

Hierarchical Gaussian Process Regression

Sunho Park

*Department of Computer Science
Pohang University of Science and Technology, Korea*

TITAN@POSTECH.AC.KR

Seungjin Choi

*Department of Computer Science
Division of IT Convergence Engineering
Pohang University of Science and Technology, Korea*

SEUNGJIN@POSTECH.AC.KR

Editor: Masashi Sugiyama and Qiang Yang

Abstract

We address an approximation method for Gaussian process (GP) regression, where we approximate covariance by a block matrix such that diagonal blocks are calculated exactly while off-diagonal blocks are approximated. Partitioning input data points, we present a two-layer hierarchical model for GP regression, where prototypes of clusters in the upper layer are involved for coarse modeling by a GP and data points in each cluster in the lower layer are involved for fine modeling by an individual GP whose prior mean is given by the corresponding prototype and covariance is parameterized by data points in the partition. In this hierarchical model, integrating out latent variables in the upper layer leads to a block covariance matrix, where diagonal blocks contain similarities between data points in the same partition and off-diagonal blocks consist of approximate similarities calculated using prototypes. This particular structure of the covariance matrix divides the full GP into a pieces of manageable sub-problems whose complexity scales with the number of data points in a partition. In addition, our hierarchical GP regression (HGPR) is also useful for cases where partitions of data reveal different characteristics. Experiments on several benchmark datasets confirm the useful behavior of our method.

Keywords: Gaussian process regression, Kernel methods, Sparse approximations

1. Introduction

Gaussian process (GP) is a powerful non-parametric Bayesian method for supervised learning. It provides probabilistic predictions and allows standard Bayesian approaches to model selection. However, direct application of GP to regression problems is limited due to its unfavorable scaling: $\mathcal{O}(N^3)$ in time and $\mathcal{O}(N^2)$ in space, where N is the number of training samples. To overcome this limitation, various approximation methods have been proposed in the literature (Csató and Opper, 2002; Tresp, 2000; Williams and Seeger, 2001; Seeger et al., 2003; Smola and Bartlett, 2001; Lawrence et al., 2003; Snelson and Ghahramani, 2006; Walder et al., 2008; Lazaro-Gredilla and Figueiras-Vidal, 2009). Most of these methods fall within *sparse approximation* where a low-rank approximation is applied to the covariance matrix of the GP prior using a smaller subset of R ($\ll N$) inducing variables (Quiñonero-Candela and Rasmussen, 2005). In this case time complexity and memory space are reduced to $\mathcal{O}(R^2N)$ and $\mathcal{O}(RN)$, respectively.

The selection of input points corresponding to the inducing variables is a main issue in sparse approximation. Typical solution is to choose them among the training data by using specific criteria (Smola and Bartlett, 2001; Seeger et al., 2003; Lawrence et al., 2003). The restriction that the inducing inputs must be selected among the training data is relaxed in (Snelson and Ghahramani, 2006), where the inducing inputs are treated as auxiliary pseudo-inputs, and, are estimated in a joint continuous optimization with unknown hyperparameters. This idea can increase flexibility in fitting the approximation model to the data. Since tuning both pseudo-inputs and hyperparameters involves high-dimensional optimization, it can also lead to overfitting.

Instead of finding the inducing inputs, we directly approximate the covariance between the pairs of the latent variables based on the clustered structure in the input data. Partitioning input data points, the covariance evaluated at the pairs of training data can be exactly calculated if two points are in the same partition, otherwise it can be approximated by using the prototype vectors of the partitions. This idea leads us to propose a two-layer hierarchical model for GP regression. In the upper layer, prototype vectors of clusters are involved for coarse modeling by a GP. In the lower layer, the data points in each cluster are involved for fine modeling by an individual GP whose prior mean is given by the corresponding prototype and covariance is parameterized by data points in the partition. Integrating out latent variables in the upper layer leads to the block covariance matrix such that diagonal blocks are calculated exactly while off-diagonal blocks are approximated.

The block structure of the covariance matrix can divide the full GP into a pieces of manageable sub-problems whose complexity scales with the number of training inputs in a partition. This approach reduces both time complexity and memory space for training, i.e., if the number of training points in all partitions is roughly equal to \tilde{N} , the time complexity and memory space are reduced to $\mathcal{O}(\tilde{N}^2N)$ and $\mathcal{O}(\tilde{N}N)$, respectively. In contrast to the sparse approximation methods, our method does not need to optimize inducing points. In the case of high-dimensional input space, this fact makes our method robust to overfitting. Furthermore our method is related to the hierarchical (multilevel) regression models that have been proposed to analyze the datasets with a hierarchical or clustered structure (Raudenbush and Bryk, 2002). They can explain the variation between clusters by introducing latent variables, e.g., random intercepts or slope, that vary between clusters. Thus our method based on the hierarchical modeling can capture different characteristics among the partitions. Experimental results confirm the useful behavior of our method.

2. Gaussian Process Regression

In this section we briefly introduce GP regression. We are given a data set \mathcal{D} consisting of N training input points $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$, where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^D$, and corresponding real valued targets $\mathbf{y} = [y_1, \dots, y_N]^\top$.

GP defines a distribution over functions of the form $f : \mathcal{X} \mapsto \mathbb{R}$, which is completely described by its mean and covariance function (Rasmussen and Williams, 2006)

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}')), \quad \mathbf{x}, \mathbf{x}' \in \mathcal{X}. \quad (1)$$

Usually the mean function $m(\cdot)$ is set to a zero function, and the covariance function $\kappa(\mathbf{x}, \mathbf{x}') \triangleq \langle f(\mathbf{x}), f(\mathbf{x}') \rangle$ is modeled as a squared exponential kernel such that

$$\kappa(\mathbf{x}, \mathbf{x}') = \ell_f \exp \left\{ -\frac{1}{2\rho} (\mathbf{x} - \mathbf{x}')^\top \mathbf{L} (\mathbf{x} - \mathbf{x}') \right\}, \quad (2)$$

where $\mathbf{L} = \text{diag}(\boldsymbol{\ell})$, $[\boldsymbol{\ell}]_i^{-1}$ is a hyper-parameter to determine a relevance of the i th input dimension.

In standard GP regression, the noisy observation is modeled as the noiseless latent function added independent noise, i.e., $y = f(\mathbf{x}) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. When we evaluate the latent function at \mathbf{X} , we have a set of latent function values, $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_N)]^\top$, which follows a multivariate Gaussian distribution

$$p(\mathbf{f}|\mathbf{X}, \theta) = \mathcal{N}(\mathbf{f}|0, \mathbf{K}), \quad (3)$$

where $[\mathbf{K}]_{i,j} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$. With the observational model, the likelihood is $p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{y}|\mathbf{f}, \sigma^2 \mathbf{I})$, where \mathbf{I} is a $N \times N$ identity matrix.

