
Transductive Log Opinion Pool of Gaussian Process Experts

Yanshuai Cao David J. Fleet
Department of Computer Science, University of Toronto

Abstract

We introduce a framework for analyzing transductive combination of Gaussian process (GP) experts for large scale problems where independently trained GP experts are combined in a way that depends on test point location. The framework provides some theoretical justification for the generalized product of GP experts (gPoE-GP) which was previously shown to work well in practice [2, 3] but lacks theoretical basis. Based on the proposed framework, an improvement over gPoE-GP is introduced and empirically validated.

1 Introduction

Gaussian processes (GP) are nonparametric Bayesian models that give probabilistic predictions required in many tasks where uncertainty is important. However, training and inference scale cubically and quadratically in the number of training data. In order to scale GP for Big Data, a number of approaches have been proposed. Most methods assume some form of conditional independence given some “inducing points”, yielding sparse GPs [5, 6, 7, 8]. While sparse GPs are much faster than the full GP, the hyperparameter optimization in these sparse GPs are still very slow in practice.

An alternative approach is to learn separate GPs on subsets of training data, and combine their predictions at test time, e.g. the Bayesian committee machine[4]. We refer to this class of methods as transductive combination of GP experts. This approach has low computational cost and are easy to parallelize. In recent works [2, 3], a new heuristic that reweights expert predictions, based on the change in entropy of expert GPs at test points, allows transductive combination of GP experts to achieve state-of-art results among large scale GP methods on a number of regression datasets. Nevertheless, there has been little theoretical justification for this new idea. In this work, we introduce a new framework for transductive combination of GP experts that explains recent empirical successes. The framework also suggests further improvement, which we validate empirically. The improved model, called diversified log opinion pool of GP experts (dLOP-GP), is highly scalable, has a more solid theoretical basis, and also improves over the previous state-of-art results.

2 Transductive combination of GP experts

The idea behind transductive combination of GPs is to train separate GP experts on subsets of data independently. Then, at test time, expert predictions are combined in a way that depends on the test point location. This scheme nicely balances the training and testing computational cost with overall model expressiveness. Learning is data-parallelizable, while for prediction, the computation nicely follows the map-reduce paradigm. Each expert computes its prediction independently during the map phase and their predictions are combined together at the reduce phase. Compared to local transductive GP, where local GPs are trained on subsets of training points selected based on test locations, transductive combination of GP experts requires no GP training at test time, hence it scales better for prediction. Furthermore, because the way expert predictions are combined depends on test location, rather than being fixed, the resulting model is much more expressive than simply taking an average of predictions. This approach can also potentially model heteroscedasticity and nonstationarity, even though individual expert GPs use relatively simple stationary kernels.

BCM The Bayesian committee machine [4] is the first well known transductive combination of GPs. Given input-output pairs, $\{(x_i, f_i)\}_i$, a GP prior $p(f)$, and a partition of training data into

disjoint subsets $\{D_i\}$, i.e. $D_j \cap D_i = \emptyset$. If we assume that when conditioned on a test point, (x_*, f_*) , the subsets of data are independent, then the posterior distribution at the test point is:

$$p(f_*|x_*, D) = \frac{1}{Z} \frac{\prod_i^K p(f_*|x_*, D_i)}{p^{K-1}(f_*|x_*)} \quad (1)$$

The resulting predictive mean and variance are: $m_{bcm}(x_*) = (\sum_i m_i(x_*)T_i(x_*)) \sigma_{bcm}^2(x_*)$ and $\sigma_{bcm}^2(x_*) = (\sum_i T_i(x_*) + (1 - M)T_{**})^{-1}$, where $T_i(x_*) = \sigma_i^{-2}(x_*)$ is the precision of the i -th Gaussian expert at point x_* , and T_{**} is the prior precision. Because BCM is derived from the conditional independence assumption and a common prior $p(f)$, using Bayes' rule, the GP on each subset of data D_i has to share the same kernel and kernel hyperparameters. Therefore, it does not naturally lend to a model where different regions of space are modelled using different kernels or hyperparameters. Beyond this lack of expressiveness, empirically, BCM has been shown to underperform by a large margin[3].

gPoE-GP The generalized product of GP experts [2] combines experts as follows:

$$p(f_*|x_*) = \frac{1}{Z} \prod_i p_i^{\alpha_i(x_*)}(f_*|x_*) \quad (2)$$

where $\alpha_i(x_*) \geq 0$ and $\sum_i \alpha_i(x_*) = 1$. Here, $\alpha_i(x_*)$ is a measure of reliability, which re-weights each expert's prediction. Taking $\alpha_i(x_*)$ to be proportional to change between prior and posterior entropy of the Gaussian distribution of the i th expert at point x_* , is simple and effective in practice [2].

