, \tilde{v})$  is lower bounded.

#### 3.1 Lower Bound on the Objective Energy

Consider, in addition to constraints (6), the following inequality constraints on  $\tilde{\mathbf{v}}_2$ ,  $\hat{\mathbf{v}}_2$  and  $\mathbf{v}_2$ :

$$\tilde{\mathbf{v}}_2 \succeq \boldsymbol{\varepsilon}, \quad \hat{\mathbf{v}}_2 \succeq \boldsymbol{\varepsilon}, \quad \mathbf{v}_2 \succeq 3\boldsymbol{\varepsilon}, \quad (11)$$

where  $\boldsymbol{\varepsilon}$  is a  $d$ -dimensional vector whose components are all equal to a small positive constant  $\varepsilon$ . Then, the objective (5) is bounded from below. For any  $\mathbf{v}_1$  and  $\mathbf{v}_2$  such that  $\mathbf{v}_2 \succeq 3\boldsymbol{\varepsilon}$ , we can choose  $\tilde{\mathbf{v}}_1 = \hat{\mathbf{v}}_1 = 0.5\mathbf{v}_1$  and  $\tilde{\mathbf{v}}_2 = \hat{\mathbf{v}}_2 = 0.5\mathbf{v}_2$ . These parameter values satisfy the required equality and inequality constraints and give a lower objective than the maximizers of  $E(v, \hat{v}, \tilde{v})$  when  $v$  is held fixed. This is more compactly written as  $\hat{v} = 0.5v$  and  $\tilde{v} = 0.5v$ , where  $0.5v = \{0.5\mathbf{v}_1, 0.5\mathbf{v}_2\}$ . For these values of  $\hat{v}$  and  $\tilde{v}$  we can lower bound each term of the EP energy:

$$-\log Z(0.5v) \geq \frac{n}{2} \log(2\pi\sigma^2) - \frac{d}{2} \log(4\pi) - \sum_{i=1}^d \left( \frac{v_{1i}^2}{4v_{2i}} - \frac{1}{2} \log v_{2i} \right), \quad (12)$$

$$\log \tilde{Z}(v) \geq \frac{d}{2} \log(2\pi) + \sum_{i=1}^d \left( \frac{v_{1i}^2}{2v_{2i}} - \frac{1}{2} \log v_{2i} \right), \quad (13)$$

$$-\log \hat{Z}(0.5v) \geq -\frac{d}{2} \log(4\pi) - \sum_{i=1}^d \left( \frac{v_{1i}^2}{4v_{2i}} - \frac{1}{2} \log v_{2i} \right) - \sum_{i=1}^d \log K_i, \quad (14)$$

with  $-\log K_i \geq 1/2 \log(4\pi) - 1/2 \log v_{2i}$ . Since  $\mathbf{v}_2 \succeq 3\boldsymbol{\varepsilon}$ , all these bounds are real and positive and their sum is  $n/2 \log(2\pi\sigma^2) - d/2 \log(2)$ , which is a lower bound on the energy function minimized by the convergent EP algorithm. Note that if the components of  $\tilde{\mathbf{v}}_2$ ,  $\hat{\mathbf{v}}_2$  and  $\mathbf{v}_2$  could be negative, some of the normalization constants (8), (9) and (10) might be infinite and the EP energy would be unbounded. The inequality constraints in (11) guarantee that this cannot occur.

### 3.2 Provably Convergent Algorithm

The first step of the provably convergent EP algorithm (PC-EP) maximizes  $E(v, \hat{v}, \tilde{v})$  with respect to  $\hat{v}$  and  $\tilde{v}$  given the current value of  $v$ , which is denoted by  $v^{(t)} = \{\mathbf{v}_1^{(t)}, \mathbf{v}_2^{(t)}\}$ , that is,

$$E(v^{(t)}) = \max_{\hat{v}, \tilde{v}} E(v^{(t)}, \hat{v}, \tilde{v}). \quad (15)$$

s.t (6), (11)

