

Multi-view Regularized Gaussian Processes

Qiuyang Liu and Shiliang Sun

Department of Computer Science and Technology, East China Normal University,
3663 North Zhongshan Road, Shanghai 200062, P. R. China
qiuyangliu2014@gmail.com
s1sun@cs.ecnu.edu.cn

Abstract. Gaussian processes (GPs) have been proven to be powerful tools in various areas of machine learning. However, there are very few applications of GPs in the scenario of multi-view learning. In this paper, we present a new GP model for multi-view learning. Unlike existing methods, it combines multiple views by regularizing marginal likelihood with the consistency among the posterior distributions of latent functions from different views. Moreover, we give a general point selection scheme for multi-view learning and improve the proposed model by this criterion. Experimental results on multiple real world data sets have verified the effectiveness of the proposed model and witnessed the performance improvement through employing this novel point selection scheme.

Keywords: Gaussian process, Multi-view learning, Posterior consistency, Co-regularization, Supervised learning

1 Introduction

Gaussian processes (GPs) [10] are flexible and popular Bayesian nonparametric tools for probabilistic modeling. Without giving concrete functional forms, they can be employed to define distributions over functions. As a kind of effective probabilistic models, they provide estimations of the uncertainty of predictions. With many convenient properties, GPs are widely used for various applications in machine learning and statistics. For instance, GPs have progressed dramatically in semi-supervised learning [9, 23, 14], active learning [8, 22], multi-task learning [19, 3], reinforcement learning [5, 11], and time series modeling [4, 21].

Standard GPs only deal with single view data. However, in real world, many data involve multi-view information, which may come from different feature extractors or different domains. For example, in web-page classification, a web-page can be described by its content and its hyperlink structure. In image classification, an image can be represented by its color, texture, shape, and so on. Therefore, recently, multi-view learning has aroused wide concern in machine learning. There are increasing number of algorithms proposed for multi-view learning, which can mainly be divided into two major categories [15]: co-training style algorithm [2, 16] and co-regularization style algorithms [20, 6]. However, GPs, as efficient and elegant methods in machine learning, have very few applications

in multi-view learning [20, 13]. Our work extends the GPs to the scenario of multi-view learning.

Existing multi-view learning methods which involve the GPs can be classified into two groups: Bayesian co-training [20] and subspace learning [13, 17]. The Bayesian co-training approach [20] is a Bayesian undirected graphical model, which pays attention to semi-supervised multi-view learning. The conditional independence between the output y of each data and latent functions f_j for each view is ensured by involving a latent function f_c [18]. On the other hand, the subspace learning methods [13, 17] use the GPs as tools to construct a latent variable model which could tackle the task of non-linear dimensional reduction. In contrast to those existing methods, inspired by the thought of co-regularization, our work focuses on directly extending the GPs to the context of the multi-view learning via the posterior consistency regularization, leading to elegant inference and optimization.

Our method models the classifier of each separated view as a Gaussian process. We optimize hyperparameters of the GPs by maximizing weighted average of marginal likelihood on each view and minimizing the discrepancy among the posterior distributions of the latent function on each view. As the Kullback-Leibler (KL) divergence [7] is a frequently used measure for describing the difference between two probability distributions, we employ it to characterize the discrepancy among the posterior distributions. Moreover, as data sets in real word are complex and may be seriously contaminated by noises, the sufficiency assumption, i.e., each view is sufficient for classification on its own, and the compatibility assumption, i.e., the target functions of all the views predict the same labels with a high probability [18], may fail in some cases. In consideration of these situations, we improve our model by a selective regularization idea, which is different from previous multi-view methods. In the experiments, we have compared the improved method with the original model to verify the effectiveness of the idea of the selective regularization on real word data sets.

The highlights of our work can be summarized as follows. First, we present a new GP model for multi-view learning, which extends the GPs to the scenario of multi-view learning by simple and elegant posterior consistency regularization. Secondly, our models automatically learn which views of the data should be trusted more when predicting class labels. Finally, we give a general point selection scheme for multi-view learning to deal with the situations where the sufficiency and compatibility assumptions fail, and propose the multi-view GPs with selective posterior consistency inspired by this criterion.

The remainder of this paper is organized as follows. Section 2 reviews the Gaussian processes. In Section 3, we present the multi-view GPs with posterior consistency (MvGP1), our first algorithm, covering the principles and detailed inference and learning in the proposed model. Moreover, we improve the MvGP1 to the multi-view GPs with selective posterior consistency (MvGP2) based on a general idea of the consistent set in Section 4. Experimental results are provided in Section 5. Finally, we conclude this paper and discuss the future work in Section 6.

2 Gaussian Processes

This section briefly reviews the Gaussian process (GP) model.

