

# A Data Parallel Approach for Large-Scale Gaussian Process Modeling

*Arindam Choudhury, Prasanth B. Nair, and Andy J. Keane\**

## Abstract

This paper proposes an enabling data parallel local learning methodology for handling large data regression through the Gaussian Process (GP) modeling paradigm. The proposed model achieves parallelism by employing a specialized compactly supported covariance function defined over spatially localized clusters. The associated load balancing constraints arising from data parallelism are satisfied using a novel greedy clustering algorithm, GeoClust producing balanced clusters localized in space. Further, the use of the proposed covariance function as a building block for GP models is shown to decompose the maximum likelihood estimation problem into smaller decoupled subproblems. The attendant benefits which include a significant reduction in training complexity, as well as sparse predictive models for the posterior mean and variance make the present scheme extremely attractive. Experimental investigations on real and synthetic data demonstrate that the current approach can consistently outperform the state-of-the-art Bayesian Committee Machine (BCM) which employs a random data partitioning strategy. Finally, extensive evaluations over a grid-based computational infrastructure using the NetSolve distributed computing system show that the present approach scales well with data and could potentially be used in large-scale data mining applications.

## Keywords

Gaussian Processes, Regression, Kernel Machines, Parallel Architectures, Grid Computing.

---

\*Arindam Choudhury, Prasanth B. Nair and Andy J. Keane {A.Choudhury, P.B.Nair, Andy.Keane}@soton.ac.uk, are with the Computational Engineering and Design Center, School of Engineering Sciences, University of Southampton

local learning strategies construct models for spatially localized partitions of the data. The motivation for this arises from the fact that local learning techniques are known to perform better than meta-learning techniques which suffer from high variance and poor generalization. Vapnik (1996) proposed a local learning technique which creates the model at runtime after the testing point is known. This is in fact a lazy learning strategy analogous to the k-nearest neighbor algorithm which selects a subset of points from the training data closest to the testing point to make a prediction. Even though, this approach has provided breakthrough performance in some application domains, the associated computational cost precludes its application in practice.

In what follows, we approach the local learning problem explicitly from the point of view of the covariance function, which is the fundamental building block of the Gaussian Stochastic Process. Specifically, we show that a compactly supported covariance function defined over spatially localized clusters naturally leads to a data parallel scheme for modeling large-scale data. It is shown that using this new covariance function, the maximum likelihood estimation (MLE) problem can be decomposed into smaller decoupled subproblems. This leads to a significant speedup in the training phase of GP models. Further, the present approach also leads to sparse<sup>1</sup> yet effective predictive models which consistently outperform Bayesian Committee machines (BCM) – a state-of-the-art approach for model aggregation. Since equally sized or balanced partitioning of data is crucial to the efficiency of such data parallel approaches, we next develop a novel clustering scheme GeoClust, which outputs clusters of (nearly) equal sizes. Finally although GP models serve as a good testbed for testing our algorithms, we stress that our data parallel local learning strategy may also be readily applied to other kernel machines such as SVMs.

Detailed experimental simulations presented later in the text support our intuitions about the algorithms presented and show conclusively that our data parallel approach scales favorably with data size both in sequential mode where the models are trained sequentially and in parallel where they maximally leverage the presence of multiple processors. The parallel scalability studies have been conducted using NetSolve [21], a computational platform which facilitates grid-based heterogeneous computing in a transparent and efficient manner. The Southampton Computational Grid was utilized in part to conduct these studies. It is generally recognized that data mining in the future would shift to such grid based computing platforms and architectures[22].

