

Distributed Gaussian Process Regression for Mobile Sensor Networks Under Localization Uncertainty

Sungjoon Choi, Mahdi Jadalih, Jongeun Choi, and Songhwai Oh

Abstract—In this paper, we propose distributed Gaussian process regression for resource-constrained mobile sensor networks under localization uncertainty. The proposed distributed algorithm, which combines Jacobi over-relaxation (JOR) and discrete-time average consensus (DAC), can effectively handle localization uncertainty as well as limited communication ranges and computation capabilities of mobile sensor networks. The performance of the proposed method is verified in numerical simulations against the centralized maximum a posteriori solution and the quick-and-dirty solution. We show that the proposed method outperforms the quick-and-dirty solution and achieves an accuracy comparable to the centralized solution.

I. INTRODUCTION

Recent dramatic climatic changes, due to increasing greenhouse gas concentrations, stratospheric ozone depletion, and tropical deforestation, are great threats to the environment and our society [1]. The advances in embedded processors and mobile sensor networks (MSNs) technologies allow a number of important and successful applications in environmental monitoring, such as monitoring complex interactions in wildlife habitats and disaster management of harmful algal blooms in water bodies [2].

In this paper, we develop a distributed estimation method which is robust against localization and measurement errors using Gaussian process regression (GPR) [3], a nonparametric Bayesian regression method. GPR can provide both an estimate for an unseen input and the uncertainty (variance) about the estimate. A capability of assessing the uncertainty about its prediction is an attractive feature of GPR. Hence, GPR has been adapted to make statistical inferences on geostatistical and environmental data and applied to robot navigation [4] and sensor placement [5]. However, GPR suffers from two major drawbacks: a heavy computational cost and the difficulty of handling localization errors.

The computational complexity of the original form of GPR is $O(n^3)$, where n is the number of measurements, for the computation of an inverse of the covariance matrix. A number of approximation schemes have been proposed to reduce the computational complexity of GPR, including [6], [7]. In particular, Xu et al. [8] proposed Gaussian process

regression based on truncated observations for mobile sensor networks with limited memory and computational power.

For environmental monitoring, better mapping of the environment is possible when accurate sensing locations are available. However, there can be many situations where sensing locations cannot be measured accurately, e.g., GPS denied areas. Also, there can be a significant localization error with inexpensive GPS receivers. A number of localization algorithms have been proposed to address this issue, including [9], [10]. Oguz-Ekim et al. [9] iteratively maximized likelihoods of position estimates given measurements and Karlsson et al. [10] applied particle filtering for surface and underwater navigation as a supplement to the GPS.

Jadaliha et al. [11] incorporated the localization uncertainty into the Gaussian process regression framework. Since the proposed predictive mean and variance estimators have no closed-form solutions, they suggested two approximation schemes, Laplace approximation and Monte Carlo importance sampling. They also proposed a simple Laplace approximation method which uses maximum a posteriori (MAP) estimates of noisy position reports. In this paper, we extend [11] by developing a distributed version of GPR, which can handle both localization and measurement uncertainties, such that it can be suitable for resource-constrained mobile sensor networks.

The remainder of this paper is organized as follows. In Section II, Gaussian process regression is briefly described. In Section III and IV, we provide a formulation which incorporates both localization and measurement errors into a single Bayesian framework. We also introduce an approximation algorithm for computing the predictive mean estimator (PME) and predictive variance estimator (PVE). In Section V, we propose a distributed Gaussian process algorithm for computing PME and PVE. The performance of the proposed prediction algorithm is extensively evaluated in Section VI.

II. GAUSSIAN PROCESS REGRESSION

A Gaussian process (GP) is completely specified by its mean function and covariance function [3]. Let us denote the mean function by $m(x)$ and the covariance function by $k(x, x')$ for a Gaussian process $f(x)$ describing an environmental parameter. $f(x)$ can be represented as:

$$f(x) \sim GP(m(x), k(x, x')). \quad (1)$$

Suppose that $x \in \mathbb{R}^{n_x}$ is an input and $y \in \mathbb{R}$ is an output (or a target), such that $y = f(x) + w$, where w is a white Gaussian noise. When the target y is continuous

S. Choi and S. Oh are with the Department of Electrical and Computer Engineering and ASRI, Seoul National University, Seoul 151-744, Korea (e-mail: {sungjoon.choi, songhwai.oh}@cpslab.snu.ac.kr).

M. Jadalih is with the Department of Mechanical Engineering, Michigan State University, MI 48824, USA (e-mail: jadalih@egr.msu.edu).

J. Choi is with the Department of Mechanical Engineering and Electrical and Computer Engineering, Michigan State University, MI 48824, USA (e-mail: jchoi@egr.msu.edu).

This research was in part supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2013R1A1A2009348).

(respectively, discrete), we have a regression (respectively, classification) problem. In this paper, we assume y takes a continuous value. Suppose that there are n observations, $\{(x_i, y_i) | i = 1, \dots, n\}$. Given observations, Gaussian process regression (GPR) can predict an output y_* for a new input vector x_* .

For notational simplicity, we assume a zero-mean Gaussian process. A Gaussian process with a nonzero mean can be treated by a change of variables. Even without a change of variables, this is not a drastic limitation, since the mean of the posterior process is not confined to zero [3].