The hyperparameters related to the covariance function and noise variance, $\theta = \{\boldsymbol{\ell}, \rho, \ell_f, \sigma^2\}$, can be estimated by maximizing the marginal-likelihood:

$$p(\mathbf{y}|\mathbf{X}, \theta) = \mathcal{N}(\mathbf{y}|0, \boldsymbol{\Sigma}_N), \quad (4)$$

where $\boldsymbol{\Sigma}_N \triangleq \mathbf{K} + \sigma^2 \mathbf{I}$ denotes a noisy covariance matrix. Given the observations and the estimated hyperparameters $\hat{\theta}$, the predictive distribution of the target value at a test point \mathbf{x}_* is given by

$$p(y_*|\mathbf{x}_*, \mathcal{D}, \hat{\theta}) = \mathcal{N}(y_*|\mathbf{k}^\top \boldsymbol{\Sigma}_N^{-1} \mathbf{y}, \kappa(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}^\top \boldsymbol{\Sigma}_N^{-1} \mathbf{k} + \sigma^2), \quad (5)$$

where \mathbf{k} is a N -length vector with $[\mathbf{k}]_i = \kappa(\mathbf{x}_*, \mathbf{x}_i)$.

3. Hierarchical Model for Gaussian Process Regression

This section presents the hierarchical model for GP regression (HGPR) based on the clustered structure in the input data. We assume that the input space \mathcal{X} is divided into Q partitions $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_Q$. It can be realized by clustering the training input data $\mathbf{X} = \{\{\mathbf{x}_i^{(1)}\}_{i=1}^{N_1}, \dots, \{\mathbf{x}_i^{(Q)}\}_{i=1}^{N_Q}\}$, where N_j is a number of training data in the j th cluster. The partitions of input data can be provided by the information of the dataset or by using clustering methods, e.g., k-means algorithm. Let \mathbf{c}_j be a prototype vector of the j partition, and $\mathbf{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_Q\}$ be a set of prototype vectors. When we use the k-means algorithm to partition the input data, the input space is divided into Q partitions:

$$\mathcal{X}_j = \{\mathbf{x} \in \mathcal{X} \mid d_j^2 < d_q^2, \text{ for } q \neq j, q = 1, \dots, Q\}, \quad (6)$$

where $d_j^2 = \|\mathbf{x} - \mathbf{c}_j\|^2$.

We model underlying latent functions for noisy observations in a hierarchical way: $g : \mathbf{C} \mapsto \mathbb{R}$ is a function in the upper layer (partition level), while $f_j : \mathcal{X}_j \mapsto \mathbb{R}$ is a function in

1. The bracket with subscripts denote the elements of matrices and vectors, and, colon subscript represents an entire column of the matrix.

the lower layer (data point level), defined on the j th partition. We now place a GP prior on the latent functions in such a way,

$$f_j(\mathbf{x}) | g \sim \mathcal{GP}(g(\mathbf{c}_j), \kappa_j(\mathbf{x}, \mathbf{x}')), \quad (7)$$

$$g(\mathbf{c}) \sim \mathcal{GP}(0, \kappa_g(\mathbf{c}, \mathbf{c}')), \quad (8)$$

where k_j and k_g are covariance functions defined on each partition and the set of prototypes, respectively. Especially in the upper layer $g(\mathbf{c}_j)$ is used for a constant mean function of the GP prior in the lower layer function f_j . Based on the hierarchical GP priors, the noisy observation on the j th partition is defined as

$$y_i^{(j)} = f_j(\mathbf{x}_i^{(j)}) + \epsilon_i^{(j)}, i = 1, \dots, N_j, \quad (9)$$

where $\epsilon_i^{(j)} \sim \mathcal{N}(0, \sigma_j^2)$, and σ_j^2 is a noisy variance of the j th partition.

We therefore have two levels of multivariate Gaussian distributions associated with the functions g and $\{f_j\}$. In the upper layer, let $\mathbf{g} = [g_1, \dots, g_Q]^\top$, where $g_j = g(\mathbf{c}_j)$, be a set of latent values evaluated at \mathbf{C} . The GP prior over \mathbf{g} is given by

$$p(\mathbf{g} | \mathbf{C}) = \mathcal{N}(\mathbf{g} | 0, \mathbf{K}_g), \quad (10)$$

where \mathbf{K}_g is a $Q \times Q$ matrix with $[\mathbf{K}_g]_{ij} = \kappa_g(\mathbf{c}_i, \mathbf{c}_j)$.

In the lower layer, we can define a set of latent function values evaluated at all training data \mathbf{X} . By abuse of notation, let us define $\mathbf{f} = [\mathbf{f}_1^\top, \dots, \mathbf{f}_Q^\top]^\top \in \mathbb{R}^N$, where $\mathbf{f}_j = [f_1^{(j)}, \dots, f_{N_j}^{(j)}]^\top$ ($f_i^{(j)} = f_j(\mathbf{x}_i^{(j)})$). Note that given the upper layer latent values \mathbf{g} , \mathbf{f}_j and \mathbf{f}_k , where $j \neq k$, are conditionally independent. Thus the training conditional is expressed as

$$p(\mathbf{f} | \mathbf{g}, \mathbf{X}) = \mathcal{N}(\mathbf{f} | \mathbf{H}\mathbf{g}, \mathbf{D}), \quad (11)$$

where \mathbf{H} is a $N \times Q$ binary matrix indicating the positions of the data points in the j th partition, so its j th column is defined as

$$[\mathbf{H}]_{:,j} = [0, \dots, 0, \overbrace{1, \dots, 1}^{N_j}, 0, \dots, 0]^\top, \quad (12)$$

and, \mathbf{D} is a $N \times N$ block diagonal matrix in which the j th block \mathbf{D}_j is a $N_j \times N_j$ matrix of the covariances of all pairs of the training data in j th partition, i.e.,

$$[\mathbf{D}_j]_{lm} = \kappa_j(\mathbf{x}_l^{(j)}, \mathbf{x}_m^{(j)}).$$

The likelihood of the targets corresponding to the j th partition is given by

$$p(\mathbf{y}_j | \mathbf{f}_j) = \mathcal{N}(\mathbf{y}_j | \mathbf{f}_j, \sigma_j^2 \mathbf{I}_j), \quad (13)$$

where $\mathbf{y}_j = [y_1^{(j)}, \dots, y_{N_j}^{(j)}]^\top$ and \mathbf{I}_j is a $N_j \times N_j$ identity matrix. The likelihood of all targets, i.e., $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_Q^\top]^\top$, is then expressed by

$$p(\mathbf{y} | \mathbf{f}) = \mathcal{N}(\mathbf{y} | \mathbf{f}, \mathbf{\Lambda}), \quad (14)$$

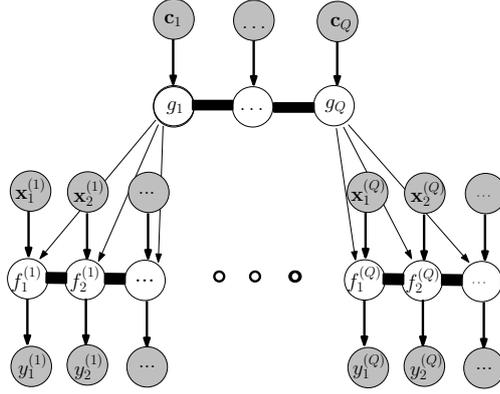


Figure 1: The graphical model for HGPR.

where $\mathbf{\Lambda} = \text{diag}(\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_Q)$, and $\boldsymbol{\lambda}_j = [\sigma_j^2, \dots, \sigma_j^2]$ is a N_j -dimensional row vector. Figure 1 describes the probabilistic graphical model for HGPR.

