



NSF-Based Mixture of Gaussian Processes and Its Variational EM Algorithm

Xiangyang Guo, Daqing Wu, T guoxy,wudq,paul.ht pku.edu.cn, jwma math.pku

. Mi e f Ga ia ce e (MGP) i a e f de
g i h da a i h i dai . H e e , i di ib i
a ce (GP) c e f a MGP a e de igged be
, hich ca de c e di ib i ha f e
da a e . I ha bee e ha e a i e (NSF)
a a f a i e ba e di ib i a a c e a ge
h gh he change f a i a e f a . I hi a e , e
NSF-ba ed i e de f Ga ia ce e (NMGP),
ed he c e i a MGP b ig di ib i de ed
f he i a i a e i e ad f Ga ia . I addi i , a
EM a g i h i e a bi hed e i a e he a a e e f a
i de a ed b he e e i a e ha ed
de e f c a ic MGP .

ls: Mi e f Ga ia ce e

a g i h · Va ia i a i f e e ce · Ne a i e · EM

1 Introduction

Gaussian process (GP) has proven to be an effective tool for a wide variety of applications in machine learning, such as modeling the inverse dynamics of a robot arm [1]. Unfortunately, GPs fail to handle multimodal datasets and, given N training samples, training a GP suffers from an expensive time complexity scaling as $O(N^3)$. To overcome these two limitations, Tresp [2] proposed the mixture of Gaussian processes (MGP), in which more than one GP is combined through a gating network. The gating network of an MGP is set to be a Gaussian mixture model (GMM) [2–5], in which each Gaussian determines the input distribution of a GP component. However, Gaussians cannot model complex distributions which frequently appear in datasets. Therefore, Yuan et al. [6] suggested replacing the Gaussians with GMMs. Although GMMs can alleviate the problem, we must determine the number of Gaussians in each GMM in advance. Inappropriate numbers may lead to bad predictive accuracy. In addition, Li et al.

[7] proposed using Student’s t -distributions as the input distributions to deal with multimodal data with overlaps between different modalities. Both GMMs and Student’s t -distributions are parametric, which are not flexible enough to model a complex distribution.

One of the most important tasks in machine learning and statistics is to model an unknown probability distribution given samples drawn from that distribution, and normalizing flow models, proposed by Tabak et al. [8] and Tabak et al. [9], are widely applied to this problem [10]. Normalizing flow is an invertible differentiable function, which can transform a simple base distribution, e.g. $\mathcal{N}(\mathbf{0}, \mathbf{I})$, into any complex distribution through the change of variables formula.

In this paper, we propose a neural-spline-flow-based mixture model of Gaussian processes (NMGP), where the neural spline flows (NSF), to be introduced in Sect. 3, are a kind of normalizing flows. In an NMGP model, the input distributions are designed to be distributions generated via NSFs instead of Gaussians. A variational EM algorithm is developed to train our proposed NMGP models, and experimental results show that NMGPs not only produce smaller errors compared to the conventional MGPs but also explore the true input distributions.

The remainder of this paper is organized as follows. We briefly introduce GP and MGP models in Sect. 2. Section 3 describes the construction of NMGPs. The variational EM algorithm is described in Sect. 4. Then, we present the experimental results in Sect. 5 and Sect. 6 concludes this paper.

2 GP and MGP Models

First, a short description of GP models is given. A GP, expressed as $\{f(\mathbf{x})|\mathbf{x} \in \mathbb{X}\}$, is a collection of random variables, any finite number of which follow a multivariate Gaussian distribution [1], and it is completely specified by its mean function $m(\mathbf{x})$ and covariance function $c(\mathbf{x}, \mathbf{x}')$, where $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ and $c(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$. In practice, we usually use a noisy GP, $\{y(\mathbf{x})|\mathbf{x} \in \mathbb{X}\}$, obtained by adding i.i.d. Gaussian noise to each $f(\mathbf{x})$, i.e. $y(\mathbf{x}) = f(\mathbf{x}) + \epsilon(\mathbf{x})$ with $\epsilon(\mathbf{x}) \sim \mathcal{N}(0, r^{-1})$, where r^{-1} denotes the noise level. We write a noisy GP as $y \sim \mathcal{GP}(m(\mathbf{x}), c(\mathbf{x}, \mathbf{x}') + r^{-1}\delta(\mathbf{x}, \mathbf{x}'))$, where $\delta(\mathbf{x}, \mathbf{x}')$ is the Kronecker delta function.