This step is implemented using standard optimization methods such as Quasi-Newton techniques, optimizing only on either  $\tilde{v}$  or  $\hat{v}$  since (6) allows us to identify one of these tuples given the other and the current value of  $v$ . The Lagrangian for this optimization problem is

$$\begin{aligned} \mathcal{L}(v^{(t)}, \hat{v}, \tilde{v}, \lambda, \mu) &= E(v^{(t)}, \hat{v}, \tilde{v}) + \boldsymbol{\lambda}_1^T (\mathbf{v}_1^{(t)} - \hat{\mathbf{v}}_1 - \tilde{\mathbf{v}}_1) + \boldsymbol{\lambda}_2^T (\mathbf{v}_2^{(t)} - \hat{\mathbf{v}}_2 - \tilde{\mathbf{v}}_2) \\ &\quad + \boldsymbol{\mu}_1^T (\boldsymbol{\epsilon} - \hat{\mathbf{v}}_2) + \boldsymbol{\mu}_2^T (\boldsymbol{\epsilon} - \tilde{\mathbf{v}}_2), \end{aligned} \quad (16)$$

where  $\lambda = \{\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2\}$ ,  $\mu = \{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2\}$  and  $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \boldsymbol{\mu}_1$  and  $\boldsymbol{\mu}_2$  are  $d$ -dimensional vectors of Lagrange multipliers that satisfy  $\boldsymbol{\mu}_1 \preceq \mathbf{0}$  and  $\boldsymbol{\mu}_2 \preceq \mathbf{0}$ . The corresponding dual function is

$$g(v^{(t)}, \lambda, \mu) = \max_{\hat{v}, \tilde{v}} \mathcal{L}(v^{(t)}, \hat{v}, \tilde{v}, \lambda, \mu) \quad (17)$$

The dual problem minimizes this latter function with respect to the Lagrange multipliers:

$$g(v^{(t)}) = \min_{\lambda, \mu} g(v^{(t)}, \lambda, \mu) \quad (18)$$

s.t.  $\boldsymbol{\mu}_1 \preceq \mathbf{0}, \boldsymbol{\mu}_2 \preceq \mathbf{0}$

Let  $\boldsymbol{\lambda}_1^*, \boldsymbol{\lambda}_2^*, \boldsymbol{\mu}_1^*$  and  $\boldsymbol{\mu}_2^*$  be the components of  $\lambda$  and  $\mu$  that minimize the objective in (18) under constraints  $\boldsymbol{\mu}_1 \preceq \mathbf{0}$  and  $\boldsymbol{\mu}_2 \preceq \mathbf{0}$ . The second step of the convergent algorithm solves the convex problem

$$\begin{aligned} \min_v \quad & \boldsymbol{\lambda}_1^{*T} \mathbf{v}_1 + \boldsymbol{\lambda}_2^{*T} \mathbf{v}_2 + \log \tilde{Z}(v). \\ \text{s.t.} \quad & \mathbf{v}_2 \succeq 3\boldsymbol{\epsilon} \end{aligned} \quad (19)$$

Let  $v^{(t+1)} = \{\mathbf{v}_1^{(t+1)}, \mathbf{v}_2^{(t+1)}\}$  be the solution to this optimization problem. When the constraint on  $\mathbf{v}_2$  is not tight, we have that  $\mathbf{v}_2^{(t+1)} = (2\boldsymbol{\lambda}_2^* - \boldsymbol{\lambda}_1^* \circ \boldsymbol{\lambda}_1^*)^{-1}$  and  $\mathbf{v}_1^{(t+1)} = -\boldsymbol{\lambda}_1^* \circ \mathbf{v}_2^{(t+1)}$ , where the operator " $\circ$ " denotes the Hadamard element-wise product and the inverse of a vector is defined as a new vector whose components are the inverse of the components of the original vector. When the constraint is tight, the solution is still given by these formulas. However, the components of  $\mathbf{v}_2$  which result to be smaller than  $3\boldsymbol{\epsilon}$  (those for which the constraint is active) must now be equal to  $3\boldsymbol{\epsilon}$ .

We now prove that the two steps of PC-EP, that is, (15) and (19), always generate a reduction in the value of  $E(v^{(t)}, \hat{v}, \tilde{v})$ . The proof shown here is similar to the one given in [3] for the unconstrained case. The objective in (15) is concave and the inequality constraints are affine. Thus, the weak form of Slater's condition is satisfied and strong duality holds. This means that the gap between the dual and the primal problems is zero, that is,  $g(v^{(t)}) = E(v^{(t)})$ . See [11] for further references. Let  $\lambda^* = \{\boldsymbol{\lambda}_1^*, \boldsymbol{\lambda}_2^*\}$  and  $\mu^* = \{\boldsymbol{\mu}_1^*, \boldsymbol{\mu}_2^*\}$ . Then, we have

$$E(v^{(t)}) = \max_{\hat{v}, \tilde{v}} \mathcal{L}(v^{(t)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*) \geq \max_{\hat{v}, \tilde{v}} \mathcal{L}(v^{(t+1)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*), \quad (20)$$

