
Asynchronous Distributed Variational Gaussian Process for Regression

Hao Peng¹ Shandian Zhe¹ Xiao Zhang¹ Yuan Qi²

Abstract

Gaussian processes (GPs) are powerful non-parametric function estimators. However, their applications are largely limited by the expensive computational cost of the inference procedures. Existing stochastic or distributed synchronous variational inferences, although have alleviated this issue by scaling up GPs to millions of samples, are still far from satisfactory for real-world large applications, where the data sizes are often orders of magnitudes larger, say, billions. To solve this problem, we propose ADVGP, the first Asynchronous Distributed Variational Gaussian Process inference for regression, on the recent large-scale machine learning platform, PARAMETERSERVER. ADVGP uses a novel, flexible variational framework based on a weight space augmentation, and implements the highly efficient, asynchronous proximal gradient optimization. While maintaining comparable or better predictive performance, ADVGP greatly improves upon the efficiency of the existing variational methods. With ADVGP, we effortlessly scale up GP regression to a real-world application with billions of samples and demonstrate an excellent, superior prediction accuracy to the popular linear models.

1. Introduction

Gaussian processes (GPs) (Rasmussen & Williams, 2006) are powerful non-parametric Bayesian models for function estimation. Without imposing any explicit parametric form, GPs merely induce a smoothness assumption via the definition of covariance function, and hence can flexibly infer various, complicated functions from data. In addition, GPs are robust to noise, resist overfitting and produce uncertainty estimations. However, a crucial bottleneck of GP

models is their expensive computational cost: exact GP inference requires $\mathcal{O}(n^3)$ time complexity and $\mathcal{O}(n^2)$ space complexity (n is the number of training samples), which limits GPs to very small applications, say, a few hundreds of samples.

To mitigate this limitation, many approximate inference algorithms have been developed (Williams & Seeger, 2001; Seeger et al., 2003; Quiñero-Candela & Rasmussen, 2005; Snelson & Ghahramani, 2005; Deisenroth & Ng, 2015). Most methods use sparse approximations. Basically, we first introduce a small set of inducing points; and then we develop an approximation that transfers the expensive computations from the entire large training data, such as the covariance and inverse covariance matrix calculations, to the small set of the inducing points. To this end, a typical strategy is to impose some simplified modeling assumption. For example, FITC (Snelson & Ghahramani, 2005) makes a fully conditionally independent assumption. Recently, Titsias (2009) proposed a more principled, variational sparse approximation framework, where the inducing points are also treated as variational parameters. The variational framework is less prone to overfitting and often yields a better inference quality (Titsias, 2009; Bauer et al., 2016). Based on the variational approximation, Hensman et al. (2013) developed a stochastic variational inference (SVI) algorithm, and Gal et al. (2014) used a tight variational lower bound to develop a distributed inference algorithm with the MAPREDUCE framework.

While SVI and the distributed variational inference have successfully scaled up GP models to millions of samples ($\mathcal{O}(10^6)$), they are still insufficient for real-world large-scale applications, in which the data sizes are often orders of magnitude larger, say, over billions of samples ($\mathcal{O}(10^9)$). Specifically, SVI (Hensman et al., 2013) sequentially processes data samples and requires too much time to complete even one epoch of training. The distributed variational algorithm in (Gal et al., 2014) uses the MAPREDUCE framework and requires massive synchronizations during training, where a large amount of time is squandered when the MAPPERS or REDUCERS are waiting for each other, or the failed nodes are restarted.

To tackle this problem, we propose Asynchronous Distributed Variational Gaussian Process inference (AD-

¹Purdue University, West Lafayette, IN, USA ²Ant Financial Service Group. Correspondence to: Hao Peng <pengh@alumni.purdue.edu>.

VGP), which enables GP regression on applications with (at least) billions of samples. To the best of our knowledge, this is the first variational inference that scales up GPs to this level. The contributions of our work are summarized as follows: first, we propose a novel, general variational GP framework using a weight space augmentation (Section 3). The framework allows flexible constructions of feature mappings to incorporate various low-rank structures and to fulfill different variational model evidence lower bounds (ELBOs). Furthermore, due to the simple standard normal prior of the random weights, the framework enables highly efficient, asynchronous proximal gradient-based optimization, with convergence guarantees as well as fast, element-wise and parallelizable variational posterior updates. Second, based on the new framework, we develop a highly efficient, asynchronous variational inference algorithm in the recent distributed machine learning platform, PARAMETERSERVER (Li et al., 2014a) (Section 4). The asynchronous algorithm eliminates an enormous amount of waiting time caused by the synchronous coordination, and fully exploits the computational power and network bandwidth; as a result, our new inference, ADVGP, greatly improves on both the scalability and efficiency of the prior variational algorithms, while still maintaining a similar or better inference quality. Finally, in a real-world application with billions of samples, we effortlessly train a GP regression model with ADVGP and achieve an excellent prediction accuracy, with 15% improvement over the popular linear regression implemented in Vowpal Wabbit (Agarwal et al., 2014), the state-of-the-art large-scale machine learning software widely used in industry.

2. Gaussian Processes Review

In this paper, we focus on Gaussian process (GP) regression. Suppose we aim to infer an underlying function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ from an observed dataset $\mathcal{D} = \{\mathbf{X}, \mathbf{y}\}$, where $\mathbf{X} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top]^\top$ is the input matrix and \mathbf{y} is the output vector. Each row of \mathbf{X} , namely \mathbf{x}_i ($1 \leq i \leq n$), is a d -dimensional input vector. Correspondingly, each element of \mathbf{y} , namely y_i , is an observed function value corrupted by some random noise. Note that the function f can be highly nonlinear. To estimate f from \mathcal{D} , we place a GP prior over f . Specifically, we treat the collection of all the function values as one realization of the Gaussian process. Therefore, the finite projection of f over the inputs \mathbf{X} , i.e., $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]$ follows a multivariate Gaussian distribution: $\mathbf{f} \sim \mathcal{N}(\mathbf{f}|\bar{\mathbf{f}}, \mathbf{K}_{nn})$, where $\bar{\mathbf{f}} = [\bar{f}(\mathbf{x}_1), \dots, \bar{f}(\mathbf{x}_n)]$ are the mean function values and \mathbf{K}_{nn} is the $n \times n$ covariance matrix. Each element of \mathbf{K}_{nn} is a covariance function $k(\cdot, \cdot)$ of two input vectors, i.e., $[\mathbf{K}_{nn}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. We can choose any symmetric positive semidefinite kernel as the covariance function, e.g., the ARD kernel: $k(\mathbf{x}_i, \mathbf{x}_j) = a_0^2 \exp(-\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^\top \text{diag}(\boldsymbol{\eta})(\mathbf{x}_i - \mathbf{x}_j))$,

where $\boldsymbol{\eta} = [1/a_1^2, \dots, 1/a_d^2]$. For simplicity, we usually use the zero mean function, namely $\bar{f}(\cdot) = 0$.