The hierarchical modeling provides an approximate covariance matrix, forming block matrix, for the GP prior. Integrating out the latent values in the upper level, the prior distribution over the latent values in the lower layer is given by

$$p(\mathbf{f}|\mathbf{X}, \mathbf{C}) = \mathcal{N}(\mathbf{f}|0, \mathbf{K}_f), \quad (15)$$

where $\mathbf{K}_f = \mathbf{H}\mathbf{K}_g\mathbf{H}^\top + \mathbf{D}$. We obtain a block covariance matrix in which the elements of each off-diagonal block are approximated to a constant, i.e., the covariance between the corresponding prototype vectors. Due to the block structure of the covariance matrix, the memory space for the training data is $\mathcal{O}(\sum_{j=1}^Q N_j^2 + Q^2)$, much smaller than $\mathcal{O}(N^2)$. This scheme also provides the computational benefits for inference and learning hyperparameters which will be explained later.

4. Inference and Prediction

We describe inference and prediction in our HGPR, which involve calculating posterior distributions over latent functions and predictive distributions given test data points.

4.1 Calculation of Posterior Distribution

We first derive the posterior distribution over the latent values in the lower layer, \mathbf{f} , which is calculated by using Bayes rule on (14) and (15)

$$p(\mathbf{f}|\mathbf{y}, \mathbf{X}, \mathbf{C}) = \mathcal{N}(\mathbf{f}|\boldsymbol{\Sigma}_f\boldsymbol{\Lambda}^{-1}\mathbf{y}, \boldsymbol{\Sigma}_f), \quad (16)$$

where $\boldsymbol{\Sigma}_f = \mathbf{K}_f\boldsymbol{\Sigma}_N^{-1}\boldsymbol{\Lambda}$ and $\boldsymbol{\Sigma}_N$ is a noisy covariance matrix defined by

$$\boldsymbol{\Sigma}_N = \mathbf{H}\mathbf{K}_g\mathbf{H}^\top + \tilde{\mathbf{D}}. \quad (17)$$

where $\tilde{\mathbf{D}} = \mathbf{D} + \boldsymbol{\Lambda}$ is also a block diagonal matrix.

We then calculate the posterior distribution over the latent values in the upper layer, \mathbf{g} . It provides the intermediate quantities for further computations. The likelihood of all targets given \mathbf{g} is given by

$$\begin{aligned} p(\mathbf{y}|\mathbf{g}, \mathbf{X}) &= \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{X}, \mathbf{g})d\mathbf{f} \\ &= \mathcal{N}(\mathbf{y}|\mathbf{H}\mathbf{g}, \tilde{\mathbf{D}}). \end{aligned} \quad (18)$$

As a result, the posterior distribution over \mathbf{g} is obtained by using Bayes rule on (14) and (10)

$$p(\mathbf{g}|\mathbf{y}, \mathbf{C}) = \mathcal{N}(\mathbf{g}|\tilde{\boldsymbol{\mu}}, \mathbf{A}^{-1}), \quad (19)$$

where $\tilde{\boldsymbol{\mu}} = \mathbf{A}^{-1}\mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y}$ and $\mathbf{A} = \mathbf{K}_g^{-1} + \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{H}$.

The structure of the binary matrix \mathbf{H} yields simple computational forms for $\tilde{\boldsymbol{\mu}}$ and \mathbf{A} . To show this, we define a Q dimensional vector $\boldsymbol{\mu} \triangleq \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y}$ and $Q \times Q$ diagonal matrix $\boldsymbol{\Delta} \triangleq \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{H}$. Their elements are easily computed by

$$[\boldsymbol{\mu}]_j = \sum_{i=1}^{N_j} [\boldsymbol{\alpha}_j]_i, \quad \text{and} \quad [\boldsymbol{\Delta}]_{j,j} = \sum_i^{N_j} [d_j]_i,$$

where $\boldsymbol{\alpha}_j = \tilde{\mathbf{D}}_j^{-1} \mathbf{y}_j$ and $[d_j]_i = \sum_l^{N_j} [\tilde{\mathbf{D}}_j^{-1}]_{i,l}$. In other words $[d_j]_i$ and $[\boldsymbol{\Delta}]_{j,j}$ are the summations of the i th row elements and all elements of $\tilde{\mathbf{D}}_j^{-1}$, respectively. The computations of $\tilde{\boldsymbol{\mu}}$ and \mathbf{A} are simplified as

$$\tilde{\boldsymbol{\mu}} = \mathbf{A}^{-1} \boldsymbol{\mu}, \quad \text{and} \quad \mathbf{A} = \mathbf{K}_g^{-1} + \boldsymbol{\Delta}.$$

4.2 Calculation of Predictive Distribution

We wish to find the predictive distribution of the target at a new input point $\mathbf{x}_* \in \mathcal{X}_j$, where the partition of the new point is given by the dataset or is determined by a following rule $j = \arg \min_i \|\mathbf{x}_* - \mathbf{c}_i\|^2$. We first define the likelihood conditioned on the latent variable in the lower layer, \mathbf{f} :

$$p(y_*|\mathbf{f}) = \mathcal{N}(y_* | \mathbf{k}^\top \mathbf{K}_f^{-1} \mathbf{f}, \kappa_j(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}^\top \mathbf{K}_f^{-1} \mathbf{k} + \sigma_j^2), \quad (20)$$

where \mathbf{k} is a vector of covariance between the test input point and all training input points. Here the evaluation of \mathbf{k} also involves the approximation induced by the partitioned input space assumption: only the covariances between test point and the training input points in the j th partition are exactly calculated, i.e.,

$$\mathbf{k} = \underline{\mathbf{k}} + \mathbf{H}[\mathbf{K}_g]_{:,j}, \quad (21)$$

where $\underline{\mathbf{k}} = [0, \dots, \mathbf{k}_j, \dots, 0]^\top$ and \mathbf{k}_j is a N_j -dimensional vector, $[\mathbf{k}_j]_i = \kappa_j(\mathbf{x}_*, \mathbf{x}_i^{(j)})$.