In this paper, we set $m(\mathbf{x})$ to be zero for simplicity and use the squared-exponential covariance function, defined as $c(\mathbf{x}, \mathbf{x}'|\boldsymbol{\sigma}) = \sigma_0^2 \exp\{-1/2 \sum_{i=1}^d (x_i - x'_i)^2 / \sigma_i^2\}$, where $\sigma_i > 0$ and d is the dimensionality of \mathbf{x} . In such a setting, given a training set $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ from a GP, we have

$$\mathbf{y}|\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{C} + r^{-1}\mathbf{I}_N), \quad (1)$$

where $\mathbf{C}(m, n) = c(\mathbf{x}_m, \mathbf{x}_n|\boldsymbol{\sigma})$.

Next, we briefly introduce MGP models. Here, suppose that \mathcal{D} is drawn from an MGP composed of K GP components. We introduce an indicator variable τ_n for each sample (\mathbf{x}_n, y_n) . The generative process of \mathcal{D} is given by:

$$p(\tau_n = k) = \pi_k, \quad \mathbf{x}_n|\tau_n = k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (2)$$

and

$$y_n|\tau_n = k \sim \mathcal{GP}(0, c(\mathbf{x}, \mathbf{x}'|\boldsymbol{\sigma}_k) + r_k^{-1}\delta(\mathbf{x}, \mathbf{x}')). \quad (3)$$

3 NSF-Based Mixture of Gaussian Processes

Let $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ be the base variable and \mathbf{x} the target variable following the distribution we want to model. Suppose that $h(x|\boldsymbol{\theta})$ is a monotonic rational-quadratic spline [11] parameterized by $\boldsymbol{\theta}$, whose inverse and derivative are easily computed. Then a one-layer NSF¹ $\mathbf{z} = \mathbf{h}(\mathbf{x})$ is given by $z_1 = h(x_1|\boldsymbol{\theta}_1)$ and $z_i = h(x_i|\boldsymbol{\theta}_i(\mathbf{x}_{1:i-1}))$, $i = 2, \dots, d$, where $\boldsymbol{\theta}_i(\mathbf{x}_{1:i-1})$, $i = 2, \dots, d$, are neural networks taking $\mathbf{x}_{1:i-1}$ as inputs [12]. The inverse, $\mathbf{x} = \mathbf{h}^{-1}(\mathbf{z})$, is $x_1 = h^{-1}(z_1|\boldsymbol{\theta}_1)$ and $x_i = h^{-1}(z_i|\boldsymbol{\theta}_i(\mathbf{x}_{1:i-1}))$, $i = 2, \dots, d$. Recursively, a J -layer NSF is defined as $\mathbf{z} = \boldsymbol{\varphi}(\mathbf{x}) = \mathbf{h}_J(\mathbf{h}_{J-1}(\dots \mathbf{h}_1(\mathbf{x}) \dots))$. It is obtained that $p_{\mathbf{x}}(\mathbf{x}) = p_{\mathbf{z}}(\boldsymbol{\varphi}(\mathbf{x})) \prod_{j=1}^J |\det(\partial \mathbf{x}_j / \partial \mathbf{x}_{j-1})|$, where $\mathbf{x}_j = \mathbf{h}_j(\mathbf{h}_{j-1}(\dots \mathbf{h}_1(\mathbf{x}) \dots))$, $j = 1, \dots, J$, $\mathbf{x}_J = \mathbf{z}$, and $\mathbf{x}_0 = \mathbf{x}$. Computing the determinants is cheap because the Jacobians are lower triangular. Therefore, we can easily calculate the likelihood of \mathbf{x} .