Given \mathbf{f} , we use an isotropic Gaussian model to sample the observed noisy output \mathbf{y} : $p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{y}|\mathbf{f}, \beta^{-1}\mathbf{I})$. The joint probability of GP regression is

$$p(\mathbf{y}, \mathbf{f}|\mathbf{X}) = \mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K}_{nn})\mathcal{N}(\mathbf{y}|\mathbf{f}, \beta^{-1}\mathbf{I}). \quad (1)$$

Further, we can obtain the marginal distribution of \mathbf{y} , namely the model evidence, by marginalizing out \mathbf{f} :

$$p(\mathbf{y}|\mathbf{X}) = \mathcal{N}(\mathbf{y}|\mathbf{0}, \mathbf{K}_{nn} + \beta^{-1}\mathbf{I}). \quad (2)$$

The inference of GP regression aims to estimate the appropriate kernel parameters and noise variance from the training data $\mathcal{D} = \{\mathbf{X}, \mathbf{y}\}$, such as $\{a_0, \boldsymbol{\eta}\}$ in ARD kernel and β^{-1} . To this end, we can maximize the model evidence in (2) with respect to those parameters. However, to maximize (2), we need to calculate the inverse and the determinant of the $n \times n$ matrix $\mathbf{K}_{nn} + \beta^{-1}\mathbf{I}$ to evaluate the multivariate Gaussian term. This will take $\mathcal{O}(n^3)$ time complexity and $\mathcal{O}(n^2)$ space complexity and hence is infeasible for a large number of samples, i.e., large n .

For prediction, given a test input \mathbf{x}^* , since the test output f^* and training output \mathbf{f} can be treated as another GP projection on \mathbf{X} and \mathbf{x}^* , the joint distribution of f^* and \mathbf{f} is also a multivariate Gaussian distribution. Then by marginalizing out \mathbf{f} , we can obtain the posterior distribution of f^* :

$$p(f^*|\mathbf{x}^*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(f^*|\alpha, v), \quad (3)$$

where

$$\alpha = \mathbf{k}_{n*}^\top (\mathbf{K}_{nn} + \beta^{-1}\mathbf{I})^{-1} \mathbf{y}, \quad (4)$$

$$v = k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}_{n*}^\top (\mathbf{K}_{nn} + \beta^{-1}\mathbf{I})^{-1} \mathbf{k}_{n*}, \quad (5)$$

and $\mathbf{k}_{n*} = [k(\mathbf{x}^*, \mathbf{x}_1), \dots, k(\mathbf{x}^*, \mathbf{x}_n)]^\top$. Note that the calculation also requires the inverse of $\mathbf{K}_{nn} + \beta^{-1}\mathbf{I}$ and hence takes $\mathcal{O}(n^3)$ time complexity and $\mathcal{O}(n^2)$ space complexity.

3. Variational Framework Using Weight Space Augmentation

Although GPs allow flexible function inference, they have a severe computational bottleneck. The training and prediction both require $\mathcal{O}(n^3)$ time complexity and $\mathcal{O}(n^2)$ space complexity (see (2), (4) and (5)), making GPs unrealistic for real-world, large-scale applications, where the number of samples (i.e., n) are often billions or even larger. To address this problem, we propose ADVGP that performs highly efficient, asynchronous distributed variational inference and enables the training of GP regression on extremely large data. ADVGP is based on a novel variational GP framework using a weight space augmentation, which is discussed below.

First, we construct an equivalent augmented model by introducing an $m \times 1$ auxiliary random weight vector \mathbf{w} ($m \ll n$). We assume \mathbf{w} is sampled from the standard normal prior distribution: $p(\mathbf{w}) = \mathcal{N}(\mathbf{w}|\mathbf{0}, \mathbf{I})$. Given \mathbf{w} , we sample an $n \times 1$ latent function values \mathbf{f} from

$$p(\mathbf{f}|\mathbf{w}) = \mathcal{N}(\mathbf{f}|\Phi\mathbf{w}, \mathbf{K}_{nn} - \Phi\Phi^\top), \quad (6)$$

where Φ is an $n \times m$ matrix: $\Phi = [\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n)]^\top$. Here $\phi(\cdot)$ represents a feature mapping that maps the original d -dimensional input into an m -dimensional feature space. Note that we need to choose an appropriate $\phi(\cdot)$ to ensure the covariance matrix in (6) is symmetric positive semidefinite. Flexible constructions of $\phi(\cdot)$ enable us to fulfill different variational model evidence lower bounds (ELBO) for large-scale inference, which we will discuss more in Section 5.

Given \mathbf{f} , we sample the observed output \mathbf{y} from the isotropic Gaussian model $p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{y}|\mathbf{f}, \beta^{-1}\mathbf{I})$. The joint distribution of our augmented model is then given by

$$p(\mathbf{y}, \mathbf{f}, \mathbf{w}|\mathbf{X}) = \mathcal{N}(\mathbf{w}|\mathbf{0}, \mathbf{I})\mathcal{N}(\mathbf{f}|\Phi\mathbf{w}, \mathbf{K}_{nn} - \Phi\Phi^\top)\mathcal{N}(\mathbf{y}|\mathbf{f}, \beta^{-1}\mathbf{I}). \quad (7)$$

This model is equivalent to the original GP regression—when we marginalize out \mathbf{w} , we recover the joint distribution in (1); we can further marginalize out \mathbf{f} to recover the model evidence in (2). Note that our model is distinct from the traditional weight space view of GP regression (Rasmussen & Williams, 2006): the feature mapping $\phi(\cdot)$ is not equivalent to the underlying (nonlinear) feature mapping induced by the covariance function (see more discussions in Section 5). Instead, $\phi(\cdot)$ is defined for computational purpose only—that is, to construct a tractable variational evidence lower bound (ELBO), shown as follows.