Given the new input $\mathbf{x}_* \in \mathcal{X}_j$, the predictive distribution is then obtained by integrating the likelihood (20) with respect to the posterior (16)

$$\begin{aligned} p(y_*|\mathbf{y}, \mathbf{X}, \mathbf{C}) &= \int p(y_*|\mathbf{f})p(\mathbf{f}|\mathbf{y}, \mathbf{X}, \mathbf{C})d\mathbf{f} \\ &= \mathcal{N}(y_*|\bar{f}(\mathbf{x}_*), \bar{\sigma}(\mathbf{x}_*)), \end{aligned} \quad (22)$$

where

$$\begin{aligned}\bar{f}(\mathbf{x}_*) &= \mathbf{k}^\top \boldsymbol{\Sigma}_N^{-1} \mathbf{y}, \\ \bar{\sigma}(\mathbf{x}_*) &= k_j(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}^\top \boldsymbol{\Sigma}_N^{-1} \mathbf{k} + \sigma_j^2.\end{aligned}$$

The computational bottleneck in HGPR is to calculate the inverse of the noisy covariance matrix $\boldsymbol{\Sigma}_N$. Fortunately its block structure enables us to efficiently calculate the inverse matrix based on matrix inversion lemma

$$\boldsymbol{\Sigma}_N^{-1} = \tilde{\mathbf{D}}^{-1} - \tilde{\mathbf{D}}^{-1} \mathbf{H} \mathbf{A}^{-1} \mathbf{H}^\top \tilde{\mathbf{D}}^{-1}. \quad (23)$$

Note that the inverse of $\boldsymbol{\Sigma}_N$ is given by the inverse of $\{\tilde{\mathbf{D}}_j\}_{j=1}^Q$ and \mathbf{K}_g , which involve much smaller computational effort than full GP. With (23) and some manipulations, the predictive mean and the predictive variance in (22) are rewritten by

$$\bar{f}(\mathbf{x}_*) = \mathbf{k}_j^\top \left(\boldsymbol{\alpha}_j - [\tilde{\boldsymbol{\mu}}]_j \mathbf{d}_j \right) + ([\mathbf{K}_g]_{:,j})^\top \left(\boldsymbol{\mu} - \boldsymbol{\Delta} \tilde{\boldsymbol{\mu}} \right), \quad (24)$$

$$\begin{aligned}\bar{\sigma}(\mathbf{x}_*) &= \sigma_j^2 + k_i(\mathbf{x}_*, \mathbf{x}_*) + \mathbf{k}_j^\top \tilde{\mathbf{D}}_j^{-1} \mathbf{k}_j - v_j^2 [\mathbf{A}^{-1}]_{j,j} + 2v_j [\mathbf{K}_g]_{j,j} \\ &\quad - ([\mathbf{K}_g]_{:,j})^\top \left\{ 2v_j \boldsymbol{\Delta} [\mathbf{A}^{-1}]_{:,j} - \boldsymbol{\Delta} [\mathbf{K}_g]_{:,j} + \boldsymbol{\Delta} \mathbf{A}^{-1} \boldsymbol{\Delta} [\mathbf{K}_g]_{:,j} \right\},\end{aligned} \quad (25)$$

where $v_j = \mathbf{k}_j^\top \mathbf{d}_j$. The time complexity for evaluating the predictive mean (24) and variance (25) are $\mathcal{O}(N_j)$ and $\mathcal{O}(N_j^2)$, respectively.

Note that only the training data in the partition including a test point are directly involved with the prediction, while other training data are indirectly used through the covariance function in the upper layer. This approach is useful when each partition of data shows a different property because only training points in the same partition are relevant to the prediction. If the dataset are well partitioned, HGPR naturally can capture these differences among the partitions.

5. Learning Hyperparameters

In our model assumption the covariance function for each partition in the lower layer can have its own parameterizations, i.e., we separately tune the hyperparameters for each partition. Since this flexibility can increase model complexity, we place the single squared exponential kernel in (2) on all partitions. In addition, we also assume that the noise variances of all partition are the same, that is $\sigma^2 = \sigma_1^2 = \dots = \sigma_Q^2$. Thus the hyperparameters related to the lower layer are $\theta_f = \{\ell, \rho, \ell_f, \sigma^2\}$. For the covariance function for the upper layer, we use 'rbf' kernel for \mathbf{c} and \mathbf{c}'

$$\kappa_g(\mathbf{c}, \mathbf{c}') = \ell_g \exp \left(-\frac{1}{2\ell_c} \|\mathbf{c} - \mathbf{c}'\|^2 \right). \quad (26)$$

We denote hyperparameters related to the upper layer by $\theta_g = \{\ell_g, \ell_c\}$.

The optimal hyperparameters $\theta = \{\theta_g, \theta_f\}$ are estimated by maximizing the marginal likelihood

$$\begin{aligned}p(\mathbf{y}|\mathbf{X}, \mathbf{C}, \theta) &= \int p(\mathbf{y}|\mathbf{f}) p(\mathbf{f}|\mathbf{X}, \mathbf{C}) d\mathbf{f} \\ &= \mathcal{N}(\mathbf{y}|0, \boldsymbol{\Sigma}_N).\end{aligned} \quad (27)$$

Equivalently we can find the optimal solution by minimizing the negative log-marginal likelihood

$$\begin{aligned}
 \mathcal{L} &= \frac{1}{2} \log |\boldsymbol{\Sigma}_N| + \frac{1}{2} \mathbf{y}^\top \boldsymbol{\Sigma}_N^{-1} \mathbf{y} \\
 &= \frac{1}{2} \left(\log |\tilde{\mathbf{D}}| + \log |\mathbf{K}_g| + \log |\mathbf{A}| + \mathbf{y}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y} - \mathbf{y}^\top \tilde{\mathbf{D}}^{-1} \mathbf{H} \mathbf{A}^{-1} \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y} \right) \\
 &= \frac{1}{2} \left(\sum_{j=1}^Q \log |\tilde{\mathbf{D}}_j| + \log |\mathbf{K}_g| + \log |\mathbf{A}| + \sum_{j=1}^Q \boldsymbol{\alpha}_j^\top \mathbf{y}_j - \boldsymbol{\mu}^\top \tilde{\boldsymbol{\mu}} \right), \tag{28}
 \end{aligned}$$

where an irrelevant constant is ignored.

The computational complexity of evaluating (28) is $\mathcal{O}(\sum_{j=1}^Q N_j^3)$ due to inverse of $\{\tilde{\mathbf{D}}_j\}$. If the size of each partition is almost identical to \tilde{N} , the training complexity of our model will be $\mathcal{O}(\tilde{N}^2 N)$. It has roughly equal computational cost to sparse approximation methods with the same number of inducing variables. We can optimize the objective function (28) by using the gradient based methods, e.g., quasi-Newton. The gradients are provided in Appendix.