The predictive mean and variance are: $m_{gpoe}(x_*) = (\sum_i m_i(x_*)\alpha_i(x_*)T_i(x_*))\sigma_{gpoe}^2(x_*)$, $\sigma_{gpoe}^2(x_*) = (\sum_i \alpha_i(x_*)T_i(x_*))^{-1}$. Although with the constraint that $\sum_i \alpha_i = 1$, the predictive distribution falls back to prior when far from training data, the overall model is ad-hoc and was previously only justified by its empirical performance [2, 3].

rBCM The robust Bayesian committee machine [3] inherits the theoretical basis from BCM, and it incorporates the reweighting by the change in entropy, like the gPoE-GP (the $\alpha_i(x_*)$). However, the BCM framework does not provide an explanation for the reweighting factor, and the rBCM is limited to all the restrictions of BCM, that is, subsets of data for the different experts need to be disjoint, and experts need to share the same kernel specification and hyperparameters.

3 Transductive log opinion pool framework

In what follows we formulate the log opinion pool of GP experts. The model is strongly motivated by the log opinion pool in [1], but adjusted for the transductive case.

Assume there are K GP experts $\{p_i\}_{i=1}^K$, each of which gives a posterior distribution $p_i(f_*|x_*)$ at the test point x_* . Also, assume we have some measure of the relative reliability of each expert at the test point, denoted by $\alpha_i(x_*)$, satisfying $\sum_i \alpha_i(x_*) = 1$ and $\alpha_i(x_*) \geq 0$. Then, we want to find a distribution $\tilde{p}(f_*|x_*)$ that is close to all experts in a weighted KL sense:

$$\tilde{p}(f_*|x_*) = \arg \min_{p(f_*|x_*)} \sum_i \alpha_i KL(p||p_i) \quad (3)$$

where the only unknowns are the $\alpha_i(x_*)$'s. Under the constraint that $\int p(f_*|x_*)df_* = 1$, the solution is exactly the gPoE in Eq. 2. Hence, this new framework provides a sound theoretical basis for gPoE, i.e. it results from model averaging in the sense of weighted KL divergence.

In [1], α_i 's are fixed for each expert, whereas they vary with location x_* here. Another major difference is that in [1], the unknown p_* can be replaced by the training inputs-outputs pairs (empirical distribution) so that α_i 's can be learned via optimization, but in the transductive setting, outputs are not observed. Hence, inference for $\alpha_i(x_*)$ needs to be done differently.

4 Selecting weights in transductive log opinion pool of GPs

We now explore alternative ways to approximately infer the weights $\alpha_i(x_*)$'s. For notational convenience, we drop the argument and write α_i directly henceforth. We need to analyze the effect of

$\alpha_i(x_*)$ on the KL divergence of $\tilde{p}(f_*|x_*)$ from the unknown ground-truth distribution $p_*(f_*|x_*)$. In [1], it was proved that:

$$KL(p_*||\tilde{p}) = \sum_i \alpha_i KL(p_*||p_i) - \sum_i \alpha_i KL(\tilde{p}||p_i) \quad (4)$$

Because \tilde{p} implicitly depends on α_i 's as well, it is not clear from Eq.4 what is the best way to select α_i 's. However, as shown [1], $KL(p_*||\tilde{p})$ can be well approximated by:

$$KL(p_*||\tilde{p}) \approx \sum_i \alpha_i KL(p_*||p_i) - \frac{1}{4} \sum_{i,j} \alpha_i \alpha_j (KL(p_i||p_j) + KL(p_j||p_i)) \equiv E - C \quad (5)$$

where E and C are defined to be the two terms in the approximation. Eq.4 still cannot be used as an objective to optimize for α_i 's numerically, because the unknown p_* cannot be approximated by samples in the transductive setting. However, Eq.4 does provide useful insight explored below:

Explaining the entropy change heuristic of gPoE-GP To reduce $KL(p_*||\tilde{p})$, one needs to decrease E and increase C . Decreasing E entails setting higher weights α_i for experts that predict well, i.e. small $KL(p_*||p_i)$, and vice versa. Again because p_* is unknown, one cannot actually compute the terms $KL(p_*||p_i)$. But we know that, without seeing any training data, the prior distribution is unlikely to be close to p_* on average. So we can decrease the weights of the experts whose prediction are not significantly influenced by training data. With Gaussian process experts, the change in entropy between prior and posterior at point α_i provides such a measure. gPoE-GP ignores the second term C in Eq.4, and attempts to decrease the first term by lowering relative weights of experts whose predictions are less based on data. When benefits of diversification is small, e.g. if subsets of points are uniformly randomly sampled and all experts share kernel hyper-parameters, ignoring C could be a reasonable approximation.