The rest of the paper is organized as follows: In section 2, the basic GP regression formulation is discussed briefly along with the associated computational hurdles. Section 3 present the essential ingredients of the proposed data parallel local learning strategy. Section 4 elaborates on load balancing requirements arising out of the data parallel paradigm and proposes a new clustering algorithm tailored to satisfy such requirements. Section 5 discusses aggregating predictions from multiple models and shows how it can be done efficiently in the present case without resorting to a model aggregation strategy. Section 6 presents a discussion on the training

---

<sup>1</sup>The term *sparse* is used here loosely to denote models which only use a subset of the training dataset to make predictions.

where,

$$\hat{t}_{n+1} = \mathbf{k}_{n+1}^T(\mathbf{x})\mathbf{C}_n^{-1}\mathbf{t}_n \quad (5)$$

$$\sigma^2 = k(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}; \theta) - \mathbf{k}_{n+1}^T(\mathbf{x})\mathbf{C}_n^{-1}\mathbf{k}_{n+1} \quad (6)$$

where,  $\hat{t}_{n+1}$  and  $\sigma^2$  is the prediction for the posterior mean and the variance, respectively, and  $\mathbf{k}_{n+1} = \{k(x_{n+1}, x_1), k(x_{n+1}, x_2), \dots, k(x_{n+1}, x_n)\} \in \mathbb{R}^n$ .

## 2.1 Determining the Optimal Hyperparameters

From a computational perspective, the search for an optimal GP regressor under the evidence maximization framework [4] involves solving the following nonlinear maximum likelihood estimation (MLE) problem to determine the most probable hyperparameters  $\theta_{MP}$  for the given data.

$$\text{Maximize}_{\theta} L(\theta) = -\frac{1}{2} \log \det \mathbf{C}_n - \frac{1}{2} \mathbf{t}_n^T \mathbf{C}_n^{-1} \mathbf{t}_n - \frac{n}{2} \log 2\pi \quad (7)$$

where  $L(\theta)$  denotes the log likelihood for a Gaussian process.

Since computing  $L(\theta)$  and its gradient generally involves computing and inverting a dense  $n \times n$  covariance matrix (requiring  $\mathcal{O}(n^3)$  resources) at each iteration, training the GP model can be prohibitively expensive even for moderately sized data (e.g., say a few thousand data points). This is amply demonstrated in Figure 1 where the training time is shown for a typical 10 variable dataset. Unfortunately, other schemes such as Markov chain Monte Carlo sampling for approximating the predictive distribution [3, 2, 5] are also plagued by similar problems involving escalating training costs with data size. In any case, it has been shown that evidence maximization is faster and performs better than its counterparts on large datasets[3].

Since the main obstacle to efficient maximization of (7) is the presence of the dense covariance matrix, it would seem natural to choose a covariance function which leads to a sparse or banded covariance matrix. Similar concerns have been raised and addressed elsewhere in the literature [17] in the context of radial basis function networks. In particular, the use of compactly supported radial basis functions has been shown to bring about appreciable speedups in the training as well as prediction phases. In what follows, we pursue a similar line of argument but for the Gaussian process model. It is easily shown that such a line of approach naturally leads to a data parallel local learning scheme for GP regression.

## 3 Data Parallel Local Learning

Since a Gaussian stochastic process is completely specified by its covariance function, training a GP involves considering a parameterized covariance function and determining its hyperparameters  $\theta$  such that the log likelihood  $L(\theta)$  of the data is maximized (see Eqn. 7). In this section, we propose a compactly supported covariance function to facilitate data parallel local learning.

where  $\mathbf{K}_{ii} \in \mathbb{R}^{n_i \times n_i}$  contains correlation terms explicitly from the  $i$ th cluster which consists of  $n_i$  points. Since in this case the determinant of the covariance matrix  $\mathbf{K}$  can be written as the product of determinants of the blocks  $\mathbf{K}_{11}$  and  $\mathbf{K}_{22}$ , the log likelihood can be split into individual log likelihoods for the two partitions, i.e.,

$$L(\theta) = L(\theta_1) + L(\theta_2) \quad (11)$$

It is important to note at this point that while the conventional GP regression model uses a single set of hyperparameters to model the full data, the data parallel version uses different sets of hyperparameters to model the covariance structures of different regions of the data. Hence, maximizing the overall log likelihood for the data is equivalent to maximizing the log likelihood for each partition, since the individual objective functions depend on separate sets of hyperparameters. It may seem at first sight that the proposed compactly supported covariance function would result in a less accurate oversimplified global model. However, the results in the later sections do not seem to bear out this observation. Indeed, it can be argued that given enough flexibility in the choice of the local covariance models, the performance degradation can be somewhat circumvented.