GPs are powerful tools for classification and regression. A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution [10]. The GP is widely used to describe a distribution over functions, and can be completely specified by its mean function and covariance function. Formally, suppose that the training set has N examples $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$, where $\mathbf{x}_i \in R^M$ is the i th input, and $y_i \in R$ is the corresponding label. Denote $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T$, and $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$. Following standard notations for GPs, the prior distribution for the latent functions \mathbf{f} is assumed to be Gaussian, $\mathbf{f}|\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{K})$, with a zero mean and a covariance matrix \mathbf{K} , whose element K_{ij} is determined by the covariance function $k(\mathbf{x}_i, \mathbf{x}_j)$. Diverse covariance functions can be employed in GPs. In this paper, we select a commonly used covariance function, the squared exponential kernel,

$$k(\mathbf{x}_i, \mathbf{x}_j) = s_f^2 \exp\left(-\frac{1}{2l^2} \sum_{d=1}^N (x_{id} - x_{jd})^2\right), \quad (1)$$

where s_f^2 is the signal variance, and l is the length-scale of the covariance.

The Gaussian likelihood for regression can be written as $\mathbf{y}|\mathbf{f} \sim \mathcal{N}(\mathbf{f}, \sigma^2 \mathbf{I})$, and after integrating out the hidden variables, the marginal likelihood is $\mathbf{y}|\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{K} + \sigma^2 \mathbf{I})$.

Under these settings, the posterior of the latent functions should be

$$\mathbf{f}|\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (2)$$

where $\boldsymbol{\mu} = \mathbf{K}(\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$ and $\boldsymbol{\Sigma} = \mathbf{K} - \mathbf{K}(\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}$ are the mean and the covariance of the posterior distribution, respectively.

We use Θ to denote the hyperparameters in the Gaussian process regression model, that is, $\Theta = \{s_f^2, l, \sigma^2\}$. These hyperparameters can be obtained by generalized maximum likelihood. In generalized maximum likelihood, we calculate the negative logarithmic marginal likelihood of the samples, $L(\Theta) = -\log p(\mathbf{y}|\mathbf{X}, \Theta)$, and then minimize $L(\Theta)$ with respect to Θ .

For a new point \mathbf{x}^* , the prediction is also Gaussian,

$$f_*|\mathbf{X}, \mathbf{y}, \mathbf{x}^* \sim \mathcal{N}(\bar{f}_*, \text{cov}(f_*)), \quad (3)$$

where $\bar{f}_* = \mathbf{k}_*^T [\mathbf{K} + \sigma^2 \mathbf{I}]^{-1} \mathbf{y}$, $\text{cov}(f_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T [\mathbf{K} + \sigma^2 \mathbf{I}]^{-1} \mathbf{k}_*$. Here, k is the covariance function, and \mathbf{k}_* is the vector of covariance function values between \mathbf{x}_* and the training data \mathbf{X} .

Standard GPs only handle single view data. However, collected data sets in the real world can often be represented by multiple views which may come from different feature extractors or various measurement modalities. As GPs are popular tools in machine learning, we propose to develop the GPs to multi-view learning.

3 Multi-view GPs with Posterior Consistency

In this section, we present the formulation of the multi-view GPs with posterior consistency (MvGP1) and show the corresponding inference and optimization in the proposed model. We first pay attention to two views learning tasks, and then give illustrations about the extensions to the scenario having more than two views.

3.1 Model Representation

Assume that the two views training set \mathbf{D} has N examples $\{(\mathbf{x}_i, \mathbf{z}_i, y_i)\}_{i=1}^N$, where $\mathbf{x}_i \in R^{M_1}$ is the i th input on the first view, $\mathbf{z}_i \in R^{M_2}$ is the i th input on the second view, and $y_i \in \{+1, -1\}$ is the corresponding label. Denote $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$, $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_N]^T$, and $\mathbf{y} = [y_1, \dots, y_N]^T$.

First, on account of leveraging the information in the separated single view, we simply assume that each view of data is modeled by a GP. That is, the prior distribution for the latent functions \mathbf{f}_1 on the first view and \mathbf{f}_2 on the second view are supposed to be Gaussian, i.e. $p(\mathbf{f}_1|\mathbf{X}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_1)$, and $p(\mathbf{f}_2|\mathbf{Z}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_2)$, where \mathbf{K}_1 and \mathbf{K}_2 are covariance matrixes determined by the corresponding covariance functions of two views, respectively. In our model, the covariance function is the squared exponential kernel as mentioned in (1). Although the Gaussian noise model is originally developed for regression, it has also been proved effective for classification, and its performance is typically comparable to the more complex probit and logit likelihood models used in classification problems [1]. Therefore, we use the Gaussian regression likelihood for our classification task to enjoy the elegant exact inference. The Gaussian likelihood for regression on the first view is $p(\mathbf{y}|\mathbf{f}_1) = \mathcal{N}(\mathbf{f}_1, \sigma_1^2 \mathbf{I})$, and the likelihood on the second view is $p(\mathbf{y}|\mathbf{f}_2) = \mathcal{N}(\mathbf{f}_2, \sigma_2^2 \mathbf{I})$. Secondly, we also need leverage the consistence between two views. The KL divergence [7] can measure the discrepancy between two distributions. Thus, we use the KL divergence between the posterior distributions on two views to regularize the objective function in MvGP1, enforcing the consistence between two views.