Diversified log opinion pool of GP experts Increasing the second term C in Eq.4 entails increasing weights for pairs of predictions that are further apart in the sense of symmetric KL divergence. This encourages more diversified predictions. Because C does not include p_* , it can be computed. In fact, C is a quadratic form in the weight (row) vector $\alpha = [\alpha_i]_{i=1}^K$: $C = \frac{1}{4} \alpha Q \alpha^\top$, where, Q is a $K \times K$ symmetric matrix whose ij -entry is $KL(p_i||p_j) + KL(p_j||p_i)$.

A simple way to incorporate C into the determination of $\{\alpha_i\}_{i=1}^K$ is to start with the α_i 's of gPoE-GP as initialization, and modify them in a way that increases C . Of course, this does not always guarantee to lower Eq.4, but we find empirically that simply taking a single normalized gradient step of the C term then re-normalizing $\sum_i \alpha_i$ to 1 works well. Since Eq. is already an approximate objective, and we cannot truly optimize it due to the E term, minimizing the C term more accurately is not worth the effort. Because the resulting model encourages more diverse expert predictions than gPoE-GP, we refer to it as diversified log opinion pool of GP experts (dLOG-GP). In other words, at a test point x_* , the weight vector $\alpha(x_*)$ for experts in dLOP-GP is computed as following: first, set $\alpha = \alpha^{gpoe}$, where α^{gpoe} is set to the changes in entropy of the experts, and re-normalized to sum to 1; then set $\hat{\alpha} = \alpha^{gpoe} + \lambda \frac{\nabla C}{\|\nabla C\|}$; and finally, $\alpha = \hat{\alpha} / \sum_i \hat{\alpha}_i$ where $\hat{\alpha}_i = \hat{\alpha}[i]$. We found that a range of step sizes λ all work well in practice, as well as taking more than one gradient step, so we simply use $\lambda = 1.0$ and use one update step in this work. The forms of predictive mean and variance are the same as gPoE-GP.

5 Experiment

Next we report empirical results on the dLOP-GP along with several related models. To this end we use three standard regression datasets: KIN40K (8D feature space, 10K training points, 30K test points), SARCOS (21D, 44, 484 training points, 4, 449 test points), and the UK apartment price dataset (2D, 64, 910 training points, 10K test points) used in [8].

For each model on each dataset, we explore five different ways to select subsets of data used by experts, and show that dLOP-GP generally improves over previous methods. The five subset selection approaches are: *i*) **(DS)** points are randomly partitioned into disjoint subsets, experts share the same hyperparameters learned from a random subset of data GP; *ii*) **(SoD-Shared-Hyp)** subsets of data are randomly sampled with replacement (so not necessarily disjoint), experts share the same hyperparameters as in DS; *iii*) **(SoD)** same as the previous SoD-Shared-Hyp scheme but experts have different hyperparameters that are independently learned; *iv*) **(tree)** subsets are based on a tree

partition: a ball tree [9] built on training set recursively partitions the space; and on each level of the tree, a random subset of data is drawn to build a GP; kernel hyperparameters are independently learned; ν) (**tree-rand-kern**) same as the previous tree based construction, but the kernels for experts are randomly specified as square exponential ARD kernel, Matern32, Matern52 or sum of any of these three. gPoE-GP and dLOP-GP both work with all construction schemes. For rBCM, only the DS construction is valid in theory, as the sets are disjoint with the same hyperparameters, but we are not able to replicate results from [3], especially when predictive uncertainty is taken into account in benchmarking. rBCM does not truly support the other schemes, since experts do not share the same prior anymore, but for comparison, we run it by taking the average of all experts’ prior variance, instead of a common prior variance in these other scenarios.