6. Related Work

Most existing methods for the scalable GP regression have been proposed in the frame of sparse approximation. A key assumption in sparse approximation is that the latent values \mathbf{f} evaluated at the training inputs are conditionally independent with the latent values \mathbf{f}_* evaluated at the test inputs, given R ($R \leq N$) inducing variables \mathbf{u}

$$\begin{aligned}
 p(\mathbf{f}, \mathbf{f}_* | \mathbf{u}) &= p(\mathbf{f} | \mathbf{u}) p(\mathbf{f}_* | \mathbf{u}) \\
 &= \mathcal{N}(\mathbf{f} | \mathbf{K}_{f,u} \mathbf{K}_{u,u}^{-1} \mathbf{u}, \mathbf{K}_{f,f} - \mathbf{Q}_{f,f}) \mathcal{N}(\mathbf{f}_* | \mathbf{K}_{*,u} \mathbf{K}_{u,u}^{-1} \mathbf{u}, \mathbf{K}_{*,*} - \mathbf{Q}_{*,*}), \tag{29}
 \end{aligned}$$

where $p(\mathbf{f} | \mathbf{u})$ is referred as a training conditional, $p(\mathbf{f}_* | \mathbf{u})$ as a test conditional. Here, $\mathbf{K}_{f,f}$, $\mathbf{K}_{*,*}$ and $\mathbf{K}_{u,u}$ are covariance matrices for prior distributions of \mathbf{f} , \mathbf{f}_* and \mathbf{u} , respectively. Each element of these matrices is given by a single covariance function, e.g., $[\mathbf{K}_{f,f}]_{i,j} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$, and \mathbf{Q} s are low-rank matrices, e.g., $\mathbf{Q}_{f,f} = \mathbf{K}_{f,u} \mathbf{K}_{u,u}^{-1} \mathbf{K}_{f,u}^\top$, where $\mathbf{K}_{f,u}$ is a $N \times R$ matrix of covariances between \mathbf{f} and \mathbf{u} . We can apply further approximations to $p(\mathbf{f} | \mathbf{u})$ and $p(\mathbf{f}_* | \mathbf{u})$. By restring the form of covariances matrix in $p(\mathbf{f} | \mathbf{u})$, several different sparse approximation methods can be derived (Quiñonero-Candela and Rasmussen, 2005):

- deterministic training conditional (DTC) (Seeger et al., 2003)

$$q(\mathbf{f} | \mathbf{u}) = \mathcal{N}(\mathbf{f} | \mathbf{K}_{f,u} \mathbf{K}_{u,u}^{-1} \mathbf{u}, 0),$$

- fully independent training conditional (FITC) (Snelson and Ghahramani, 2006)

$$q(\mathbf{f} | \mathbf{u}) = \mathcal{N}(\mathbf{f} | \mathbf{K}_{f,u} \mathbf{K}_{u,u}^{-1} \mathbf{u}, \text{diag}(\mathbf{K} - \mathbf{Q})),$$

- partially independent training conditional (PITC)

$$q(\mathbf{f} | \mathbf{u}) = \mathcal{N}(\mathbf{f} | \mathbf{K}_{f,u} \mathbf{K}_{u,u}^{-1} \mathbf{u}, \text{block-diag}(\mathbf{K} - \mathbf{Q})).$$

Despite the similar form of the training conditionals in PITC and our HGPR (11), HGPR differs from the sparse approximation methods in the aspect that the test latent value $f(\mathbf{x}_*)$, where $\mathbf{x}_* \in \mathcal{X}_j$, directly communicates to the training latent values for j th partition, \mathbf{f}_j . This fact enables us to model the different characteristics among the partitions. In the case of sparse approximation, information from \mathbf{f} can only be transmitted to $f(\mathbf{x}_*)$ through the inducing variables (see 29). Since, however, the inducing variables do not reflect any local property, the sparse approximation methods would fail to model the local characteristics.

Our method is related to partially independent conditional (PIC) (Snelson and Ghahramani, 2007), which is a variant of PITC and directly incorporates the partitioned structure of input data into inference. PIC assumes a more relaxed assumption than (29), where the test latent value $f(\mathbf{x}_*)$, for $\mathbf{x}_* \in \mathcal{X}_j$, is grouped with the j th partition. The approximate conditional is given by

$$p(\mathbf{f}, f(\mathbf{x}_*)|\mathbf{u}) = p(\mathbf{f}_j, f(\mathbf{x}_*)|\mathbf{u}) \prod_{i \neq j}^Q p(\mathbf{f}_i|\mathbf{u}), \quad (30)$$

where \mathbf{f}_i is a set of latent values for the i th partition as described in Section 3. Here HGPR has also the same form of conditional if \mathbf{u} is replaced with \mathbf{g} . As shown in (30), the test latent value can directly communicate \mathbf{f}_j , the training latent values for the j th partition, as in our HGPR model. In this point of view, PIC also can model the local characteristics among the partitions. It is worth for considering the covariance function to understand the difference between PIC and HGPR. When $\mathbf{x} \in \mathcal{X}_i$ and $\mathbf{x}' \in \mathcal{X}_j$, the covariance function in PIC is given by

$$\text{cov}(\mathbf{x}, \mathbf{x}') = Q(\mathbf{x}, \mathbf{x}') + \delta(i, j)[k(\mathbf{x}, \mathbf{x}) - Q(\mathbf{x}, \mathbf{x})], \quad (31)$$

where $\delta(i, j)$ is a delta function such that it is equal to 1 if $i=j$, otherwise 0, and $Q(\mathbf{x}, \mathbf{x}') = \mathbf{k}_{\mathbf{x}, \mathbf{u}}^\top \mathbf{K}_u^{-1} \mathbf{k}_{\mathbf{x}, \mathbf{u}}$, where $\mathbf{k}_{\mathbf{x}, \mathbf{u}}$ is a R dimensional vector of the covariances evaluated at all pairs of \mathbf{x} and a set of inducing points that are input points corresponding to the inducing variables. In our HGPR model, the covariance function is defined by

$$\text{cov}(\mathbf{x}, \mathbf{x}') = k_g(\mathbf{c}_i, \mathbf{c}_j) + \delta(i, j)k_f(\mathbf{x}, \mathbf{x}'). \quad (32)$$

Two methods are different in the way that how the covariance between two points from different partitions is evaluated. In PIC, the covariance is approximated by the low-rank covariance function Q , while in HGPR it is approximated by the covariance between two prototypes. Since in HGPR the covariances between all pairs of input points from two different partitions are set to a constant, i.e., the covariance between two prototypes, PIC provides more accurate covariances. However PIC should optimize the inducing inputs, which involves an optimization of $R \times D$ dimensional parameters. In the case of high-dimensional input space, PIC may suffer from overfitting. In this case, we expect that HGPR gives better results. We confirm this point in Section 7.1, where our HGPR is compared with PIC in the case of high-dimensional input data.