Now, we derive the tractable ELBO based on the weight space augmented model in (7). The derivation is similar to (Titsias, 2009; Hensman et al., 2013). Specifically, we first consider the conditional distribution $p(\mathbf{y}|\mathbf{w})$. Because $\log p(\mathbf{y}|\mathbf{w}) = \log \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{w})d\mathbf{f} = \log \langle p(\mathbf{y}|\mathbf{f}) \rangle_{p(\mathbf{f}|\mathbf{w})}$, where $\langle \cdot \rangle_{p(\theta)}$ denotes the expectation under the distribution $p(\theta)$, we can use Jensen’s inequality to obtain a lower bound:

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{w}) &= \log \langle p(\mathbf{y}|\mathbf{f}) \rangle_{p(\mathbf{f}|\mathbf{w})} \geq \langle \log p(\mathbf{y}|\mathbf{f}) \rangle_{p(\mathbf{f}|\mathbf{w})} \\ &= \sum_{i=1}^n \log \mathcal{N}(y_i|\phi^\top(\mathbf{x}_i)\mathbf{w}, \beta^{-1}) - \frac{\beta}{2}\tilde{k}_{ii}, \end{aligned} \quad (8)$$

where \tilde{k}_{ii} is the i th diagonal element of $\mathbf{K}_{nn} - \Phi\Phi^\top$.

Next, we introduce a variational posterior $q(\mathbf{w})$ to construct

the variational lower bound of the log model evidence,

$$\begin{aligned} \log p(\mathbf{y}) &= \log \left\langle \frac{p(\mathbf{y}|\mathbf{w})p(\mathbf{w})}{q(\mathbf{w})} \right\rangle_{q(\mathbf{w})} \\ &\geq \langle \log p(\mathbf{y}|\mathbf{w}) \rangle_{q(\mathbf{w})} - \text{KL}(q(\mathbf{w})\|p(\mathbf{w})). \end{aligned} \quad (9)$$

where $\text{KL}(\cdot\|\cdot)$ is the Kullback–Leibler divergence. Replacing $\log p(\mathbf{y}|\mathbf{w})$ in (9) by the right side of (8), we obtain the following lower bound,

$$\begin{aligned} \log p(\mathbf{y}) &\geq \mathcal{L} = -\text{KL}(q(\mathbf{w})\|p(\mathbf{w})) \\ &+ \sum_{i=1}^n \langle \log \mathcal{N}(y_i|\phi^\top(\mathbf{x}_i)\mathbf{w}, \beta^{-1}) \rangle_{q(\mathbf{w})} - \frac{\beta}{2}\tilde{k}_{ii}. \end{aligned} \quad (10)$$

Note that this is a variational lower bound: the equality is obtained when $\Phi\Phi^\top = \mathbf{K}_{nn}$ and $q(\mathbf{w}) = p(\mathbf{w}|\mathbf{y})$. To achieve equality, we need to set $m = n$ and have $\phi(\cdot)$ map the d -dimensional input into an n -dimensional feature space. In order to reduce the computational cost, however, we can restrict m to be very small and choose any family of mappings $\phi(\cdot)$ that satisfy $\mathbf{K}_{nn} - \Phi\Phi^\top \succeq \mathbf{0}$. The flexible choices of $\phi(\cdot)$ allows us to explore different approximations in a unified variational framework. For example, in our practice, we introduce an $m \times d$ inducing matrix $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_m]^\top$ and define

$$\phi(\mathbf{x}) = \mathbf{L}^\top \mathbf{k}_m(\mathbf{x}), \quad (11)$$

where $\mathbf{k}_m(\mathbf{x}) = [k(\mathbf{x}, \mathbf{z}_1), \dots, k(\mathbf{x}, \mathbf{z}_m)]^\top$ and \mathbf{L} is the lower triangular Cholesky factorization of the inverse kernel matrix over \mathbf{Z} , i.e., $[\mathbf{K}_{mm}]_{i,j} = k(\mathbf{z}_i, \mathbf{z}_j)$ and $\mathbf{K}_{mm}^{-1} = \mathbf{L}\mathbf{L}^\top$. It can be easily verified that $\Phi\Phi^\top = \mathbf{K}_{nm}\mathbf{K}_{mm}^{-1}\mathbf{K}_{nm}^\top$, where \mathbf{K}_{nm} is the cross kernel matrix between \mathbf{X} and \mathbf{Z} , i.e., $[\mathbf{K}_{nm}]_{i,j} = k(\mathbf{x}_i, \mathbf{z}_j)$. Therefore $\mathbf{K}_{nn} - \Phi\Phi^\top$ is always positive semidefinite, because it can be viewed as a Schur complement of \mathbf{K}_{nn} in the block matrix $\begin{bmatrix} \mathbf{K}_{mm} & \mathbf{K}_{nm}^\top \\ \mathbf{K}_{nm} & \mathbf{K}_{nn} \end{bmatrix}$. We discuss other choices of $\phi(\cdot)$ in Section 5.

4. Delayed Proximal Gradient Optimization for ADVGP

A major advantage of our variational GP framework is the capacity of using the asynchronous, delayed proximal gradient optimization supported by PARAMETERSERVER (Li et al., 2014b), with convergence guarantees and scalability to huge data. PARAMETERSERVER is a well-known, general platform for asynchronous machine learning algorithms for extremely large applications. It has a bipartite architecture where the computing nodes are partitioned into two classes: server nodes store the model parameters and worker nodes the data. PARAMETERSERVER assumes the learning procedure minimizes a non-convex loss function

with the following composite form:

$$L(\boldsymbol{\theta}) = \sum_{k=1}^r G_k(\boldsymbol{\theta}) + h(\boldsymbol{\theta}) \quad (12)$$

where $\boldsymbol{\theta}$ are the model parameters. Here $G_k(\boldsymbol{\theta})$ is a (possibly non-convex) function associated with the data in worker k and therefore can be calculated by worker k independently; $h(\boldsymbol{\theta})$ is a convex function with respect to $\boldsymbol{\theta}$.