Suppose the posterior distribution of the latent function \mathbf{f}_1 on the first view is

$$p_1 = p(\mathbf{f}_1|\mathbf{X}, \mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1), \quad (4)$$

and the posterior distribution of the latent function \mathbf{f}_2 on the second view is

$$p_2 = p(\mathbf{f}_2|\mathbf{Z}, \mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2). \quad (5)$$

Based on the above setting, our objective function of MvGP1 is

$$\min \mathbb{L}_1 = \min\{-[a \log p(\mathbf{y}|\mathbf{X}) + (1-a) \log p(\mathbf{y}|\mathbf{Z})] + \frac{b}{2}[KL(p_1||p_2) + KL(p_2||p_1)]\}, \quad (6)$$

where

$$\log p(\mathbf{y}|\mathbf{X}) = -\frac{1}{2}\mathbf{y}^T(\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1}\mathbf{y} - \frac{1}{2} \log |\mathbf{K}_1 + \sigma_1^2 \mathbf{I}| - \frac{N}{2} \log 2\pi \quad (7)$$

is the marginal likelihood on the first view,

$$\log p(\mathbf{y}|\mathbf{Z}) = -\frac{1}{2}\mathbf{y}^T(\mathbf{K}_2 + \sigma_2^2\mathbf{I})^{-1}\mathbf{y} - \frac{1}{2}\log|\mathbf{K}_2 + \sigma_2^2\mathbf{I}| - \frac{N}{2}\log 2\pi \quad (8)$$

is the marginal likelihood on the second view, and the KL divergences between the posterior distributions \mathbf{f}_1 and \mathbf{f}_2 are

$$\begin{aligned} KL(p_1||p_2) &= \frac{1}{2}[\log|\boldsymbol{\Sigma}_2| - \log|\boldsymbol{\Sigma}_1| + tr(\boldsymbol{\Sigma}_2^{-1}\boldsymbol{\Sigma}_1) \\ &\quad + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T\boldsymbol{\Sigma}_2^{-1}(\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) - N], \end{aligned} \quad (9)$$

and

$$\begin{aligned} KL(p_2||p_1) &= \frac{1}{2}[\log|\boldsymbol{\Sigma}_1| - \log|\boldsymbol{\Sigma}_2| + tr(\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_2) \\ &\quad + (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T\boldsymbol{\Sigma}_1^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) - N]. \end{aligned} \quad (10)$$

Since the KL divergence is not a symmetrical quantity, that is, $KL(p_1||p_2) \neq KL(p_2||p_1)$, and we have no general method to determine which one is better for measuring the discrepancy between the two posterior distribution p_1 and p_2 , we construct a symmetrical quantity based on the above KL divergences, i.e., $\frac{1}{2}[KL(p_1||p_2) + KL(p_2||p_1)]$.

Parameters $\boldsymbol{\mu}_1$, $\boldsymbol{\mu}_2$, $\boldsymbol{\Sigma}_1$, and $\boldsymbol{\Sigma}_2$ in (9) and (10) are provided as follows. The mean and covariance of the posterior distribution for the latent function \mathbf{f}_1 are $\boldsymbol{\mu}_1 = \mathbf{K}_1(\mathbf{K}_1 + \sigma_1^2\mathbf{I})^{-1}\mathbf{y}$ and $\boldsymbol{\Sigma}_1 = \mathbf{K}_1 - \mathbf{K}_1(\mathbf{K}_1 + \sigma_1^2\mathbf{I})^{-1}\mathbf{K}_1$. The mean and covariance of the posterior distribution for the latent function \mathbf{f}_2 are $\boldsymbol{\mu}_2 = \mathbf{K}_2(\mathbf{K}_2 + \sigma_2^2\mathbf{I})^{-1}\mathbf{y}$ and $\boldsymbol{\Sigma}_2 = \mathbf{K}_2 - \mathbf{K}_2(\mathbf{K}_2 + \sigma_2^2\mathbf{I})^{-1}\mathbf{K}_2$.