where the first equality follows from strong duality and the following inequality is obtained because in the second step of PC-EP we are always minimizing and consequently,  $\mathcal{L}(v^{(t)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*)$  will never increase when we replace  $v^{(t)}$  by  $v^{(t+1)}$ . Continuing with the derivations, we obtain

$$\begin{aligned} \max_{\hat{v}, \tilde{v}} \mathcal{L}(v^{(t+1)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*) &\geq \max_{\hat{v}, \tilde{v}} \mathcal{L}(v^{(t+1)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*) \geq E(v^{(t+1)}), \\ &\text{s.t (6), (11)} \end{aligned} \quad (21)$$

where the first inequality is obtained because adding extra constraints in a maximization problem can never produce higher values of the target function. The second inequality is obtained by expanding

$\mathcal{L}(v^{(t+1)}, \hat{v}, \tilde{v}, \lambda^*, \mu^*)$  according to (16), using the fact that  $\boldsymbol{\mu}_1^T(\boldsymbol{\epsilon} - \tilde{\mathbf{v}}_2)$  and  $\boldsymbol{\mu}_2^T(\boldsymbol{\epsilon} - \hat{\mathbf{v}}_2)$  can only be positive and finally, using the definition (15).

Note that we still need the vectors  $\boldsymbol{\lambda}_1^*$  and  $\boldsymbol{\lambda}_2^*$  for the practical implementation of the second step of PC-EP. Let  $\hat{v}^* = \{\hat{v}_1^*, \hat{v}_2^*\}$  and  $\tilde{v}^* = \{\tilde{v}_1^*, \tilde{v}_2^*\}$  be the solution to (15). Then, the gradient of  $\mathcal{L}$  with respect to the elements in  $\tilde{v}$  and  $\hat{v}$  should be zero at  $\tilde{v} = \tilde{v}^*$ ,  $\hat{v} = \hat{v}^*$ ,  $\lambda = \lambda^*$  and  $\mu = \mu^*$ . This generates the equations

$$\lambda_{1i}^* = -\mathbb{E}_{\mathcal{Q}}[w_i] = -\mathbb{E}_{\mathcal{P}}[w_i], \quad \lambda_{2i}^* + \mu_{1i}^* = \frac{1}{2}\mathbb{E}_{\mathcal{P}}[w_i^2], \quad \lambda_{2i}^* + \mu_{2i}^* = \frac{1}{2}\mathbb{E}_{\mathcal{Q}}[w_i^2], \quad (22)$$

for  $i = 1, \dots, d$ , where  $\mathbb{E}_{\mathcal{Q}}$  denotes expectation with respect to the normalized version of  $\mathcal{Q}$  and  $\mathbb{E}_{\mathcal{P}}$  denotes expectation with respect to  $\mathcal{P}(w_i) \propto \exp\{\hat{v}_{1i}w_i - \frac{1}{2}\hat{v}_{2i}w_i^2\}f_i(w_i)$ . The inequalities  $\tilde{\mathbf{v}}_2 \succeq \boldsymbol{\epsilon}$  and  $\hat{\mathbf{v}}_2 \succeq \boldsymbol{\epsilon}$  can only be active on either  $\hat{v}_{2i}^*$  or  $\tilde{v}_{2i}^*$ , but not on both at the same time. The reason for this is that  $\hat{\mathbf{v}}_2$  and  $\tilde{\mathbf{v}}_2$  must satisfy  $\hat{\mathbf{v}}_2 + \tilde{\mathbf{v}}_2 = \mathbf{v}_2$  and  $\mathbf{v}_2$  satisfies  $\mathbf{v}_2 \succeq 3\boldsymbol{\epsilon}$ . This means that, for each  $i$ , only either  $\mu_{1i}^*$  or  $\mu_{2i}^*$  can be different from zero, but not both at the same time. Therefore, when  $\tilde{v}_{2i}^* = \boldsymbol{\epsilon}$ , we have that  $\mu_{1i}^* = 0$  and  $\lambda_{2i}^* = 1/2\mathbb{E}_{\mathcal{P}}[w_i^2]$ . On the other hand, when  $\hat{v}_{2i}^* = \boldsymbol{\epsilon}$ , we have that  $\mu_{2i}^* = 0$  and  $\lambda_{2i}^* = 1/2\mathbb{E}_{\mathcal{Q}}[w_i^2]$ . Finally, when no inequality constraints are active in  $\hat{\mathbf{v}}_2^*$  and  $\tilde{\mathbf{v}}_2^*$ , that is,  $\boldsymbol{\mu}_1^* = \mathbf{0}$  and  $\boldsymbol{\mu}_2^* = \mathbf{0}$ , we have that  $\lambda_{2i}^* = 1/2\mathbb{E}_{\mathcal{Q}}[w_i^2] = 1/2\mathbb{E}_{\mathcal{P}}[w_i^2]$  for  $i = 1, \dots, d$ . The expectations with respect to  $\mathcal{Q}$  are given by