All subsets have 512 points (except for one of the subsets in the DS scheme, whenever the total number of data is not divisible by 512). And with the exception of *DS*, all schemes use 128 experts. The only overhead that dLOP-GP adds compared to gPoE-GP is in the one gradient step to determined α_i ’s. Empirically, this presents only a 5% to 10% increases in test computation time. In our experiments, on KIN40K under SoD scheme, training takes about 20 minutes when parallelized on 24 cores, while testing takes about 400 seconds for gPoE-GP and 420 seconds for dLOP-GP.

	gPoE	dLOP	rBCM
DS	(-0.685 , 0.266)	(-0.677, 0.270)	(1.363, 0.244)
SoD-Shared-Hyp	(-0.688 , 0.265)	(-0.683, 0.267)	(20.031, 0.254)
SoD	(-0.851, 0.193)	(-0.929 , 0.172)	(3.558, 0.172)
Tree	(-0.795, 0.216)	(-0.893 , 0.172)	(-0.125, 0.168)
Tree-Rand-Kern	(-0.912, 0.179)	(-1.040 , 0.148)	(0.830, 0.126)

(a) KIN40K

	gPoE	dLOP	rBCM
DS	(-1.480 , 0.088)	(-1.452, 0.081)	(7.712, 0.073)
SoD-Shared-Hyp	(-0.955, 0.110)	(-1.003 , 0.091)	(46.78, 1.48)
SoD	(-1.662, 0.046)	(-1.668 , 0.044)	(2.695, 0.040)
Tree	(-2.191, 0.028)	(-2.315 , 0.020)	(0.505, 0.020)
Tree-Rand-Kern	(-2.800, 0.026)	(-3.099 , 0.016)	(-1.716, 0.019)

(b) SARCOS

	gPoE	dLOP	rBCM
DS	(-0.209 , 0.002)	(-0.209 , 0.002)	(1201.099, 3.952)
SoD-Shared-Hyp	(-0.210, 0.002)	(-0.212 , 0.002)	(1219.394, 4.050)
SoD	(-0.213, 0.002)	(-0.215 , 0.002)	(1248.594, 4.047)
Tree	(-0.192 , 0.002)	(-0.171, 0.002)	(211.803, 1.094)
Tree-Rand-Kern	(-0.173 , 0.002)	(-0.042, 0.003)	(2967.578, 21.299)

(c) UK-APT

Table 1: Each tuple of result is (SNLP, SMSE). Cases where dLOP is better than or equal to gPoE are in blue.

Results are measured by standardized negative log probability (SNLP) and standardized mean square error (SMSE) in Table.1a - 1c. SNLP is a more informative metric as it takes into account the uncertainty in prediction. Insights from the theoretical frameworks are confirmed: dLOP-GP almost always improves over gPoE-GP (blue text); and as individual experts become more diverse from SoD-Shared-Hyp to SoD, both gPoE-GP and dLOP-GP improve; and on KIN40K and SARCOS, Tree-Rand-Kern, which gives the most diverse experts, yields the significantly better results. Furthermore, results also demonstrate the need for different rather than shared hyperparameters, which, only the two models under the transductive log opinion pool framework truly support.

6 Conclusion

We presented a theoretical basis for the gPoE-GP model which was previously believed to be theoretically unfounded but yields surprisingly good results in practice. Intuitions from the new framework, along with a new model, dLOP-GP, that is an improvement over gPoE-GP are validated empirically on three datasets.

References

- [1] T. Heskes. Selecting weighting factors in logarithmic opinion pools. *NIPS*, 1998.
- [2] Y. Cao, D. J. Fleet. Generalized Product of Experts for Automatic and Principled Fusion of Gaussian Process Predictions. *Modern Nonparametrics 3: Automating the Learning Pipeline workshop at NIPS 2014*, 2014.
- [3] M. P. Deisenroth, J. W. Ng. Distributed Gaussian processes. *ICML*, 2015.
- [4] V. Tresp. A Bayesian committee machine. *Neural Computation*, 2000.
- [5] E. Snelson and Z. Ghahramani. Sparse gaussian processes using pseudo-inputs. *NIPS 18*, pp. 1257–1264. 2006.
- [6] M. K. Titsias. "Variational learning of inducing variables in sparse Gaussian processes." *AISTATS*, 2009.
- [7] Y. Cao, M. Brubaker, D. J. Fleet, A. Hertzmann. Efficient Optimization for Sparse Gaussian Process Regression. *IEEE TPAMI*, 2015.
- [8] James Hensman, Nicolo Fusi, Neil D. Lawrence. Gaussian Processes for Big Data. *NIPS*, 2013.
- [9] Omohundro, Stephen M. Five Balltree Construction Algorithms International Computer Science Institute Technical Report, 1989.