To efficiently minimize the loss function in (12), PARAMETERSERVER uses a delayed proximal gradient updating method to perform asynchronous optimization. To illustrate it, let us first review the standard proximal gradient descent. Specifically, for each iteration t , we first take a gradient descent step according to $\sum_k G_k(\boldsymbol{\theta})$ and then perform a proximal operation to project $\boldsymbol{\theta}$ toward the minimum of $h(\cdot)$, i.e., $\boldsymbol{\theta}^{(t+1)} = \text{Prox}_{\gamma_t}[\boldsymbol{\theta}^{(t)} - \gamma_t \sum_k \nabla G_k(\boldsymbol{\theta}^{(t)})]$, where γ_t is the step size. The proximal operation is defined as

$$\text{Prox}_{\gamma_t}[\boldsymbol{\theta}] = \underset{\boldsymbol{\theta}^*}{\text{argmin}} \quad h(\boldsymbol{\theta}^*) + \frac{1}{2\gamma_t} \|\boldsymbol{\theta}^* - \boldsymbol{\theta}\|_2^2. \quad (13)$$

The standard proximal gradient descent guarantees to find a local minimum solution. However, the computation is inefficient, even in parallel: in each iteration, the server nodes wait until the worker nodes finish calculating each $\nabla G_k(\boldsymbol{\theta}^{(t)})$; then the workers wait for the servers to finish the proximal operation. This synchronization wastes much time and computational resources. To address this issue, PARAMETERSERVER uses a delayed proximal gradient updating approach to implement asynchronous computation.

Specifically, we set a delay limit $\tau \geq 0$. At any iteration t , the servers do not enforce all the workers to finish iteration t ; instead, as long as each worker has finished an iteration no earlier than $t - \tau$, the servers will proceed to perform the proximal updates, i.e., $\boldsymbol{\theta}^{(t+1)} = \text{Prox}_{\gamma_t}[\boldsymbol{\theta}^{(t)} - \gamma_t \sum_k \nabla G_k(\boldsymbol{\theta}^{(t_k)})]$ ($t - \tau \leq t_k \leq t$), and notify all the workers with the new parameters $\boldsymbol{\theta}^{(t+1)}$. Once received the updated parameters, the workers compute and push the local gradient to the servers immediately. Obviously, this delay mechanism can effectively reduce the wait between the server and worker nodes. By setting different τ , we can adjust the degree of the asynchronous computation: when $\tau = 0$, we have no asynchronization and return to the standard, synchronous proximal gradient descent; when $\tau = \infty$, we are totally asynchronous and there is no wait at all.

A highlight is that given the composite form of the non-convex loss function in (12), the above asynchronous delayed proximal gradient descent guarantees to converge according to Theorem 1.

Theorem 1. (Li et al., 2013) *Assume the gradient of the function G_k is Lipschitz continuous, that is, there is a constant C_k such that $\|\nabla G_k(\boldsymbol{\theta}) - \nabla G_k(\boldsymbol{\theta}')\| \leq C_k \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$*

for any $\boldsymbol{\theta}, \boldsymbol{\theta}'$, and $k = 1, \dots, r$. Define $C = \sum_{k=1}^r C_k$. Also, assume we allow a maximum delay for the updates by τ and a significantly-modified filter on pulling the parameters with threshold $\mathcal{O}(t^{-1})$. For any $\epsilon > 0$, the delayed proximal gradient descent converges to a stationary point if the learning rate γ_t satisfies $\gamma_t \leq ((1 + \tau)C + \epsilon)^{-1}$.

Now, let us return to our variational GP framework. A major benefit of our framework is that the negative variational evidence lower bound (ELBO) for GP regression has the same composite form as (12). Thereby we can apply the asynchronous proximal gradient descent for GP inference on PARAMETERSERVER. Specifically, we explicitly assume $q(\mathbf{w}) = \mathcal{N}(\mathbf{w}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and obtain the negative variational ELBO (see (10))

$$-\mathcal{L} = \sum_{i=1}^n g_i + h \quad (14)$$

where

$$\begin{aligned} g_i &= -\log \mathcal{N}(y_i | \boldsymbol{\phi}^\top(\mathbf{x}_i)\boldsymbol{\mu}, \beta^{-1}) + \frac{\beta}{2} \boldsymbol{\phi}^\top(\mathbf{x}_i)\boldsymbol{\Sigma}\boldsymbol{\phi}(\mathbf{x}_i) \\ &\quad + \frac{\beta}{2} \tilde{k}_{ii}, \\ h &= \frac{1}{2} (-\ln |\boldsymbol{\Sigma}| - m + \text{tr}(\boldsymbol{\Sigma}) + \boldsymbol{\mu}^\top \boldsymbol{\mu}). \end{aligned} \quad (15)$$

Instead of directly updating $\boldsymbol{\Sigma}$, we consider \mathbf{U} , the upper triangular Cholesky factor of $\boldsymbol{\Sigma}$, i.e., $\boldsymbol{\Sigma} = \mathbf{U}^\top \mathbf{U}$. This not only simplifies the proximal operation but also ensures the positive definiteness of $\boldsymbol{\Sigma}$ during computation. The partial derivatives of g_i with respect to $\boldsymbol{\mu}$ and \mathbf{U} are

$$\frac{\partial g_i}{\partial \boldsymbol{\mu}} = \beta (-y_i \boldsymbol{\phi}(\mathbf{x}_i) + \boldsymbol{\phi}(\mathbf{x}_i)\boldsymbol{\phi}^\top(\mathbf{x}_i)\boldsymbol{\mu}), \quad (16)$$

$$\frac{\partial g_i}{\partial \mathbf{U}} = \beta \text{triu}[\mathbf{U}\boldsymbol{\phi}(\mathbf{x}_i)\boldsymbol{\phi}^\top(\mathbf{x}_i)], \quad (17)$$

where $\text{triu}[\cdot]$ denotes the operator that keeps the upper triangular part of a matrix but leaves any other element zero. It can be verified that the partial derivatives of g_i with respect to $\boldsymbol{\mu}$ and \mathbf{U} are Lipschitz continuous and h is also convex with respect to $\boldsymbol{\mu}$ and \mathbf{U} . According to Theorem 1, minimizing $-\mathcal{L}$ (i.e., maximizing \mathcal{L}) with respect to the variational parameters, $\boldsymbol{\mu}$ and \mathbf{U} , using the asynchronous proximal gradient method can guarantee convergence. For other parameters, such as kernel parameters and inducing points, h is simply a constant. As a result, the delayed proximal updates for these parameters reduce to the delayed gradient descent optimization such as in (Agarwal & Duchi, 2011).