Our method is also related to a multistrategy approach that integrates two or more inferential strategies to solve more complex problems (Michalski and Tecuci, 1994). In our case, first strategy is clustering and second strategy is inference. Unfortunately since both

strategies are completely separated, the results from the inference can not make clustering better. In that sense, our method is a naive version of mixture of GPs (Tresp, 2001; Rasmussen and Ghahramani, 2002) in which clustering and inference are done simultaneously. However due to unfavorable scaling of the mixture of GPs, our method still has attractions to handle with the large datasets.

7. Numerical Experiments

We evaluate the performance of our method, HGPR, on several large datasets, compared to four sparse approximation methods, DTC, FITC, PITC, and PIC described in Section 6. We denote R as the number of the inducing variables for these methods. To implement these sparse approximation methods, we use the GP software toolbox². Additionally we consider a full GP as a baseline method in comparison. For all cases, the covariance function is the squared exponential kernel in (2) and the hyperparameters are trained separately for each model.

This section consists of two parts: the comparison of PIC and HGPR in the case of high-dimensional input data; the performance evaluation of the methods on the large datasets. As mentioned before, PIC need to optimize the inducing inputs, which involves the high-dimensional optimization. In the experiment with the high-dimensional input data, we show that PIC suffers from overfitting. We next compare the performance of all methods on the real-world datasets, some of which have the clustered structure in the input data.

In order to evaluate the predictive performance of each method, we use a normalized mean square error (NMSE) that is defined by

$$\text{NMSE} = \langle (y^* - \bar{f}(\mathbf{x}_*))^2 \rangle / \langle (y^* - \bar{y})^2 \rangle, \quad (33)$$

where $\langle \cdot \rangle$ averages over the test data and \bar{y} is the mean value of the test targets.

7.1 Comparison with PIC

We examine the performance of PIC and our method in the case where input points are lying on high-dimensional space, and are already clustered. We generate input points from 50 number of 100-dimensional Gaussian distributions, the mean vectors of which are well separated. We draw the same number of input points from each Gaussian distribution, i.e., $Q = 50$ and $N_1 = N_2, \dots, = N_Q = \tilde{N}$. The target variable for \mathbf{x} is defined by

$$y(\mathbf{x}) = \exp \left\{ -5e^{-5} \sum_{l=1}^{20} [\mathbf{x}]_l^2 \right\} + 0.05\varepsilon \quad (34)$$

where $\varepsilon \sim \mathcal{N}(0, 1)$ and we assume that only first 20 elements of input point are relevant. The test data are generated by the same manner (1000 samples for each test data). The label of Gaussian distributions are used to partition the input points.

We conduct the experiments with \tilde{N} varying from 10 to 100, thus the number of training data N is varying from 500 to 5000. For each \tilde{N} , experiments are repeated ten times, and we report the averages. To impose the same computational complexity, we set the number

2. see <http://www.cs.manchester.ac.uk/~neill/software.html>

of inducing variables R of PIC to \tilde{N} . As mentioned in Section 6, PIC should optimize the inducing points to complete inference. In this case, the size of parameters to be estimated is $|\theta| + RD = |\theta| + 100\tilde{N}$, where $|\theta|$ is a number of the hyperparameters in PIC. This high-dimensional optimization often yields overfitting problem. On the other hands, HGPR does not need the optimization of the inducing inputs: the parameters to be estimated are only the hyperparameters related to the covariance functions, $\{\theta_g, \theta_f\}$. Although PIC provides more accurate covariance function, it may suffer from overfitting in the case of the high-dimensional input data. We can confirm this fact from Figure 2, which shows the predicative performance, in terms of NMSE, of PIC and HGPR.

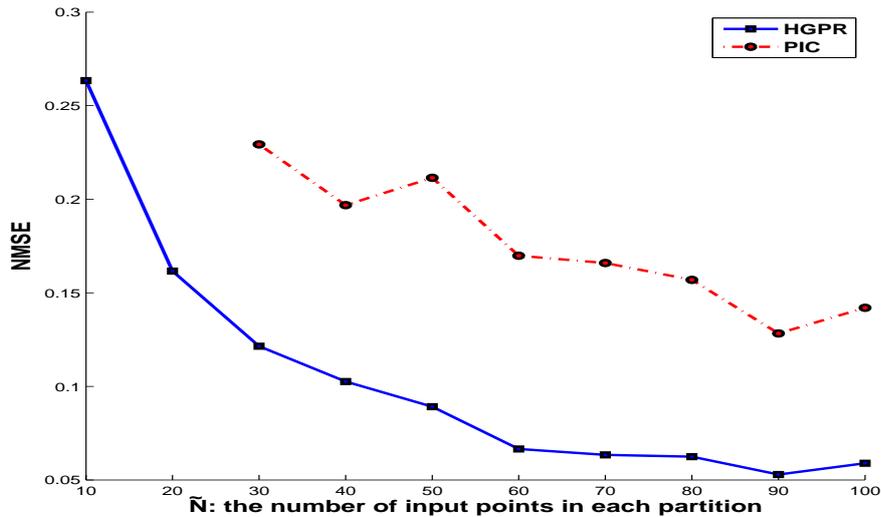


Figure 2: The predicative error (NMSE) of PIC and HGPR, when the number of input points in each partition, \tilde{N} , varies from 10 to 100, in the case where the dimension of input data is 100 and the input points are clustered. Note that, PIC fails to work at $\tilde{N} = 10, 20$.

7.2 Performance Evaluation

We use six datasets³ whose descriptions are summarized in Table 1. Kin-40k and pumadyn-32nm are artificially generated datasets describing the dynamics of a robot arm. Elevators is a prediction problem related to control of the elevators of an F16 aircraft. California housing (Cal-housing) and house-price-16H are collected from 1990 US Census, and are concerned with predicting the median price of the house in a small survey region. In contrast to Cal-housing collected in California, house-price-16H includes census records of all states in US.

3. Kin-40k: see <http://ida.first.fraunhofer.de/~anton/data.html>. Pumadyn-32nm and house-price-16H: see <http://www.cs.toronto.edu/delve/data/datasets.html>. Cal-housing: see <http://lib.stat.cmu.edu>. Elevators: see <http://www.liaad.up.pt/~ltorgo/Regression/DataSets.htm>. School-exam: see <http://multilevel.ioe.ac.uk/intro/datasets.html>.

School-exam consists of examination records from 139 secondary schools in 1985, 1986 and 1987. Actually it includes four features for each student and four features for each school. In this paper we only use the features for students as the inputs, and consider a regression problem to predict exam score for each student from these values.

For preprocessing, we normalize each input and target to have zero mean and unit variance. Only the attributes related to currency, e.g., the median price of house in Cal-housing dataset, are firstly transformed by the logarithmic function. Experiments are repeated ten times, and we report their statistics, mean and standard deviation.

Table 1: Description of data sets.