3.2 Inference and Optimization

In our model, we consider the hybrid prediction function $sign(a\mathbf{f}_1 + (1-a)\mathbf{f}_2)$, $a \in [0, 1]$. For a new point $\{\mathbf{x}^*, \mathbf{z}^*\}$, the prediction distribution \mathbf{f}_{1*} of the first view and the prediction distribution \mathbf{f}_{2*} of the second view are also Gaussian. The mean of \mathbf{f}_{1*} is $\mathbf{k}_{1*}^T(\mathbf{K}_1 + \sigma_1^2\mathbf{I})^{-1}\mathbf{y}$, and the covariance is $k_1(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_{1*}^T(\mathbf{K}_1 + \sigma_1^2\mathbf{I})^{-1}\mathbf{k}_{1*}$. Here, k_1 is the covariance function, and \mathbf{k}_{1*} is the vector of covariance function values between \mathbf{x}_* and the training data \mathbf{X} . The mean of \mathbf{f}_{2*} is $\mathbf{k}_{2*}^T(\mathbf{K}_2 + \sigma_2^2\mathbf{I})^{-1}\mathbf{y}$, and the covariance is $k_2(\mathbf{z}_*, \mathbf{z}_*) - \mathbf{k}_{2*}^T(\mathbf{K}_2 + \sigma_2^2\mathbf{I})^{-1}\mathbf{k}_{2*}$. Here, k_2 is the covariance function, and \mathbf{k}_{2*} is the vector of covariance function values between \mathbf{z}_* and the training data \mathbf{Z} . According to the hybrid prediction function, we give our prediction on the new data point $\{\mathbf{x}^*, \mathbf{z}^*\}$.

As GPs are nonparametric models, the related hyperparameters need to be determined. In the context of a single view, we often obtain the hyperparameters via generalized maximum likelihood. In the case of two views, we optimize the hyperparameters collaboratively by two views. Co-regularization approaches often expect the predictions for the same observation of different views to be consistent to optimize the parameters. Similarly, in the context of Bayesian learning, we want the posterior distributions of the latent functions of the same observation across different views to be consistent to obtain the hyperparameters, which can be realized via minimizing the above objective function shown in (6).

The hyperparameters in our model can be divided into two classes: the trade-off hyperparameters a , b and the GP related hyperparameters, which include the hyperparameters s_f^2 , l in the covariance functions and the noise hyperparameters σ^2 in the likelihood. We use $\Theta = \{s_{f_1}^2, l_1, \sigma_1^2, s_{f_2}^2, l_2, \sigma_2^2\}$ to denote hyperparameters in the second group, where $s_{f_1}^2, l_1, \sigma_1^2$ are hyperparameters related to the first view, and others are hyperparameters related to the second view. The GP related hyperparameters Θ are optimized by the gradient descent method. As our model has an elegant formulation, the gradients with respect to Θ also have graceful forms, and the code can be easily implemented by the existing toolbox [12]. Following the above parameter notations, the gradient w.r.t the s_{f_1} is

$$\begin{aligned}
\frac{\partial \mathbf{L}_1}{\partial s_{f_1}} = & \frac{a}{2} \left\{ -\mathbf{y}^T (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{y} + \text{tr} \left[(\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} \right] \right\} \\
& + \frac{b}{2} \text{tr} \left\{ \Sigma_2^{-1} \left[\frac{2\mathbf{K}_1}{s_{f_1}} - \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 - \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} + \right. \right. \\
& \left. \left. \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 \right] - \Sigma_1^{-1} \left[\frac{2\mathbf{K}_1}{s_{f_1}} - \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 \right. \right. \\
& \left. \left. + \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 - \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} \right] \Sigma_1^{-1} \Sigma_2 \right\} \\
& + \frac{b}{2} \left\{ \left[\frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{y} - \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{y} \right]^T \right. \\
& (\Sigma_1^{-1} + \Sigma_2^{-1}) (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \Sigma_1^{-1} \left[\frac{2\mathbf{K}_1}{s_{f_1}} - \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 \right. \\
& \left. + \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{K}_1 - \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} \right] \Sigma_1^{-1} \\
& (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\Sigma_1^{-1} + \Sigma_2^{-1}) \left[\frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{y} \right. \\
& \left. \left. - \mathbf{K}_1 (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \frac{2\mathbf{K}_1}{s_{f_1}} (\mathbf{K}_1 + \sigma_1^2 \mathbf{I})^{-1} \mathbf{y} \right] \right\} \quad (11)
\end{aligned}$$

The gradients with respect to other hyperparameters in Θ are similar with (11), and hence we omit them here. The trade-off hyperparameters a and b are obtained through grid search.

We summarize MvGP1 in Algorithm 1.