We now present the details of ADVGP implementation on PARAMETERSERVER. We first partition the data for r workers and allocate the model parameters (such as the kernel parameters, the parameters of $q(\mathbf{w})$ and the inducing points \mathbf{Z}) to server nodes. At any iteration t , the server nodes aggregate all the local gradients and perform the

proximal operation in (13), as long as each worker k has computed and pushed the local gradient on its own data subset D_k for some prior iteration t_k ($t - \tau \leq t_k \leq t$), $\nabla G_k^{(t_k)} = \sum_{i \in D_k} \nabla g_i^{(t_k)}$. Note that the proximal operation is only performed for the parameters of $q(\mathbf{w})$, namely $\boldsymbol{\mu}$ and \mathbf{U} ; since h is constant for the other model parameters, such as the kernel parameters and the inducing points, their gradient descent updates remain unchanged. Minimizing (13) by setting the derivatives to zero, we obtain the proximal updates for each element in $\boldsymbol{\mu}$ and \mathbf{U} :

$$\mu_i^{(t+1)} = \mu_i'^{(t+1)} / (1 + \gamma_t), \quad (18)$$

$$U_{ij}^{(t+1)} = U_{ij}'^{(t+1)} / (1 + \gamma_t), \quad (19)$$

$$U_{ii}^{(t+1)} = \frac{U_{ii}'^{(t+1)} + \sqrt{(U_{ii}'^{(t+1)})^2 + 4(1 + \gamma_t)\gamma_t}}{2(1 + \gamma_t)}, \quad (20)$$

where

$$\mu_i'^{(t+1)} = \mu_i^{(t)} - \gamma_t \sum_{k=1}^r \frac{\partial G_k^{(t_k)}}{\partial \mu_i^{(t_k)}},$$

$$U_{ij}'^{(t+1)} = U_{ij}^{(t)} - \gamma_t \sum_{k=1}^r \frac{\partial G_k^{(t_k)}}{\partial U_{ij}^{(t_k)}}.$$

The proximal operation comprises of element-wise, closed-form computations, therefore making the updates of the variational posterior $q(\mathbf{w})$ highly parallelizable and efficient. The gradient calculation for the other parameters, including the kernel parameters and inducing points, although quite complicated, is pretty standard and we give the details in the supplementary material (Appendix A). Finally, ADVGP is summarized in Algorithm 1.

Algorithm 1 Delayed Proximal Gradient for ADVGP

Worker k at iteration t_k

- 1: Block until servers have new parameters ready.
- 2: Pull the parameters from servers and update the current version (or iteration) t_k .
- 3: Compute the gradient $\nabla G_k^{(t_k)}$ on data D_k .
- 4: Push the gradient $\nabla G_k^{(t_k)}$ to servers.

Servers at iteration t

- 1: **if** Each worker k completes iteration $t_k \geq t - \tau$ **then**
 - 2: Aggregate gradients to obtain $\nabla G^{(t)} = \sum \nabla G_k^{(t_k)}$.
 - 3: Update $\boldsymbol{\mu}$ and \mathbf{U} using (18), (19) and (20).
 - 4: Update the other parameters using gradient descent.
 - 5: Notify all blocked workers of the new parameters and the version (i.e., $t + 1$).
 - 6: Proceed to iteration $t + 1$.
 - 7: **end if**
-

5. Discussion and Related Work

Exact GP inference requires computing the full covariance matrix (and its inverse), and therefore is infeasible for large data. To reduce the computational cost, many sparse GP inference methods use a low-rank structure to approximate the full covariance. For example, Williams & Seeger (2001); Peng & Qi (2015) used the Nyström approximation; Bishop & Tipping (2000) used relevance vectors, constructed from covariance functions evaluated on a small subset of the training data. A popular family of sparse GPs introduced a small set of inducing inputs and targets, viewed as statistical summary of the data, and define an approximate model by imposing some conditional independence between latent functions given the inducing targets; the inference of the inexact model is thereby much easier. Quiñero-Candela & Rasmussen (2005) provided a unified view of those methods, such as SoR (Smola & Bartlett, 2001), DTC (Seeger et al., 2003), PITC (Schwaighofer & Tresp, 2003) and FITC (Snelson & Ghahramani, 2005).

Despite the success of those methods, their inference procedures often exhibit undesirable behaviors, such as underestimation of the noise and clumped inducing inputs (Bauer et al., 2016). To obtain a more favorable approximation, Titsias (2009) proposed a variational sparse GP framework, where the approximate posteriors and the inducing inputs are both treated as variational parameters and estimated by maximizing a variational lower bound of the true model evidence. The variational framework is less prone to overfitting and often yields a better inference quality (Titsias, 2009; Bauer et al., 2016). Based on Titsias (2009)'s work, Hensman et al. (2013) developed a stochastic variational inference for GP (SVIGP) by parameterizing the variational distributions explicitly. Gal et al. (2014) reparameterized the bound of Titsias (2009) and developed a distributed optimization algorithm with MAPREDUCE framework. Further, Dai et al. (2014) developed a GPU acceleration using the similar formulation, and Matthews et al. (2017) developed GPflow library, a TensorFlow implementation that exploit GPU hardwares.

To further enable GPs on real-world, extremely large applications, we proposed a new variational GP framework using a weight space augmentation. The proposed augmented model, introducing an extra random weight vector \mathbf{w} with standard normal prior, is distinct from the traditional GP weight space view (Rasmussen & Williams, 2006) and the recentering tricks used in GP MCMC inferences (Murray & Adams, 2010; Filippone et al., 2013; Hensman et al., 2015). In the conventional GP weight space view, the weight vector is used to combine the nonlinear feature mapping induced by the covariance function and therefore can be infinite dimensional; in the recentering tricks, the weight vector is used to reparameterize the latent function values, to

dispose of the dependencies on the hyper-parameters, and to improve the mixing rate. In our framework, however, the weight vector \mathbf{w} has a fixed, much smaller dimension than the number of samples ($m \ll n$), and is used to introduce an extra feature mapping $\phi(\cdot)$ — $\phi(\cdot)$ plays the key role to construct a tractable variational model evidence lower bound (ELBO) for large scale GP inference.