In an NMGP, a base variable \mathbf{z}_k , subject to $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, and an NSF $\boldsymbol{\varphi}_k(\mathbf{x}|\boldsymbol{\omega}_k)$, where $\boldsymbol{\omega}_k$ denotes the parameters, are introduced for k th GP. Then, to build an NMGP, we only need to replace the $\mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ in Eq. (2) with the probability distribution of $\mathbf{x} = \boldsymbol{\varphi}_k^{-1}(\mathbf{z}_k|\boldsymbol{\omega}_k)$, i.e. $p(\mathbf{x}|\boldsymbol{\omega}_k) = p_{\mathbf{z}_k}(\boldsymbol{\varphi}_k(\mathbf{x}|\boldsymbol{\omega}_k)) |\det(\partial \boldsymbol{\varphi}_k(\mathbf{x}|\boldsymbol{\omega}_k) / \partial \mathbf{x})|$.

4 Variational EM Algorithm

4.1 E-Step

Similar to the variational EM algorithm for MGP models, the linear GP model [6, 13–15], an approximation of the standard GP, is used to eliminate the dependency between the outputs y_n , $n = 1, \dots, N$. To construct a linear GP, we need to choose M ($\ll N$) samples from \mathcal{D} to form a support set \mathcal{D}' . Let $\boldsymbol{\lambda}$ follow $\mathcal{N}(\mathbf{0}, \mathbf{C}_{MM}^{-1})$, where \mathbf{C}_{MM} is the $M \times M$ covariance matrix consisting of covariance functions between samples in \mathcal{D}' . Then we assume that $y_n|\boldsymbol{\lambda} \sim \mathcal{N}(c(\mathbf{x}_n, \mathcal{D}')\boldsymbol{\lambda}, r^{-1})$. It follows that $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{NM}\mathbf{C}_{MM}^{-1}\mathbf{C}_{NM}^T + r^{-1}\mathbf{I}_N)$, where \mathbf{C}_{NM} is composed of covariance functions between samples in \mathcal{D} and \mathcal{D}' . That is identical to the likelihood of a sparse Gaussian process obtained via the subset of regressors (SoR) method [16], which shows the rationality of linear models. The most important advantage of linear models is that, given $\boldsymbol{\lambda}$, y_n , $n = 1, 2, \dots, N$, are independent.

In order to develop a variational EM algorithm for NMGPs, we pick K support sets \mathcal{D}_k , $k = 1, 2, \dots, K$, and introduce corresponding variables $\boldsymbol{\lambda}_k$, $k = 1, 2, \dots, K$. Then, K linear GPs are built to take the place of K standard GPs. We give $\boldsymbol{\pi}$ a Dirichlet prior $\text{Dir}(\boldsymbol{\alpha})$ and r_k a Gamma prior $\Gamma(a, b)$. The graphical model representation for the proposed model is shown in Fig. 1.

For simplicity, we denote all the latent variables as $\boldsymbol{\Gamma} = \{\pi_k, \boldsymbol{\lambda}_k, r_k, \tau_n | k = 1, 2, \dots, K, n = 1, 2, \dots, N\}$. In the framework of linear models, the complete data log-likelihood is given by

¹ I h i a e , h a e g e i e a n f r a e i i e d .




Fig. 1. The graphical model of the generative model.

$$\ln p(\mathbf{\Gamma}, \mathbf{X}, \mathbf{y}) = \ln p(\boldsymbol{\pi}) + \sum_{k=1}^K (\ln p(\boldsymbol{\lambda}_k) + \ln p(r_k)) + \sum_{n=1}^N \ln p(\tau_n | \boldsymbol{\pi}) + \sum_{n=1}^N (\ln p(\mathbf{x}_n | \tau_n) + \ln p(y_n | \tau_n, \mathbf{x}_n, \boldsymbol{\lambda}_{\tau_n}, r_{\tau_n})). \tag{4}$$

According to variational inference based on mean field theory, we choose a variational posterior distribution of the form

$$q(\mathbf{\Gamma}) = q(\boldsymbol{\pi}) \prod_{k=1}^K \{q(\boldsymbol{\lambda}_k)q(r_k)\} \prod_{n=1}^N q(\tau_n) \tag{5}$$

to approximate the true one in the E-step. Formulas used to calculate the above variational factors iteratively are standard [15] and thus omitted here.