From the preceding discussion, it is clear that the use of a compactly supported covariance function naturally leads to a data parallel local learning approach to GP regression and hence provides a means to handle large datasets. Having made this connection, we are now confronted with a problem: that of partitioning or clustering the data into subsets. From a data parallel point of view the subset sizes chosen need to be (nearly) equal.<sup>3</sup> What we really need is an algorithm which creates clusters which are localized in the data space. Local models can then be built on each of these clusters. Finally the ensemble of local models can be used in place of the global model. As argued in Vapnik [23], for unevenly distributed data, such local learning may lead to effective capacity control which can significantly improve the performance for some cases.

Even though, we discussed the data parallel approach in the context of GP regression, our approach can also be applied to other kernel-based machine learning algorithms such as SVMs. The basic approach is elucidated as a pseudo code in Algorithm 1, where  $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_m$  denote the  $m$  clusters or partitions of data and  $M_1, M_2, \dots, M_m$  denote the corresponding local prediction models. As is clear from Algorithm 1, a complete description of the data parallel approach requires the specification of three functions viz., `Partition()`, `CreateLocalModel()` and `Aggregate()`.


---

<sup>3</sup>This is primarily because of efficiency reasons and also to reduce the waiting times when operating in a multiprocessor scenario.

cluster center  $c_3$ . At the end of each iteration, the cluster strengths  $W_1, W_2, W_3$  are reassessed. The process continues till the increments become zero or very small. The iterative process is assisted by a relaxation/learning factor  $\alpha$ . Typically,  $\alpha$  is set to a low value. In our experiments, we have found  $10^{-2}$  to be reasonable choice when the data is normalized. Note that as the cluster strengths become close, the relative movement of the centers towards each other slows down and finally ceases. Hence, the algorithm converges when the clusters have equal strengths.<sup>4</sup> Figure 2 illustrates how GeoClust partitions a toy 2-D dataset into 2, 3 and 4 clusters respectively.

*Algorithm 2: GeoClust*

<b>Inputs</b>	: Data D
	: Number of partitions desired $m$
	: Maximum number of iterations $Maxiter$
	: learning parameter $\alpha$
<b>Initialize</b>	$m$ random centers, $c_i^0, i = 1, 2, \dots, m$
For $t = 1 : Maxiter$	
•	Assign each data point to the cluster nearest to it
•	Compute number of points in each cluster: $W_1, W_2, \dots, W_m$
•	For each cluster $C_i$ , Update the cluster center as follows:
o	Compute $\delta v_i = \sum_{j=1, j \neq i}^m (\frac{W_j}{W_i} - 1)(c_j^{t-1} - c_i^{t-1})$
o	Update center $c_i^t = c_i^{t-1} + \alpha \delta v_i$
•	End
<b>End</b>	



**Figure 2.** The partitions obtained by GeoClust for a toy dataset when the desired number of partitions are increased from 2 to 4 (left to right)

*Remark* : Ideally one would like to solve a graph partitioning problem where  $m$  disjoint partitions of a set of  $N$  nodes (points) are sought so as to minimize the total weight of the edges connecting nodes from different sets. In the present case, the weights of the edges are given by the correlation between two points. Unfortunately,

<sup>4</sup>When the number of instances is not exactly divisible by the number of clusters, the cluster strengths settle to nearly equal values.

( $\mathcal{O}(n^2)$  memory) and computing its inverse ( $\mathcal{O}(n^3)$  flops). The computational advantages of the present local learning strategy over the conventional GP regression approach are significant and are straightforward to evaluate. Specifically, for  $n$  data points and  $m$  clusters, a sequential data parallel version of our algorithm requires training time of  $\mathcal{O}(m\frac{n^3}{m^3})$  only as compared to  $\mathcal{O}(n^3)$  for the global GP regression model. Further, the attendant memory requirements are reduced from  $\mathcal{O}(n^2)$  for the conventional case to  $\mathcal{O}(\frac{n^2}{m^2})$ . The parallel version of the algorithm offers even more advantages. Typically, in such a scenario, one has access to at least as many processors as the number of clusters specified.<sup>5</sup> The overall training time now reduces to  $\mathcal{O}(\frac{n^3}{m^3})$  and the memory requirements fall to  $\mathcal{O}(\frac{n^2}{m^2})$  per processor involved.