The advantages of our framework are mainly twofold. First, by using the feature mapping $\phi(\cdot)$, we are flexible to incorporate various low rank structures, and meanwhile still cast them into a principled variational inference framework. For example, in addition to (11), we can define

$$\phi(\mathbf{x}) = \text{diag}(\boldsymbol{\lambda})^{-1/2} \mathbf{Q}^\top \mathbf{k}_m(\mathbf{x}), \quad (21)$$

where \mathbf{Q} are $\boldsymbol{\lambda}$ are eigenvectors and eigenvalues of \mathbf{K}_{mm} . Then $\phi(\cdot)$ is actually a scaled Nyström approximation for eigenfunctions of the kernel used in GP regression. This actually fulfills a variational version of the EigenGP approximation (Peng & Qi, 2015). Further, we can extend (21) by combining q Nyström approximations. Suppose we have q groups of inducing inputs $\{\mathbf{Z}_1, \dots, \mathbf{Z}_q\}$, where each \mathbf{Z}_l consists of m_l inducing inputs. Then the feature mapping can be defined by

$$\phi(\mathbf{x}) = \sum_{l=1}^q q^{-1/2} \text{diag}(\boldsymbol{\lambda}_l)^{-1/2} \mathbf{Q}_l^\top \mathbf{k}_{m_l}(\mathbf{x}), \quad (22)$$

where $\boldsymbol{\lambda}_l$ and \mathbf{Q}_l are the eigenvalues and eigenvectors of the covariance matrix for \mathbf{Z}_l . This leads to a variational sparse GP based on the ensemble Nyström approximation (Kumar et al., 2009). It can be trivially verified that both (21) and (22) satisfied $\mathbf{K}_{nn} - \boldsymbol{\Phi} \boldsymbol{\Phi}^\top \succeq \mathbf{0}$ in (6).

In addition, we can also relate ADVGP to GP models with pre-defined feature mappings, for instance, Relevance Vector Machines (RVMs) (Bishop & Tipping, 2000), by setting $\phi(\mathbf{x}) = \text{diag}(\boldsymbol{\alpha}^{1/2}) \mathbf{k}_m(\mathbf{x})$, where $\boldsymbol{\alpha}$ is an $m \times 1$ vector. Note that to ensure $\mathbf{K}_{nn} - \boldsymbol{\Phi} \boldsymbol{\Phi}^\top \succeq \mathbf{0}$, we have to add some constraint over the range of each α_i in $\boldsymbol{\alpha}$.

The second major advantage of ADVGP is that our variational ELBO is consistent with the composite non-convex loss form favored by PARAMETERSERVER, therefore we can utilize the highly efficient, distributed asynchronous proximal gradient descent in PARAMETERSERVER to scale up GPs to extremely large applications (see Section 6.3). Furthermore, the simple element-wise and closed-form proximal operation enables exceedingly efficient and parallelizable variational posterior update on the server side.

6. Experiments

6.1. Predictive Performance

First, we evaluated the inference quality of ADVGP in terms of predictive performance. To this end, we used the

US Flight data¹ (Hensman et al., 2013), which recorded the arrival and departure time of the USA commercial flights between January and April in 2008. We performed two groups of tests: in the first group, we randomly chose 700K samples for training; in the second group, we randomly selected 2M training samples. Both groups used 100K samples for testing. We ensured that the training and testing data are non-overlapping.

We compared ADVGP with two existing scalable variational inference algorithms: SVIGP (Hensman et al., 2013) and DistGP (Gal et al., 2014). SVIGP employs an online training, and DistGP performs a distributed synchronous variational inference. We ran all the methods on a computer node with 16 CPU cores and 64 GB memory. While SVIGP uses a single CPU core, DistGP and ADVGP use all the CPU cores to perform parallel inference. We used ARD kernel for all the methods, with the same initialization of the kernel parameters. For SVIGP, we set the mini-batch size to 5000, consistent with (Hensman et al., 2013). For DistGP, we tested two optimization frameworks: local gradient descent (DistGP-GD) and L-BFGS (DistGP-LBFGS). For ADVGP, we initialized $\boldsymbol{\mu} = \mathbf{0}$, $\mathbf{U} = \mathbf{I}$, and used ADADELTA (Zeiler, 2012) to adjust the step size for the gradient descent before the proximal operation. To choose an appropriate delay τ , we sampled another set of training and test data, based on which we tuned τ from $\{0, 8, 16, 24, 32, 40\}$. These tuning datasets do not overlap the test data in the evaluation. Note that when $\tau = 0$, the computation is totally synchronous; larger τ results in more asynchronous computation. We chose $\tau = 32$ as it produced the best performance on the tuning datasets.

Table 1 and Table 2 report the root mean square errors (RMSEs) of all the methods using different numbers of inducing points, i.e., $m \in \{50, 100, 200\}$. As we can see, ADVGP exhibits better or comparable prediction accuracy in all the cases. Therefore, while using asynchronous computation, ADVGP maintains the same robustness and quality for inference. Furthermore, we examined the prediction accuracy of each method along with the training time, under the settings $m \in \{100, 200\}$. Figure 1 shows that during the same time span, ADVGP achieves the highest performance boost (i.e., RMSE is reduced faster than the competing methods), which demonstrates the efficiency of ADVGP. It is interesting to see that in a short period of time since the beginning, SVIGP reduces RMSE as fast as ADVGP; however, after that, RMSE of SVIGP is constantly larger than ADVGP, exhibiting an inferior performance. In addition, DistGP-LBFGS converges earlier than both ADVGP and SVIGP. However, RMSE of DistGP-LBFGS is larger than both ADVGP and SVIGP at convergence. This implies that the L-BFGS optimization converged to a sub-

¹<http://stat-computing.org/dataexpo/2009/>




Figure 1. Root mean square errors (RMSEs) for US flight data as a function of training time.

optimal solution.

Table 1. Root mean square errors (RMSEs) for 700K/100K US Flight data.