3.3 Extension to Multiple Views

In the above sections, we take two views as an example to illustrate our model MvGP1. This model can be easily extended to multiple views because of the elegant formulation. The posterior distribution of the latent function on each view is the Gaussian distribution. Moreover, the KL divergence [7] between two Gaussian distributions can be calculated analytically, i.e., for two Gaussian distributions $N_0 = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and $N_1 = \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$, we have $KL(N_0||N_1) =$

Algorithm 1 Multi-view GPs with Posterior Consistency**Input:** training data $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}_{i=1}^{N_1}$, test samples $\{\mathbf{x}_i^*, \mathbf{z}_i^*, y_i^*\}_{i=1}^{N_2}$.**Output:** accuracy acc , trade-off parameters a, b , and GP related hyperparameters Θ .

- 1: initialize Θ randomly.
- 2: **for** $k = 1$ to 10 **do**
- 3: Divide the training data $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}_{i=1}^{N_1}$ into the training set $\{\mathbf{x}_i^t, \mathbf{z}_i^t, y_i^t\}_{i=1}^{N_1^t}$ and the validation set $\{\mathbf{x}_i^v, \mathbf{z}_i^v, y_i^v\}_{i=1}^{N_1^v}$.
- 4: **for** a, b in the search grids **do**
- 5: **while** termination conditions are not satisfied **do**
- 6: Update Θ by gradient descent to minimize L_1 in (6).
- 7: **end while**
- 8: Calculate the predictions by $sign(a\mathbf{f}_1 + (1-a)\mathbf{f}_2)$ on the validation set.
- 9: Calculate the accuracy acc_v on the validation set $\{\mathbf{x}_i^v, \mathbf{z}_i^v, y_i^v\}_{i=1}^{N_1^v}$.
- 10: **if** acc_v is larger than the accuracy on the last iteration **then**
- 11: record the trade-off parameters a, b .
- 12: **end if**
- 13: **end for**
- 14: **end for**
- 15: **while** termination conditions are not satisfied **do**
- 16: Update Θ by gradient descent to minimize L_1 in (6).
- 17: **end while**
- 18: Calculate the predictions by $sign(a\mathbf{f}_1 + (1-a)\mathbf{f}_2)$ on the test samples.
- 19: Calculate the accuracy acc on the test samples $\{\mathbf{x}_i^*, \mathbf{z}_i^*, y_i^*\}_{i=1}^{N_2}$.

$\frac{1}{2}\{\log(\frac{|\Sigma_1|}{|\Sigma_0|}) + tr(\Sigma_1^{-1}\Sigma_0) + (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \Sigma_1^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0) - N\}$. Therefore, we can extend MvGP1 to multiple views by regularizing the weighted logarithm with the KL divergences between every pair of the distinct posterior distributions. Following the notations in MvGP1, given a data set which involves K views, $p(\mathbf{y}|\mathbf{X}^k)$ denotes the marginal distribution of the k th view, and p_i represents the corresponding posterior distribution of the latent function \mathbf{f}_i . The objective function in multiple views is

$$\min - \sum_{k=1}^K a_k \log p(\mathbf{y}|\mathbf{X}^k) + \sum_{i=1}^K \sum_{j>i}^K \{b_{ij}[KL(p_i||p_j) + KL(p_j||p_i)]\}. \quad (12)$$

4 Multi-view GPs with Selective Posterior Consistency

When the predictions from different views are not consistent on some data points, the multi-view sufficiency assumption and compatibility assumption can not be satisfied on the whole data set. In this context, it is not appropriate to enforce the posteriors on the whole data set across the different views as similar as possible. For instance, in some cases, the input data of some views may be largely affected by the noises, which may lead the predictions of these points in these views totally different from those in other views. In these cases, the

multi-view assumptions fail, and enforcing the predictions of all views on these data points consistent seems to be improper.

Considering the above problem, we improve MvGP1 and present multi-view GPs with selective posterior consistency (MvGP2). In this model, we modify the regularization term in the objective function to make the posterior distributions across different views on a subset of the data set other than the whole one as similar as possible. In order to find the proper subset, namely the consistent set, we first optimize the hyperparameters through MvGP1 on the training set and give the predictions for the training data on each view. Next, we select data points whose predictions on each view are all consistent and are also consistent with the true label to construct the consistent set. Finally, we optimize the hyperparameters with the analogous procedure to MvGP1 except that we regularize the posteriors only on the chosen consistent set.