Data set	input dim.	# training(test)
<i>kin-40k</i>	8	10000(30000)
<i>pumadyn-32nm</i>	32	7168(1024)
<i>Elevators</i>	18	16599(-)
<i>Cal-housing</i>	8	20640(-)
<i>house-price-16H</i>	16	22784(-)
<i>school-exam</i>	4	15362(-)

For HGPR and PIC, we need to divide the training input data into Q partitions. The partitions of input data can be provided by the information of the dataset: in the cases of house-price-16H and school-exam datasets, the examples are clustered according to the state and the school, respectively. Otherwise we can use the k-means algorithm in which the center of each cluster becomes the prototype vector. We separately consider two artificial datasets, Kin-40k and pumadyn-32nm, from other datasets because their training input data are uniformly distributed without the clustered structure.

We thus first consider the datasets, Kin-40k and pumadyn-32nm. In both cases the size of each partition generated by k-means algorithm is almost identical, i.e., $\{N_j\}_{j=1}^Q \approx \tilde{N} = N/Q$, where N_j is the number of training data in the j th partition. We investigate the relation of the predictive accuracy of HGPR and the average number of training data in the partition, \tilde{N} . Figure 3 shows the average NMSE on both datasets when \tilde{N} is increased (\tilde{N} can be determined by adjusting Q). HGPR gives reasonable predictive accuracy when each partition has enough training points. Since the time complexity of our method is proportion to the number of points in the partition, we should carefully choose the number of partitions in terms of both efficiency and accuracy. In the case where the dataset does not have the clustered structure in the input data, we can not obtain further benefits from our method, compared to other sparse approximation methods.

We now compare the predictive performance of each method on four real-world datasets. Especially in the case of school-exam dataset, the prediction of exam score for the student can be slightly different according to the school. In other word, the partitions of dataset have different characteristics. The house-price-16H dataset shows the similar phenomena. To test whether the prediction utilizing the clustered structure in the input data improves the predictive accuracy, we consider the case where the training input data are randomly divided into $Q = 30$ partitions, and HGPR is trained on this randomly partitioned dataset. This approach is denoted as r-HGPR.

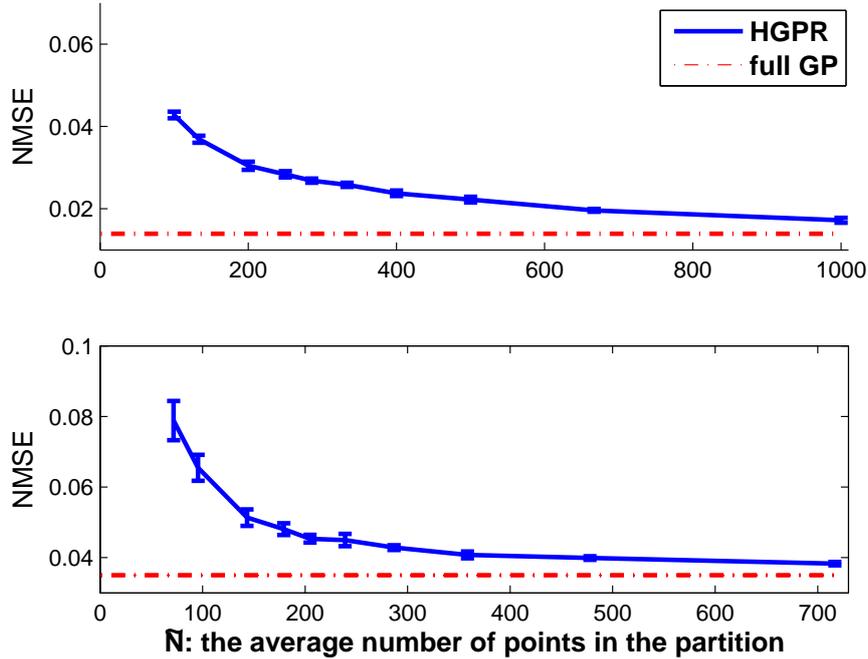


Figure 3: The predictive error (NMSE) of our method varying with respect to the average number of training data in the partitions, $\tilde{N} = N/Q$, on: (top) kin-40K dataset; (bottom) pumadyn-32mm dataset. The horizontal line is NMSE for the full GP with the hyperparameters trained on a subset of data of size 2000 for kin-40k and 1024 for pumadyn-32.

Table 2: Comparison of the predictive performance, in terms of NMSE, for all methods. The result of PIC for 'Elevators' dataset is not available, since PIC shows unstable convergence to the given dataset.

Data set	f-GP	DTC	FITC	PITC	PIC	r-HGPR	HGPR
<i>Cal-housing</i> (std.)	0.1779 (0.0099)	0.2881 (0.0388)	0.2184 (0.0082)	0.2150 (0.0043)	0.2273 (0.0046)	0.2935 (0.0037)	0.2007 (0.0051)
<i>Elevators</i> (std.)	0.0867 (0.0041)	0.1044 (0.0029)	0.1106 (0.0156)	0.1083 (0.0034)	- (-)	0.1238 (0.0042)	0.0933 (0.0023)
<i>school-exam</i> (std.)	0.6739 (0.0067)	0.6645 (0.0044)	0.7233 (0.0751)	0.6973 (0.0091)	0.6597 (0.0097)	0.7182 (0.0057)	0.6541 (0.0119)
<i>house-price-16H</i> (std.)	0.2098 (0.0063)	0.2647 (0.0103)	0.2553 (0.0071)	0.2446 (0.0042)	0.1804 (0.0046)	0.3534 (0.0048)	0.1760 (0.0029)

For Cal-housing and elevators datasets, we also use k-means algorithm. In these cases, the size of each partition can be variable and some partitions may have too small number

of points. If N_j is smaller than given threshold (200), we merged the data points in the j th partition into the nearest partition, in terms of Euclidean distance between the data point and the center of the cluster. After this procedure the number of training points in the partitions for Cal-housing less than 800, and for elevators, 600. To impose equal computational cost to sparse approximation methods, we set R to 800 for Cal-housing and 600 for elevators.

For house-price-16H and School-exam datasets, we can partition the training data based on the additional variable indicating the state or the school of the given example. Then the mean vector of each cluster is used to the prototype vector. As the number of the training inputs in the partitions for house-price-16H less than 600, and for school-exam, 200, we set R to 600 and 200 for both datasets. For other compared methods, we add this indicate variable (state/school) into the input attributes. The dataset is randomly divided into the training data ($N = 10000$) and the test data (remainder), and the hyperparameters for full GP are trained on a subset of data of size 2000.

Table 2 presents the average NMSE of each method. For all cases our method shows good predictive performance. Especially as we claimed, our method gives best performance in the cases of house-price-16H and School-exam datasets, even the predictive error of our method is lower than those of full GP. We also see that our method always outperforms rand-HGPR. This results demonstrate that utilizing the clustered structure in the input data is useful to improve the prediction performance for the dataset which reveals the different characteristics among the partitions.