Method	$m = 50$	$m = 100$	$m = 200$
Prox GP	32.9080	32.7543	32.6143
GD Dist GP	32.9411	32.8069	32.6521
LBFG Dist GP	33.0707	33.2263	32.8729
SVIGP	33.1054	32.9499	32.7802

Table 2. RMSEs for 2M/100K US Flight data.

Method	$m = 50$	$m = 100$	$m = 200$
Prox GP	36.1156	35.8347	35.7017
GD Dist GP	36.0142	35.9487	35.7971
LBFG Dist GP	35.9809	36.1676	36.0749
SVIGP	36.2019	35.9517	35.8599

We also studied how the delay limit τ affects the performance of ADVGP. Practically, when many machines are used, some worker may always be slower than the others due to environmental factors, *e.g.*, unbalanced workloads. To simulate this scenario, we intentionally introduced a latency by assigning each worker a random sleep time of 0, 10 or 20 seconds at initialization; hence a worker would pause for its given sleep time before each iteration. In our experiment, the average per-iteration running time was only 0.176 seconds; so the fastest worker could be hundreds of iterations ahead of the slowest one in the asynchronous setting. We examined $\tau = 0, 5, 10, 20, 40, 80, 160$ and plotted RMSEs as a function of time in Figure 2. Since RMSE of the synchronous case ($\tau = 0$) is much larger than the others, we do not show it in the figure. When τ is larger, ADVGP’s performance is more fluctuating. Increasing τ , we first improved the prediction accuracy due to more efficient CPU usage; however, later we observed a decline caused by the excessive asynchronization that impaired the optimization. Therefore, to use ADVGP for workers at various paces, we need to carefully choose the appropriate delay limit τ .




Figure 2. Root mean square errors (RMSEs) as a function of time for different delay limits τ

6.2. Scalability

Next, we examined the scalability of our asynchronous inference method, ADVGP. To this end, we used the 700K/100K dataset and compared with the synchronous inference algorithm DistGP (Gal et al., 2014). For a fair comparison, we used the local gradient descent version of DistGP, *i.e.*, DistGP-GD. We conducted two experiments on 4 c4.8xlarge instances of Amazon EC2 cloud, where we set the number of inducing points $m = 100$. In the first experiment, we fixed the size of the training data, and increased the number of CPU cores from 4 to 128. We examined the per-iteration running time of both ADVGP and DistGP-GD. Figure 3(A) shows that while both decreasing with more CPU cores, the per-iteration running time of ADVGP is much less than that of DistGP-GD. This demonstrates the advantage of ADVGP in computational efficiency. In addition, the per-iteration running time of ADVGP decays much more quickly than that of DistGP-GD as the number of cores approaches 128. This implies that even the communication cost becomes dominant, the asynchronous mechanism of ADVGP still effectively reduces the latency and maintains a high usage of the computational power. In the second experiment, we simultaneously increased the number of cores and the size of training data. We started from 87.5K samples and 16 cores

and gradually increased them to 700K samples and 128 cores. As shown in Figure 3(B), the average per-iteration time of DistGP-GD grows linearly; in contrast, the average per-iteration time of ADVGP stays almost constant. We speculate that without synchronous coordination, ADVGP can fully utilize the network bandwidth so that the increased amount of messages, along with the growth of the data size, affect little the network communication efficiency. This demonstrates the advantage of asynchronous inference from another perspective.

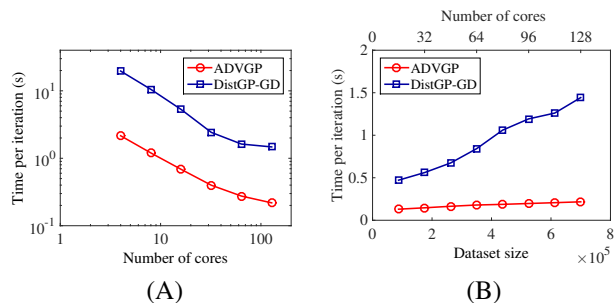


Figure 3. Scalability tests on 700K US flight data. (A) Per-iteration time as a function of available cores in log-scale. (B) Per-iteration time when scaling the computational resources proportionally to dataset size.

6.3. NYC Taxi Traveling Time Prediction

Finally, we applied ADVGP for an extremely large problem: the prediction of the taxi traveling time in New York city. We used the New York city yellow taxi trip dataset², which consist of 1.21 billions of trip records from January 2009 to December 2015. We excluded the trips that are outside the NYC area or more than 5 hours. The average traveling time is 764 seconds and the standard derivation is 576 seconds. To predict the traveling time, we used the following 9 features: time of the day, day of the week, day of the month, month, pick-up latitude, pick-up longitude, drop-off latitude, drop-off longitude, and travel distance. We used Amazon EC2 cloud, and ran ADVGP on multiple Amazon c4.8xlarge instances, each with 36 vCPUs and 60 GB memory. We compared with the linear regression model implemented in Vowpal Wabbit (Agarwal et al., 2014). Vowpal Wabbit is a state-of-the-art large scale machine learning software package and has been used in many industrial-scale applications, such as click-through-rate prediction (Chapelle et al., 2014).

We first randomly selected 100M training samples and 500K test samples. We set $m = 50$ and initialized the inducing points as the the K-means cluster centers from a subset of 2M training samples. We trained a GP regression model with ADVGP, using 5 Amazon instances with 200

²http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml

processes. The delay limit τ was selected as 20. We used Vowpal Wabbit to train a linear regression model, with default settings. We also took the average traveling time over the training data to obtain a simple mean prediction. In Figure 4(A), we report RMSEs of the linear regression and the mean prediction, as well as the GP regression along with running time. As we can see, ADVGP greatly outperforms the competing methods. Only after 6 minutes, ADVGP has improved RMSEs of the linear regression and the mean prediction by 9% and 41%, respectively; the improvements continued for about 30 minutes. Finally, ADVGP reduced the RMSEs of the linear regression and the mean prediction by 22% and 49%, respectively. The RMSEs are {ADVGP: 333.4, linear regression: 424.8, mean-prediction: 657.7}.