The prediction model obtained using the compactly supported covariance function is extremely cheap in terms of both runtime computations and memory required as compared to BCM. Since a BCM prediction involves the additional cost of computing the variance of prediction for each model using equation (6), it effectively means explicit computation and storage of multiple matrix inverses. Hence runtime computation and memory requirements both scale as  $\mathcal{O}(m\frac{n^2}{m^2})$ . In contrast, the runtime complexity of the predictive model arising from the data parallel local learning approach scales only as  $\mathcal{O}(mp + \frac{n}{m})$  flops, where  $p$  is the dimensionality of the input data. Additionally, as demonstrated in the next section, experimental investigations suggest that the current strategy can indeed outperform BCM consistently in terms of accuracy of prediction.

It is clear from the preceding discussion that the efficiency of the algorithm increases with increasing number of clusters  $m$ . Further, it has been shown elsewhere that the use of compactly supported kernels produces more stable solutions [18]. However, it is important to mention that increasing the number of clusters beyond a certain limit actually leads to a deterioration in the generalization capability of the resulting model [18]. This results from an ever reducing number of points available for modeling each cluster.<sup>6</sup>

## 7 Experimental Investigation

The aim of the experimental studies presented in this section is two-fold: (i) we seek a direct comparison of the performance of the proposed algorithm with respect to the state-of-the-art BCM algorithm suggested recently [16] for handling large data using data parallel GP models, and (ii) we carry out a full scale timing study of our technique on a large data set to see how it scales with the size of the training data. Simultaneously, we also compute training times for the full scale GP model trained on the whole data.

We shall use the following acronyms henceforth to refer to the competing schemes: *LL-GP* for the scheme which uses the proposed local learning approach and *RP-BCM* for the scheme which uses a Random Partitioning of data and a BCM

<sup>5</sup>Conversely, it is reasonable to assume that the number of clusters are prespecified by the user according to the number of processors available.

<sup>6</sup>From the point of view of compactly supported covariance functions, poor generalization arises when the support radius is reduced below a certain threshold.




Figure 4. MSE performance for 100 realizations of a synthetic function with 10 variables generated using  $F^*$  with varying cluster sizes.

Methods	Measure	No. of clusters					
		1	2	4	6	8	10
LL-GP	MSE	8.04	8.98	9.33	10.38	10.67	10.72
	Std	6.2	8.55	7.20	8.12	8.41	7.5
	$-\log(L)(10^3)$	1.25	1.26	1.25	1.26	1.26	1.26
RP-BCM	MSE	-	11.03	15.17	17.13	20.33	40.67
	Std	-	9.5	12.73	12.84	21.56	189.22
	$-\log(L)(10^3)$	-	1.30	1.34	1.36	1.39	1.40

Table 1. Comparison of the statistics of MSE and negative log likelihood ( $-\log(L)$ ) values for the Boston dataset

test set of 2000 instances. From the figures, it is clear that the present local learning strategy is superior in terms of average MSE performance when compared to the RP-BCM. Also, this experiment clearly exhibits that the mean accuracy of prediction decreases with increasing number of clusters for both the algorithms although the training times are cut down significantly with increasing cluster sizes. This reflects the familiar speed-accuracy tradeoff, which limits the number of clusters one may use in practice and hence upper bounds the speedups one can obtain.

Table 1 summarizes the statistics for the Boston data set. The measures shown are averaged over 100 runs each of which includes training over 481 instances and testing over the remaining 25. Note that for a single cluster, the LL-GP scheme

the parallel code increases approximately in a linear fashion when the number of training points are increased. For the case of 30,000 points, the grid-enabled version is nearly 12 times faster than a sequential version of our approach. Interestingly, the conventional GP rapidly became infeasible beyond a training set size of 4000 points, which itself took 140 minutes of training time. Hence, the proposed approach in both its sequential and grid-enabled form is clearly much faster than the conventional GP modeling and is an enabling approach for handling problems with large data.