4.2 M-Step

In an NMGP, the whole parameters, their current values, and the current variational posterior are denoted as Θ , Θ_t , and $q_t(\mathbf{\Gamma})$, respectively. The true Q -function, $Q(\Theta | \Theta_t) = \mathbb{E}_{p(\mathbf{\Gamma} | \mathbf{X}, \mathbf{y}, \Theta_t)} \ln p(\mathbf{\Gamma}, \mathbf{X}, \mathbf{y} | \Theta)$, cannot be computed analytically. Therefore, we use $\hat{Q}(\Theta | \Theta_t) = \mathbb{E}_{q_t(\mathbf{\Gamma})} \ln p(\mathbf{\Gamma}, \mathbf{X}, \mathbf{y} | \Theta)$ as its approximation. We maximize $\hat{Q}(\Theta | \Theta_t)$ w.r.t. Θ through gradient ascent methods. The deep learning framework PyTorch is utilized so there is no need to calculate the gradient manually.

4.3 Updating Support Sets

The first step of the variational EM algorithm is to determine K support sets randomly, and then a certain number, denoted as S , of E-steps and M-steps are performed. After that, to improve the result, the K support sets are updated according to certain criteria, and E-steps and M-steps are conducted again. We repeat the process T times determined in advance.

Fixing the parameters and variational posteriors of $\Gamma \setminus \{\boldsymbol{\lambda}_k\}_{k=1}^K$, it is reasonable to assume that the best \mathcal{D}_k maximizes the value of $q(\boldsymbol{\lambda}_k)$ at the mean [6, 13–15]. That is equivalent to maximizing

$$|\mathbf{T}_k| = |\mathbf{C}_k + \mathbb{E}_q r_k \sum_{n=1}^N q_{nk} c_k(\mathbf{x}_n) c_k(\mathbf{x}_n)^T|, \tag{6}$$

where \mathbf{C}_k is the covariance matrix composed of covariance functions between data points in \mathcal{D}_k and $c_k(\mathbf{x}_n)$ is the column vector consisting of covariance functions between \mathbf{x}_n and training samples in \mathcal{D}_k . $|\mathbf{T}_k|$ can be thought of as a function of \mathcal{D}_k .

Finding the best \mathcal{D}_k from $\binom{N}{M}$ candidates is an NP problem. Therefore, a greedy algorithm [6, 13–15] is employed to find a suboptimal solution.

4.4 Predictive Distribution

Assume that \mathbf{x}^* is a new input and y^* represents its predictive output. We need to calculate the distribution $p(y^*|\mathbf{x}^*, \mathcal{D}) = \int p(y^*|\mathbf{x}^*, \boldsymbol{\Gamma})p(\boldsymbol{\Gamma}|\mathcal{D})d\boldsymbol{\Gamma}$. Similar to the calculation of Q -function, we can replace $p(\boldsymbol{\Gamma}|\mathcal{D})$ with the variational posterior $q(\boldsymbol{\Gamma})$ obtained in the last E-step. Thus, we have

$$\begin{aligned} p(y^*|\mathbf{x}^*, \mathcal{D}) &\approx \int p(y^*|\mathbf{x}^*, \boldsymbol{\Gamma})q(\boldsymbol{\Gamma})d\boldsymbol{\Gamma} \approx p(y^*|\mathbf{x}^*, \mathbb{E}_q \boldsymbol{\Gamma}) \\ &= \sum_k \frac{\mathbb{E}_q \pi_k p(\mathbf{x}^*|\boldsymbol{\omega}_k)}{\sum_i \mathbb{E}_q \pi_i p(\mathbf{x}^*|\boldsymbol{\omega}_i)} \mathcal{N}(c_k(\mathbf{x}^*)^T \mathbb{E}_q(\boldsymbol{\lambda}_k), (\mathbb{E}_q r_k)^{-1}), \end{aligned} \tag{7}$$

which is the weighted sum of K independent Gaussian distributions.

5 Experiments

In this section, experiments on three datasets are presented. In all experiments, $\boldsymbol{\alpha}$, a and b are set to be $(1, 1, \dots, 1)^T$, 0.th43.799987792(da).2999877968699646 -1.6 -.