In fact, the key idea of MvGP2 is to construct the consistent set and constrain the multi-view assumptions on the consistent set. We construct the consistent set by selecting the data points whose label predictions on each view are all equal and the same as true labels. Formally, given two views of the input data $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$, $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_N]^T$ and the corresponding label data $\mathbf{y} = [y_1, \dots, y_N]^T$, we find a index set $T = [t_1, t_2, \dots, t_k]$, ($k \leq N$) such that for each $t \in T$, the predictions for \mathbf{X}_t , \mathbf{Z}_t and the corresponding \mathbf{y}_t are all agreed. Then the consistent set is $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}_{i \in T}$. Let $\mathbf{X}_T, \mathbf{Z}_T$, and \mathbf{y}_T denote the corresponding data matrix of the consistent set. After constructing the consistent set, we modify the objective function to only restrict the KL divergence of the posterior distributions of the latent functions on the consistent set to be minimized,

$$\begin{aligned} \min L_1' = & \{-[a \log p(\mathbf{y}|\mathbf{X}) + (1-a) \log p(\mathbf{y}|\mathbf{Z})] \\ & + \frac{b}{2}[KL(p_1'|p_2') + KL(p_2'|p_1')]\}, \end{aligned} \quad (13)$$

where $p_1' = p(\mathbf{f}_1'|\mathbf{X}_T, \mathbf{y}_T)$ is the posterior distribution of the latent function \mathbf{f}_1' on the set $\{\mathbf{X}_T, \mathbf{y}_T\}$, and $p_2' = p(\mathbf{f}_2'|\mathbf{Z}_T, \mathbf{y}_T)$ is the posterior distribution of the latent function \mathbf{f}_2' on the set $\{\mathbf{Z}_T, \mathbf{y}_T\}$.

We summarize MvGP2 in Algorithm 2. The idea of constraining the multi-view assumption on the consistent set other than the whole data set is a novel view in multi-view learning. In the real word, data are complex and noisy. It is likely that not all the data satisfy that the predictions on different views are equal. Moreover, we have verified this selective idea through the experiments on real word data sets in the following section.

5 Experiments

In this section, we performed experiments with our proposed MvGP1 and MvGP2 on multiple real word data sets. For comparison, we use three single-view methods corresponding to GPs named GP1, GP2, GP3, which use the first view, the second view, and the combination of two views, i.e., concatenating two views

Algorithm 2 Multi-view GPs with Selective Posterior Consistency

Input: training data $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}_{i=1}^{N_1}$, test samples $\{\mathbf{x}_i^*, \mathbf{z}_i^*, y_i^*\}_{i=1}^{N_2}$.**Output:** accuracy acc , trade-off parameters a, b , GP related hyperparameters Θ and the consistent index set T .

```

1:  $T = \{\}$ , initialize  $\Theta$  randomly.
2: run MvGP1.
3: for  $i = 1$  to  $N_1$  do
4:   Calculate the predictions by  $f_{i1}, f_{i2}$  on the training data  $\mathbf{x}_i$  and  $\mathbf{z}_i$ , respectively.
5:   if  $f_{i1} = f_{i2}$  &&  $f_{i2} = y_i$  then
6:     add  $i$  into the set  $T$ 
7:   end if
8: end for
9: for  $k = 1$  to 10 do
10:  Divide the training data  $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}_{i=1}^{N_1}$  into the training set  $\{\mathbf{x}_i^t, \mathbf{z}_i^t, y_i^t\}_{i=1}^{N_1^t}$  and
    the validation set  $\{\mathbf{x}_i^v, \mathbf{z}_i^v, y_i^v\}_{i=1}^{N_1^v}$ .
11:  for  $a, b$  in the search grids do
12:    while termination conditions are not satisfied do
13:      Update  $\Theta$  by gradient descent to minimize  $L'_1$  in (13).
14:    end while
15:    Calculate the predictions by  $sign(a\mathbf{f}_1 + (1-a)\mathbf{f}_2)$  on the validation set.
16:    Calculate the accuracy  $acc_v$  on the validation set  $\{\mathbf{x}_i^v, \mathbf{z}_i^v, y_i^v\}_{i=1}^{N_1^v}$ .
17:    if  $acc_v$  is larger than the accuracy on the last iteration then
18:      record the trade-off parameters  $a, b$ .
19:    end if
20:  end for
21: end for
22: while termination conditions are not satisfied do
23:  Update  $\Theta$  by gradient descent to minimize  $L'_1$  in (13).
24: end while
25: Calculate the predictions by  $sign(a\mathbf{f}_1 + (1-a)\mathbf{f}_2)$  on the test samples.
26: Calculate the accuracy  $acc$  on the test samples  $\{\mathbf{x}_i^*, \mathbf{z}_i^*, y_i^*\}_{i=1}^{N_2}$ .

```

to construct new high-dimensional data, respectively. In addition, we also compare our algorithms with a multi-view method SVM-2K [6], which is a two-view version of SVMs and is also inspired by the thought of co-regularization.