It is also worth noting that meta-learning approaches which use random partitions of the data can be deployed on grid-based computational infrastructures. However, it is often the case that the model aggregation stage cannot be fully parallelized. As a consequence, the resulting speedups may not be very significant for some cases. For example, the parallelized version of the meta-learning strategy proposed recently for SVM mixtures [12] (running on 50 compute nodes) is only 3.3 times faster than the sequential code for a problem with 100,000 training points.

## 8 Concluding Remarks

Regression techniques based on Gaussian processes have recently been applied with great success to modeling observational data. However, the impressive generalization behavior of GPs is offset by the overwhelming training cost incurred even for modest datasets. This paper proposes a novel data parallel local learning paradigm as an enabling approach for applying GP models to large-scale regression tasks. This is achieved using a combination of a compactly supported covariance function and a novel data partitioning scheme which respects load balancing considerations as well as space localization. A simple yet effective prediction strategy follows naturally from the choices effected, which is shown to consistently outperform the state-of-the-art BCM framework in terms of predictive performance. Experimental studies provide empirical evidence for the attractive scaling behavior of the proposed methodology which indicates its potential as a data mining tool.

A number of interesting lines of research are worth pursuing. From a theoretical point of view, it is of interest to examine the relationship between the number of clusters to be chosen and the resulting generalization error of the predictive model. This could result in a more principled data dependent basis for the *a priori* selection of the number of clusters to be formed to achieve reasonable generalization, while still obtaining good training and runtime speedups. While Geoclust efficiently solves an unconstrained graph partitioning problem, it remains to be seen whether imposing edge constraints improves the generalization performance any further. Lastly, it is to be noted that our data parallel paradigm can also be applied to speed-up other kernel-based machine learning algorithms.



- [13] Joachims, T., Making Large Scale Support Vector Machine Learning Possible, In *Advances in Kernel Methods*, Schölkopf, B., Burges, C., Smola, A.,(Eds.), MIT Press, 1999
- [14] Platt, J., Fast Training of Support Vector Machines using Sequential Minimal Optimization, In *Advances in Kernel Methods*, Schölkopf, B., Burges, C., Smola, A.,(Eds.), MIT Press, 1999
- [15] Ridda, A., Labbi, A., Pellegrini, C., Local Experts Combination Through Density Estimation, International Workshop on AI and Statistics, Uncertainty '99: Morgan Kaufmann, 1999
- [16] Tresp, V., A Bayesian Committee Machine, *Neural Computation*, Vol. 12, pp 2719-2741, 2000
- [17] Schaback, R., Creating surfaces from scattered data using radial basis functions, in *Mathematical Methods for Curves and Surfaces*, M. Daehlen, T. Lyche and L. Schumaker (eds.), pp 477-496 : Vanderbilt University Press, 1995
- [18] Schaback, R., Error estimates and condition numbers for radial basis function interpolation, *Advances in Computational Mathematics*, Vol. 3, pp. 251-264, 1995
- [19] Friedman, J.H., Greedy Function Approximation: A Gradient Boosting Machine, *Annals of Statistics*, To Appear
- [20] Kernighan, B.W., Lin, S., An Efficient Heuristic Procedure for Partitioning Graphs *Bell Systems Technical Journal*, Vol. 49, pp. 291-307, 1970
- [21] Casanova, H., Dongarra, J., NetSolve: A Network-enabled Server for Solving Computational Science Problems *The International Journal of Supercomputer Applications and High Performance Computing*, Vol. 11(3) pp. 212-223, 1997
- [22] Foster, I., Kesselman, C., (Eds), *The Grid: Blueprint for a New Computing Infrastructure*: Morgan Kaufmann, 1999.
- [23] Bottou L. and Vapnik V., Local learning algorithms, *Neural Computation*, Vol. 4(6), pp. 888-900, 1992.
- [24] Chan, P.K., Stolfo, S.J., Learning Arbiter and Combiner Trees from Partitioned Data for Scaling Machine Learning, *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*, pp. 39-44 : AAAI Press, 1995