$$\mathbb{E}_{\mathcal{Q}}[w_i] = m_i^*, \quad \mathbb{E}_{\mathcal{Q}}[w_i^2] = (\mathbf{A}^*)_{ii}^{-1}, \quad (23)$$

for  $i = 1, \dots, d$ , where  $\mathbf{A}^* = \sigma^{-2}\mathbf{X}^T\mathbf{X} + \text{diag}(\tilde{\mathbf{v}}_2^*)$  and  $\mathbf{m}^* = (\mathbf{A}^*)^{-1}(\tilde{\mathbf{v}}_1^* + \sigma^{-2}\mathbf{X}^T\mathbf{y})$ . For computing the expectations with respect to  $\mathcal{P}$ , we define

$$p_i = -\frac{1}{2}\log(\hat{v}_{2i}^*) - \frac{1}{2}\log((\hat{v}_{2i}^*)^{-1} + v_s) + \frac{1}{2}\frac{(\hat{v}_{1i}^*)^2}{(\hat{v}_{2i}^*)^2} [\hat{v}_{2i}^* - ((\hat{v}_{2i}^*)^{-1} + v_s)^{-1}], \quad (24)$$

$$a_i = \sigma(p_i + \rho_s) \frac{\hat{v}_{1i}^*}{1 + \hat{v}_{2i}^* v_s} + \sigma(-p_i - \rho_s) \hat{v}_{1i}^*, \quad (25)$$

$$b_i = \sigma(p_i + \rho_s) \frac{(\hat{v}_{1i}^*)^2 (\hat{v}_{2i}^*)^{-2} - (\hat{v}_{2i}^*)^{-1} - v_s}{((\hat{v}_{2i}^*)^{-1} + v_s)^2} + \sigma(-p_i - \rho_s) [(\hat{v}_{1i}^*)^2 - \hat{v}_{2i}^*]. \quad (26)$$

where  $\rho_s = \log(p_s) - \log(1 - p_s)$  and  $\sigma$  denotes the logistic function. The expectations with respect to  $\mathcal{P}$  are then given by

$$\mathbb{E}_{\mathcal{P}}[w_i] = (\hat{v}_{2i}^*)^{-1}(\hat{v}_{1i}^* - a_i), \quad \mathbb{E}_{\mathcal{P}}[w_i^2] = (\hat{v}_{2i}^*)^{-1} - (\hat{v}_{2i}^*)^{-2}(a_i^2 - b_i) + \mathbb{E}_{\mathcal{P}}[w_i]^2, \quad (27)$$

for  $i = 1, \dots, d$ . See [2] for the derivation of these formulas.

The cost of PC-EP is determined by the computation of the diagonal of the covariance matrix of the posterior approximation, that is, matrix  $\mathbf{A}^{-1}$  in (4). When the number  $d$  of features is larger than the number  $n$  of training instances (that is,  $d \gg n$ ), we can use the Woodbury formula to compute  $\mathbf{A}^{-1}$  in  $\mathcal{O}(dn^2)$  operations. This is the same computational cost of regular EP [2] since this method also has to compute the diagonal of this matrix on each iteration. However, PC-EP computes the marginal variances of the posterior approximation each time that the gradient of the objective in (15) needs to be evaluated. This means that the multiplicative constant hidden in the cost of PC-EP is about two orders of magnitude larger than the one hidden in the cost of regular EP. Nevertheless, by applying the optimization algorithm described in [9] to PC-EP, this method can actually become as fast as regular EP. This is left as future work.

## 4 Experiments

The performance of regular EP (R-EP) with damping to improve convergence is compared with the proposed provably convergent algorithm (PC-EP). For this, we consider a synthetic linear regression problem with  $d = 25$  dimensions and  $\mathbf{w}$  sampled according to (1). Only 5 components of  $\mathbf{w}$  are different from zero on average ( $p_s = 0.2$ ). These are drawn from a standard Gaussian distribution ( $v_s = 1$ ). The attribute vectors  $\mathbf{x}_i$  are uniformly sampled from the unit hyper-sphere. The targets are sampled according to  $y_i = \mathbf{w}^T \mathbf{x}_i + \epsilon_i$ , where  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$  and  $\sigma = 0.005$ . Using these settings, we generate 100 training and test sets with 10 and 1000 instances, respectively.