5.1 Data Sets

Web-Page The web-page data sets have been extensively used in multi-view learning, which consist of two-view web pages collected from computer science department web sites at four universities: Cornell university, university of Washington, university of Wisconsin, and university of Texas. The two views are words occurring in a web page and words appearing in the links pointing to that page. The documents are described by 1703 words in the content view, and by 569 links between them in the cites views. We list the statistical information about

the four data sets in Table 1. The web pages are distributed over five classes: student, project, course, staff and faculty. We set the category with the greatest size to be the positive class (denoted as class 1), and all the other categories as the negative class (denoted as class 2) in each data set.

Table 1. Statistical information of four web-page data sets.

Data set	size	view size	content dimension	cite dimension	class 1 size	class 2 size
Cornell	195	2	1703	195	83	112
Washington	230	2	1703	230	107	123
Wisconsin	265	2	1703	265	122	143
Texas	187	2	1703	187	103	84

Ionosphere Downloaded from UCI,¹ the ionosphere data set is collected by a system in Goose Bay, Labrador. This system involves a phased array of 16 high-frequency antennas with a total transmitted power on the order of 6.4 kilowatts. The targets are free electrons in the ionosphere. Those showing evidence of some type of structure in the ionosphere are good radar returns, while those not showing the above phenomenon and whose signals pass through the ionosphere are bad returns. The data set consists of 351 examples in which 225 are "good" instances and 126 are "bad" instances. There is only one view in this data set, and we generate the other view via principal component analysis, resulting in two views, which have 35 and 24 dimensions, respectively.

5.2 Experimental Setting

In the experiments, we select 60% data in each data set as the training set, and the rest as the test set. Multiple values of the hyperparameters a and b in MvGP1 and MvGP2 are explored in all the experiments. Given a division of the training and test set, we use cross validation with 10 folds and 20% training set as the validation set for the selection of the hyperparameters a and b in MvGP1 and MvGP2. The considered grid ranges are $a \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$ and $b \in \{2^{-18}, 2^{-12}, 2^{-8}, 2, 2^3, 2^8\}$. Other hyperparameters in MvGP1, MvGP2 and hyperparameters in GP1, GP2, GP3 are initialized randomly. As for the kernel function in GP1, GP2, GP3, MvGP1, and MvGP2, we all use the squared exponential kernel as mentioned in (1). We repeat the experiments for all the data sets five times and record the average accuracies and the corresponding standard deviations.

We compared our proposed MvGP1, MvGP2 with multi-view method SVM-2K, and three single view methods GP1, GP2, and GP3. For SVM-2K, besides the prediction functions $sign(f_1)$ and $sign(f_2)$ from the separated views, we also consider the hybrid prediction function $sign((f_1 + f_2)/2)$.

¹ Data sets are available at <http://archive.ics.uci.edu/ml/>.

5.3 Results

We present the average accuracies and standard deviations of all the methods on the webpage data sets and ionosphere data set in Table 2.

It is clearly shown in Table 2 that our proposed methods MvGP1 and MvGP2 are superior to GP1, GP2, GP3 and SVM-2K. We can also observe that MvGP2 further improves the performance over MvGP1, which benefits from the idea of using the selective posterior regularization other than the posterior regularization on the whole data sets to ensure the consistency.

Table 2. The average accuracies and standard deviations (%) of six methods on real world data sets.

Data set	Cornell	Washington	Wisconsin	Texas	Ionosphere
GP1	80.26±14.52	67.61±14.71	72.64±16.02	56.01±6.32	84.75 ± 2.15
GP2	62.56±7.87	74.78±1.42	61.51±6.27	73.52±12.66	98.72 ± 1.37
GP3	77.95±14.66	73.04±14.89	75.85±17.15	62.35±16.98	97.87 ± 4.76
SVM-2K	73.68±5.04	74.78±4.38	75.28±5.75	75.15±9.44	99.72 ± 0.39
MVGP1	85.64 ±6.87	86.30±6.08	91.32 ±1.55	78.33 ±15.14	99.29 ±1.24
MVGP2	87.18 ±5.94	87.18 ±5.94	91.32 ±2.04	81.29 ±10.81	100 ± 0

6 Conclusion

In this paper, we have proposed MvGP1 which extends GPs to the scenario of learning with multiple views via the methods of posterior consistency regularization. This approach is very intuitive, resulting in an elegant objective function formulation. Experimental results on real-word web-page classification validate the effectiveness of the proposed MvGP1. Moreover, considering that the multi-view assumptions may not be met on all data points, we have proposed MvGP2, which constructs a consistent set and constrains the posterior consistency regularization on the consistent set other than the whole data set, leading to further improvements of the performance. In fact, the idea of constraining the multi-view assumptions on a selective consistent set other than the whole data set is general. It not only can be applied to GPs, but also can inspire other multi-view learning methods.