8. Conclusions

We have presented a two-layer hierarchical model for GP regression where the covariance matrix is approximated by a block matrix such that diagonal blocks are exactly calculated while off-diagonal blocks are approximated. Partitioning input data points using a clustering algorithm, the latent variables in the upper layer define the mean functions of GP priors in the lower layer, while the latent variables associated with individual GP priors in the lower layer represent the noiseless latent function values evaluated at data points in the corresponding partition. By integrating out the latent variables in the upper layer leads to the block covariance matrix that enables us to reduce both the time complexity and the memory space. With appropriate partitioning methods, our method showed the high predictive performance in the case where the partitions of data reveals the different characteristics. In the case where the information for partitioning the input data is not provided by the dataset, the result of clustering algorithm is crucial to the performance of our method. Since the clustering is completely separated from the inference, the feedback from inference can not be used to make clustering better. This limitation will be solved by adopting the idea of mixture of GPs (Tresp, 2001; Rasmussen and Ghahramani, 2002) into our model formulation. Furthermore we will extend our approximation scheme to the classification task, in which the full GP is intractable due to its limited scalability. Approximation inference such as variational Bayesian (VB) inference or expectation propagation (EP) can be incorporated into the hierarchical GP prior model. The assumption on the partitioned input space could be more appropriate to model classification datasets because they are naturally clustered according to the class.

Acknowledgments

This work was supported by National Research Foundation (NRF) of Korea (No. 2010-0014306), NIPA Program of Software Engineering Technologies Development and Experts Education, NIPA ITRC Support Program (NIPA-2010-C1090-1031-0009), and NRF WCU Program (Project No. R31-2008-000-10100-0).

Appendix

In this section we provide the gradient of the negative log-marginal likelihood (28), in Section 5, with respect to the hyperparameters $\phi \in \{\theta_f, \theta_g\}$. Objective function (28) can be rewritten as

$$\begin{aligned} \mathcal{L} &= \frac{1}{2} \left(\log |\tilde{\mathbf{D}}| + \log |\mathbf{K}_g| + \log |\mathbf{A}| \right) \\ &+ \frac{1}{2} \text{tr} \left\{ \tilde{\mathbf{D}}^{-1} \mathbf{y} \mathbf{y}^\top \right\} - \frac{1}{2} \text{tr} \left\{ \mathbf{A}^{-1} \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y} \mathbf{y}^\top \tilde{\mathbf{D}}^{-1} \mathbf{H} \right\}, \end{aligned} \quad (35)$$

Then the gradient is calculated based on the notations for the matrix derivative in (Brookes, 2005). One can refer to the derivations in (Lawrence, 2007) which discusses the similar form of objective function. We only include the final results of the derivations.

We first define

$$\begin{aligned} \mathbf{M} &\triangleq \mathbf{A}^{-1} + \mathbf{A}^{-1} \mathbf{H}^\top \tilde{\mathbf{D}}^{-1} \mathbf{y} \mathbf{y}^\top \tilde{\mathbf{D}}^{-1} \mathbf{H} \mathbf{A}^{-1} \\ &= \mathbf{A}^{-1} + \tilde{\boldsymbol{\mu}} \tilde{\boldsymbol{\mu}}^\top. \end{aligned} \quad (36)$$

In the case of the upper layer, the derivative of (35) with respect to $\phi \in \theta_g$ is given by,

$$\frac{\partial \mathcal{L}}{\partial \phi} = \frac{1}{2} \text{tr} \left\{ \left(\mathbf{K}_g^{-1} - \mathbf{K}_g^{-1} \mathbf{M} \mathbf{K}_g^{-1} \right)^\top \frac{\partial \mathbf{K}_g}{\partial \phi} \right\}. \quad (37)$$

In the case of the lower layer, the derivative of (35) with respect to $\phi \in \theta_f$ is given by

$$\frac{\partial \mathcal{L}}{\partial \phi} = \frac{1}{2} \sum_{j=1}^Q \text{tr} \left\{ \boldsymbol{\Gamma}_j^\top \frac{\partial \tilde{\mathbf{D}}_j}{\partial \phi} \right\} \quad (38)$$

where $\boldsymbol{\Gamma}_j = \tilde{\mathbf{D}}_j^{-1} - \boldsymbol{\alpha}_j \boldsymbol{\alpha}_j^\top + \mathbf{d}_j \left(2[\tilde{\boldsymbol{\mu}}]_j \boldsymbol{\alpha}_j^\top - [\mathbf{M}]_{jj} \mathbf{d}_j^\top \right)$.

References

- M. Brookes. The matrix reference manual, 2005. [online] <http://www.ee.ic.ac.uk/hp/staff/dmb/matrix/intro.html>.
- L. Csató and M Opper. Sparse on-line Gaussian processes. *Neural Computation*, 14:641–668, 2002.
- N. Lawrence. Learning for larger datasets with the Gaussian process latent variable model. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, San Juan, Puerto Rico, 2007.

- N. Lawrence, M. Seeger, and R. Herbrich. Fast sparse Gaussian process models: The informative vector machine. In *Advances in Neural Information Processing Systems (NIPS)*, volume 15. MIT Press, 2003.
- M. Lazaro-Gredilla and A. Figueiras-Vidal. Inter-domain Gaussian processes for sparse inference using inducing features. In *Advances in Neural Information Processing Systems (NIPS)*, volume 22. MIT Press, 2009.
- R. S. Michalski and G. Tecuci. *Machine Learning: A Multistrategy Approach*. Morgan Kaufmann, San Francisco, CA, 1994.
- J. Quiñonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.
- C. E. Rasmussen and Z. Ghahramani. Infite mixtures of Gaussian process experts. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14. MIT Press, 2002.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- S. Raudenbush and A. Bryk. *Hierarchical Linear Models*. Thousand Oaks: Sage Publications, 2 edition, 2002.
- M. Seeger, C. K. I. Williams, and N. D. Lawrence. Fast forward selection to speed up sparse Gaussian process regression. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2003.
- A. J. Smola and P. Bartlett. Sparse greedy Gaussian process regression. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13. MIT Press, 2001.
- E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems (NIPS)*, volume 18. MIT Press, 2006.
- E. Snelson and Z. Ghahramani. Local and global sparse Gaussian process approximations. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, San Juan, Puerto Rico, 2007.
- V. Tresp. A Bayesian committee machine. *Neural Computation*, 12(11):2719–2741, 2000.
- V. Tresp. Mixture of Gaussian processes. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13. MIT Press, 2001.
- C. Walder, K. I. Kim, and B. Schölkopf. Sparse multiscale Gaussian process regression. In *Proceedings of the International Conference on Machine Learning (ICML)*, Helsinki, Finland, 2008.
- C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13. MIT Press, 2001.