R-EP and PC-EP are run on each training set and their prediction performance is evaluated on the corresponding test set. The method R-EP is run for a maximum of 1000 iterations using the same

Figure 1: Left, average mean square error (MSE) of R-EP and PC-EP on the test sets for which R-EP does **not** converge on the associated training sets. Right, average MSE of R-EP and PC-EP on the test sets for which R-EP does converge.

R-EP does not Converge				R-EP does Converge			
PC-EP	R-EP	$\tau$	# Sets	PC-EP	R-EP	$\tau$	# Sets
0.24±0.09	0.29±0.09	0.1	6	0.041±0.007	0.040±0.009	0.1	94
0.16±0.04	0.20±0.05	0.3	10	0.041±0.009	0.045±0.012	0.3	90
0.14±0.03	0.27±0.05	0.5	13	0.040±0.010	0.042±0.012	0.5	87
0.18±0.05	0.22±0.06	0.7	14	0.032±0.007	0.029±0.008	0.7	86
0.15±0.04	0.17±0.05	0.9	17	0.033±0.007	0.024±0.008	0.9	83

constraints (11) enforced in PC-EP. Different values are considered for the damping parameter:  $\tau \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ . The table in the left-hand side of Figure 1 shows, for each value of  $\tau$ , the average mean square error (MSE) of each method on the test sets for which R-EP does not converge on the associated training sets. The number of these test sets is also displayed in the table. The other table in the right-hand side of the figure shows the same information for those sets in which R-EP does converge. These results indicate that, when R-EP does not converge, the prediction errors of this method are typically larger than when it does converge. However, more importantly, when R-EP does not converge, the posterior approximation generated by PC-EP is better than the one obtained by R-EP. On the other hand, when R-EP does converge, the prediction errors of both methods are significantly lower and both techniques obtain similar performances.

## 5 Conclusions

The LRMSSP is interesting in under-determined scenarios with  $n \ll d$  since the spike-and-slab prior can help to alleviate over-fitting in these cases. Nevertheless, approximate inference in this model can be difficult. More precisely, expectation propagation (EP) may have convergence problems in some extreme cases and one is forced to limit the maximum number of iterations of the algorithm. In this work, we have proposed a provably convergent EP algorithm (PC-EP) for the LRMSSP. Each iteration of PC-EP is proved to minimize an energy function whose stationary points coincide with the solution of regular EP (R-EP). We also introduce constraints on the parameters of the EP approximation so that this energy function is bounded, which guarantees convergence. Experiments with synthetic data illustrate the advantages of PC-EP over R-EP. When R-EP does not converge, the posterior approximation given by PC-EP is usually better. By contrast, when R-EP does converge, both methods perform similarly.

## References

- [1] Thomas Minka. *A Family of Algorithms for approximate Bayesian Inference*. PhD thesis, MIT, 2001.
- [2] José Miguel Hernández-Lobato. *Balancing Flexibility and Robustness in Machine Learning: Semi-parametric Methods and Sparse Linear Models*. PhD thesis, Computer Science Department, Universidad Autónoma de Madrid, 2010.
- [3] M. Opper and O. Winther. Expectation consistent approximate inference. *Journal of Machine Learning Research*, 6:2177–2204, 2005.
- [4] Thomas Minka. Expectation propagation for approximate Bayesian inference. In Jack S. Breese and Daphne Koller, editors, *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence*, pages 362–36, San Francisco, CA, 2001. Morgan Kaufmann.
- [5] Edward I. George and Robert E. McCulloch. Approaches for Bayesian variable selection. *Statistica Sinica*, 7(2):339–373, 1997.
- [6] I. M. Johnstone and D. M. Titterton. Statistical challenges of high-dimensional data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4237, 2009.

- [7] T. Minka and J. Lafferty. Expectation-propagation for the generative aspect model. In Adnan Darwiche and Nir Friedman, editors, *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pages 352–359. Morgan Kaufmann, 2002.
- [8] Thomas P. Minka. Expectation propagation for approximate bayesian inference. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [9] Matthias Seeger and Hannes Nickisch. Fast convergent algorithms for expectation propagation approximate Bayesian inference. In Geoffrey Gordon, David Dunson, and Miroslav Dudk, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 (draft). JMLR W&CP, 2011.
- [10] Matthias Seeger. Expectation propagation for exponential families. Technical report, Department of EECS, University of California, Berkeley, 2006.
- [11] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, March 2004.