In the future, we will attempt to extend the proposed models to the scenario of big data, which may use Nyström methods or other approximate approaches.

Acknowledgments. The corresponding author Shiliang Sun would like to thank supports from the National Natural Science Foundation of China under Projects 61673179 and 61370175.

References

1. Ashish, K., Grauman, K., Urtasun, R., Darrell, T.: Gaussian processes for object categorization. *International Journal of Computer Vision* 88(2), 169–188 (2010)
2. Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: *Proceedings of the 11th Annual Conference on Computational Learning Theory*. pp. 92–100 (1998)
3. Bonilla, E.V., Chai, K.M.A., Williams, C.K.I.: Multi-task Gaussian process prediction. In: *Proceedings of the 21th Annual Conference on Neural Information Processing Systems*. pp. 153–160 (2007)
4. Damianou, A.C., Titsias, M.K., Lawrence, N.D.: Variational Gaussian process dynamical systems. In: *Proceedings of the 25th Annual Conference on Neural Information Processing Systems*. pp. 2510–2518 (2011)
5. Engel, Y., Mannor, S., Meir, R.: Reinforcement learning with Gaussian processes. In: *Proceedings of the 22th International Conference on Machine Learning*. pp. 201–208 (2005)
6. Farquhar, J.D.R., Hardoon, D.R., Meng, H., Shawe-Taylor, J., Szedmak, S.: Two view learning: SVM-2K, theory and practice. In: *Proceedings of the 19th Annual Conference on Neural Information Processing Systems*. pp. 355–362 (2005)
7. Joyce, J.M.: *Kullback-Leibler Divergence*. Springer (2011)
8. Krause, A., Guestrin, C.: Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach. In: *Proceedings of the 24th International Conference on Machine Learning*. pp. 449–456 (2007)
9. Lawrence, N.D., Jordan, M.I.: Semi-supervised learning via Gaussian processes. In: *Proceedings of the 18th Annual Conference on Neural Information Processing Systems*. pp. 753–760 (2004)
10. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT Press (2006)
11. Rasmussen, C.E., Kuss, M.: Gaussian processes in reinforcement learning. In: *Proceedings of the 17th Annual Conference on Neural Information Processing Systems*. pp. 751–758 (2003)
12. Rasmussen, C.E., Nickisch, H.: Gaussian processes for machine learning (GPML) toolbox. *Journal of Machine Learning Research* 11(Nov), 3011–3015 (2010)
13. Shon, A.P., Grochow, K., Hertzmann, A., Rao, R.P.N.: Learning shared latent structure for image synthesis and robotic imitation. In: *Proceedings of the 19th Annual Conference on Neural Information Processing Systems*. pp. 1233–1240 (2005)
14. Sindhwani, V., Chu, W., Keerthi, S.S.: Semi-supervised Gaussian process classifiers. In: *Proceedings of the 20th International Joint Conference on Artificial Intelligence*. pp. 1059–1064 (2007)
15. Sun, S.: A survey of multi-view machine learning. *Neural Computing and Applications* 23(7-8), 2031–2038 (2013)
16. Sun, S., Jin, F.: Robust co-training. *International Journal of Pattern Recognition and Artificial Intelligence* 25(07), 1113–1126 (2011)
17. Xu, C., Tao, D., Li, Y., Xu, C.: Large-margin multi-view Gaussian process for image classification. In: *Proceedings of the 5th International Conference on Internet Multimedia Computing and Service*. pp. 7–12. ACM (2013)
18. Xu, C., Tao, D., Xu, C.: A survey on multi-view learning. *arXiv preprint arXiv:1304.5634* (2013)
19. Yu, K., Tresp, V., Schwaighofer, A.: Learning Gaussian processes from multiple tasks. In: *Proceedings of the 22th International Conference on Machine Learning*. pp. 1012–1019 (2005)

20. Yu, S., Krishnapuram, B., Rosales, R., Rao, R.B.: Bayesian co-training. *Journal of Machine Learning Research* 12, 2649–2680 (2011)
21. Zhao, J., Sun, S.: Variational dependent multi-output Gaussian process dynamical systems. *Journal of Machine Learning Research* 17, 1–36 (2016)
22. Zhou, J., Sun, S.: Gaussian process versus margin sampling active learning. *Neurocomputing* 167, 122–131 (2015)
23. Zhu, X., Ghahramani, Z., Lafferty, J.D.: Semi-supervised learning using Gaussian fields and harmonic functions. In: *Proceedings of the 20th International Conference on Machine Learning*. pp. 912–919 (2003)