The training set contains 300 points, in which points from three functions account for 30%, 40%, and 30%, respectively. The test set consists of 300 points evenly distributed in $[-15, 15]$. Gaussian noise, following $\mathcal{N}(0, 0.3^2)$, is added to all samples in this dataset.

We employ an NMGP consisting of three GP components to model the dataset and set $M = 30$ and $Q = 30$. The predictive curve, pointwise standard deviations, and three learned input distributions, similar to the true ones, are plotted in Fig. 2. The RMSEs of our NMGP and baseline models are shown in Table 1. The result demonstrates that our proposed NMGP not only caused smaller error but also produced more proper approximations of the true input distributions.

Table 1. RMSE h ee da a e

	S h e ic da a e	C a ga da a e	M c c e da a e
MGP (MCMC EM) [5]	0.5411	0.6011	24.2000
TMGP (Ha d-c EM) [7]	0.5322	0.5939	23.8888
TMGP (Va ia i a EM) [15]	0.5245	0.5854	21.6147
O NMGP	0.5139	0.5799	21.5936

5.2 Coal Gas Dataset

The proposed NMGP model was also applied to a coal gas dataset [7, 15], where the training set consists of 200 samples and the test set consists of 113 samples. Here, we set $K = 4$, $M = 20$ and $Q = 20$ following Guo et al. [15]. We display the dataset and the predictions in Fig. 2, and the RMSEs in Table 1.

5.3 Motorcycle Dataset

Finally, we employed the NMGP model to fit the motorcycle dataset [5, 15] composed of 133 samples. The whole dataset is used to train the model and compute the RMSEs. Figure 2 shows the dataset and the result of our proposed model, and the RMSEs are presented in Table 1.