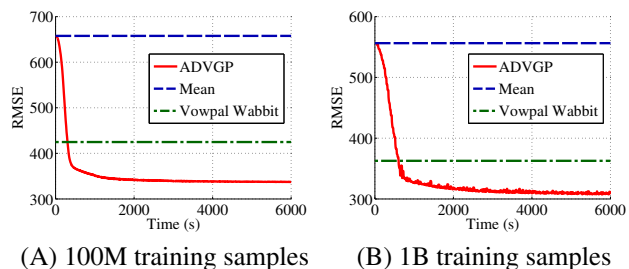


Figure 4. RMSE as a function of training time on NYC Taxi Data.

To further verify the advantage of GP regression in extremely large applications, we used 1B training and 1M testing samples. We used 50 inducing points, initialized by the K-means cluster centers from a 1M training subset. We ran ADVGP using 28 Amazon instances with 1000 processes and chose $\tau = 100$. As shown in Figure 4(B), the RMSE of GP regression outperforms the linear models by a large margin. After 12 minutes, ADVGP has improved the RMSEs of the linear regression and the mean prediction by 8% and 40%, respectively; the improvement kept growing for about 1.5 hours. At the convergence, ADVGP outperforms the linear regression and the mean prediction by 15% and 44%, respectively. The RMSEs are {ADVGP: 309.7, linear regression: 362.8, mean-prediction: 556.3}. In addition, the average per-iteration time of ADVGP is only 0.21 seconds. These results confirm the power of the nonlinear regression in extremely large real-world scenarios, comparing with linear models, while the latter are much easier to be scaled up and hence more popular.

7. Conclusion

We have presented ADVGP, an asynchronous, distributed variational inference algorithm for GP regression, which enables real-world extremely large applications. ADVGP is based on a novel variational GP framework, which allows flexible construction of low rank approximations and can relate to many sparse GP models.

References

- Agarwal, Alekh and Duchi, John C. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*, pp. 873–881. Curran Associates, Inc., 2011.
- Agarwal, Alekh, Chapelle, Olivier, Dudík, Miroslav, and Langford, John. A reliable effective terascale linear learning system. *Journal of Machine Learning Research*, 15:1111–1133, 2014.
- Bauer, Matthias, van der Wilk, Mark, and Rasmussen, Carl Edward. Understanding probabilistic sparse Gaussian process approximations. In *Advances in Neural Information Processing Systems 29*, pp. 1525–1533, 2016.
- Bishop, Christopher M. and Tipping, Michael E. Variational relevance vector machines. In *Proceedings of the 16th Conference in Uncertainty in Artificial Intelligence (UAI)*, 2000.
- Chapelle, Olivier, Manavoglu, Eren, and Rosales, Romer. Simple and scalable response prediction for display advertising. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):61:1–61:34, December 2014. ISSN 2157-6904.
- Dai, Zhenwen, Damianou, Andreas, Hensman, James, and Lawrence, Neil D. Gaussian process models with parallelization and GPU acceleration. In *NIPS Workshop on Software Engineering for Machine Learning*, 2014.
- Deisenroth, Marc and Ng, Jun W. Distributed Gaussian processes. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 1481–1490, 2015.
- Filippone, Maurizio, Zhong, Mingjun, and Girolami, Mark. A comparative evaluation of stochastic-based inference methods for gaussian process models. *Machine Learning*, 93(1):93–114, 2013.
- Gal, Yarin, van der Wilk, Mark, and Rasmussen, Carl. Distributed variational inference in sparse Gaussian process regression and latent variable models. In *Advances in Neural Information Processing Systems 27*, pp. 3257–3265, 2014.
- Hensman, James, Fusi, Nicolo, and Lawrence, Neil D. Gaussian processes for big data. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.
- Hensman, James, Matthews, Alexander G, Filippone, Maurizio, and Ghahramani, Zoubin. Mcmc for variationally sparse gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 1648–1656, 2015.
- Kumar, Sanjiv, Mohri, Mehryar, and Talwalkar, Ameet. Ensemble Nyström method. In Bengio, Y., Schuurmans, D., Lafferty, J. D., Williams, C. K. I., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems 22*, pp. 1060–1068. Curran Associates, Inc., 2009.
- Li, Mu, Andersen, David G, and Smola, Alexander J. Distributed delayed proximal gradient methods. In *NIPS Workshop on Optimization for Machine Learning*, 2013.
- Li, Mu, Andersen, David G, Park, Jun Woo, Smola, Alexander J, Ahmed, Amr, Josifovski, Vanja, Long, James, Shekita, Eugene J, and Su, Bor-Yiing. Scaling distributed machine learning with the parameter server. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pp. 583–598, 2014a.
- Li, Mu, Andersen, David G, Smola, Alexander, and Yu, Kai. Communication efficient distributed machine learning with the parameter server. In *Neural Information Processing Systems 27*, 2014b.
- Matthews, Alexander G de G, van der Wilk, Mark, Nickson, Tom, Fujii, Keisuke, Boukouvalas, Alexis, León-Villagrà, Pablo, Ghahramani, Zoubin, and Hensman, James. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, 2017.
- Murray, Iain and Adams, Ryan P. Slice sampling covariance hyperparameters of latent gaussian models. In *Advances in Neural Information Processing Systems 24*, pp. 1732–1740, 2010.
- Peng, Hao and Qi, Yuan. EigenGP: Sparse Gaussian process models with adaptive eigenfunctions. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pp. 3763–3769, 2015.
- Quiñonero-Candela, Joaquin and Rasmussen, Carl Edward. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.
- Rasmussen, Carl E. and Williams, Christopher K. I. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Schwaighofer, Anton and Tresp, Volker. Transductive and inductive methods for approximate Gaussian process regression. In *Advances in Neural Information Processing Systems 15*, pp. 953–960. MIT Press, 2003.
- Seeger, Matthias, Williams, Christopher, and Lawrence, Neil. Fast forward selection to speed up sparse Gaussian process regression. In *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, 2003.

Smola, Alexander J. and Bartlett, Peter L. Sparse greedy Gaussian process regression. In *Advances in Neural Information Processing Systems 13*. MIT Press, 2001.

Snelson, Edward and Ghahramani, Zoubin. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pp. 1257–1264, 2005.

Titsias, Michalis K. Variational learning of inducing variables in sparse Gaussian processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, pp. 567–574, 2009.

Williams, Christopher and Seeger, Matthias. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems 13*, pp. 682–688, 2001.

Zeiler, Matthew D. ADADELTA: an adaptive learning rate method. *arXiv:1212.5701*, 2012.