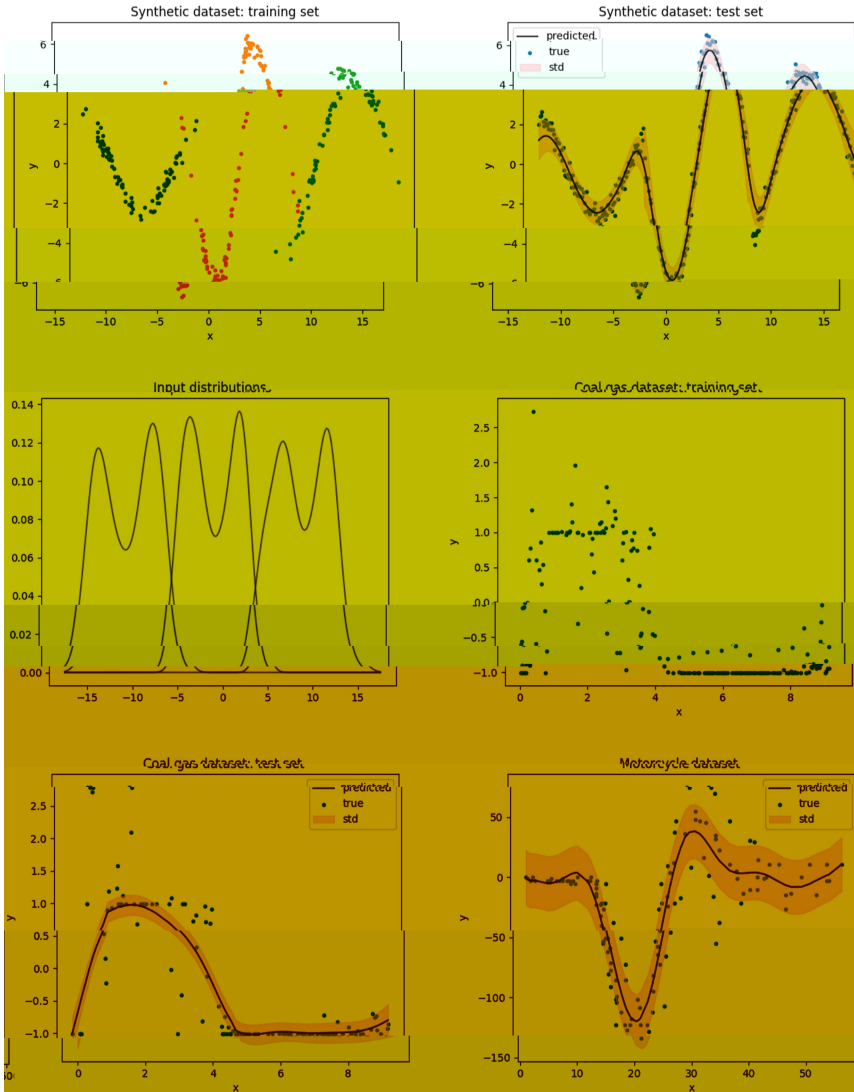


Fig. 2. The datasets

6 Conclusion

We have proposed an NSF-based mixture model of Gaussian processes and developed a variational EM algorithm for training this model. NMGPs model input distributions using NSF models instead of Gaussians. It is demonstrated by the experimental results on three datasets that our proposed NMGP models outperform MGPs. However, the variational EM algorithm cannot be applied to large datasets since the calculation of the variational posterior involves the sum

of N terms. Thus, an algorithm with lower computational complexity should be developed in the future, e.g. hard-cut EM algorithm.

Acknowledgment. This is supported by the National Key R & D Program of China (2018AAA0100205).

References

1. Rasmussen, C.E., Williams, C.K.I.: Gaussian Process for Machine Learning. MIT Press, Cambridge (2006)
2. Teller, V.: Mixed Gaussian process. In: Advances in Neural Information Processing Systems, vol. 13, pp. 654–660 (2001)
3. Yang, Y., Ma, J.: An efficient EM algorithm for learning of Gaussian process. In: Li, D., Zhang, H., Pica, M., Ai, C., He, H. (ed.) ISNN 2011, Part II. LNCS, vol. 6676, pp. 165–174. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21090-7_20
4. Chen, Z., Ma, J., Zhang, Y.: A efficient hard-cut EM algorithm for learning of Gaussian process. In: Huang, D.-S., Ji, K.-H., Wang, L. (ed.) ICIC 2014. LNCS (LNAI), vol. 8589, pp. 68–75. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09339-0_7
5. Wang, D., Chen, Z., Ma, J.: An MCMC based EM algorithm for learning of Gaussian process. In: He, X., Xia, Y., Zhang, Y., Zhang, D. (ed.) ISNN 2015. LNCS, vol. 9377, pp. 327–334. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25393-0_36
6. Yang, C., Nebel, C.: Variational learning of Gaussian process. In: Advances in Neural Information Processing Systems, vol. 21, pp. 1897–1904 (2009)
7. Li, X., Li, T., Ma, J.: The ν -hard EM algorithm for learning of Gaussian process. In: 15th IEEE International Conference on Signal Processing (ICSP), pp. 289–294 (2020)
8. Taba, E.G., Vanden-Eijnde, E.: Derivation of a closed form for the gradient of the Cramér-Rao lower bound. *Math. Sci.* **8**(1), 217–233 (2010)
9. Taba, E.G., Tane, C.V.: A fast algorithm for the derivation of the Cramér-Rao lower bound. *Math.* **66**(2), 145–164 (2013)
10. Kibbe, I., Pritchard, S.J.D., Bawa, M.A.: Non-invertible covariance matrices in the ν -hard EM algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964–3979 (2020)
11. Geman, J., DeBerg, R.: Piecewise variational approximation. *IMA J. Numer. Anal.* **2**(2), 123–130 (1982)
12. Dana, C., Bera, A., Miao, I., Papanicolaou, G.: Neural process. In: Advances in Neural Information Processing Systems, vol. 32 (2019)
13. Sun, S., Xu, X.: Variational inference for learning of Gaussian process in high-dimensional settings. *IEEE Trans. Inf. Theory* **57**(12), 466–475 (2011)
14. Li, C., Sun, S.: Variational learning of Gaussian process for causal inference. *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 4603–4609 (2017)
15. Guo, X., Li, X., Ma, J.: Variational EM algorithm for learning of Gaussian process. In: International Conference on Intelligent Computing (2021)
16. Sudderth, A.J., Bunt, P.: Saegreed Gaussian process regression. In: Advances in Neural Information Processing Systems, vol. 13, pp. 619–625 (2001)