Assuming that $y_i = f(x_i) + w_i$ and $w_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_w^2)$, the covariance between y_i and y_j can be computed as

$$\text{cov}(y_i, y_j) = k(x_i, x_j) + \sigma_w^2 \delta_{ij}, \quad (2)$$

where $k(x, x')$ is a predefined kernel function. We represent the covariance in the following matrix form

$$\text{cov}(\mathbf{y}) = K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I, \quad (3)$$

where $\mathbf{y} = (y_1, \dots, y_n)^T$, $\mathbf{x} = (x_1^T, \dots, x_n^T)^T$, and $K(\mathbf{x}, \mathbf{x})$ is the covariance matrix computed from n data points.

Let $D = \{(x_i, y_i) | i = 1, \dots, n\}$ be a set of input-output pairs. The conditional distribution of y_* at a new input x_* given data becomes

$$y_* | D \sim N(\mu_*(x_* | D), \sigma_*^2(x_* | D)) \quad (4)$$

where

$$\mu_*(x_* | D) = k(x_*, \mathbf{x})^T (K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1} \mathbf{y} \quad (5)$$

and

$$\sigma_*^2(x_* | D) = \sigma_f^2 - k(x_*, \mathbf{x})^T (K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1} k(x_*, \mathbf{x}). \quad (6)$$

Here, $k(x_*, \mathbf{x}) \in \mathbb{R}^n$ is a covariance vector between the test point x_* and data points.

III. GPR UNDER LOCALIZATION UNCERTAINTY

In a realistic situation, acquiring samples $\{x_i, y_i\}$ with exact localization for x_i is often impossible. As Gaussian processes incorporate a measurement noise naturally, it is desirable to incorporate a localization noise into Gaussian processes. In [11], Gaussian process regression was reformulated to incorporate noisy input data $\bar{\mathbf{x}} = \{\bar{x}_1, \dots, \bar{x}_n\}$, where $\bar{x}_i = x_i + v_i$ and v_i is a zero-mean white Gaussian noise with variance σ_v^2 . Let $D' = \{(\bar{x}_i, \bar{y}_i) | i = 1, \dots, n\}$, where $\bar{y}_i = f(x_i) + w_i$. Under this new formulation, the following predictive mean estimator (PME) and predictive variance estimator (PVE) are derived in [11].

$$\mathbb{E}(y_* | D') = \frac{\int_X \mu_*(x_* | D') p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{y}}) d\mathbf{x}}{\int_X p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{y}}) d\mathbf{x}} \quad (7)$$

and

$$\text{var}(y_* | D') = \frac{\int_X (\sigma_*^2(x_* | D') + \mu_*^2(x_* | D')) p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{y}}) d\mathbf{x}}{\int_X p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{y}}) d\mathbf{x}} - \mathbb{E}^2(y_* | D'), \quad (8)$$

where $\mu_*(x_* | D')$ and $\sigma_*^2(x_* | D')$ are given by (5) and (6).

However, there are no closed-form solutions for the PME and PVE given in (7) and (8), respectively. In [11], three approaches are proposed to approximate the PME and PVE and they are the fully-exponential Laplace approximation method [12]–[15], Monte Carlo importance sampling, and a simple Laplace approximation method.

IV. SIMPLE LAPLACE APPROXIMATION METHOD

The simple Laplace approximation method for estimating (7) and (8) is based on the maximum a posteriori (MAP) estimation. The simple Laplace approximation method will be called MAP-GP in the remainder of this paper. The following proposition from [11] demonstrates that MAP-GP can provide good estimates.

Proposition 1: Let $\hat{\mathbf{x}}$ be an asymptotic mode of order $O(n^{-1})$ for $-h(x)$, such that

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in X} p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{x}}) \quad (9)$$

$$h(\mathbf{x}) = -\frac{1}{n} \log(p(\bar{\mathbf{y}} | \mathbf{x}) p(\mathbf{x} | \bar{\mathbf{x}})). \quad (10)$$

Assume that $\{h, \hat{\mathbf{x}}\}$ satisfies the regularity conditions described in [11]. Consider the following approximations for $\mathbb{E}(y_* | D)$ and $\text{var}(y_* | D)$

$$\hat{\mathbb{E}}(y_* | D) = k^T(x_*, \hat{\mathbf{x}}) (k(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + \sigma_w^2 I)^{-1} \bar{\mathbf{y}}, \quad (11)$$

$$\widehat{\text{var}}(y_* | D) = \sigma_f^2 - k^T(x_*, \hat{\mathbf{x}}) (k(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + \sigma_w^2 I)^{-1} k(x_*, \hat{\mathbf{x}}), \quad (12)$$

where $k(\hat{\mathbf{x}}, \hat{\mathbf{x}}) \in \mathbb{R}^{n \times n}$ and $k(x_*, \hat{\mathbf{x}}) \in \mathbb{R}^n$ are covariance matrices obtained using $\hat{\mathbf{x}}$. Then, we have the following orders of errors.

$$\hat{\mathbb{E}}(y_* | D) = \mathbb{E}(y_* | D) + O(n^{-1})$$

$$\widehat{\text{var}}(y_* | D) = \text{var}(y_* | D) + O(n^{-1}).$$

Remark 1: Note that the MAP-GP predictive mean and variance in (11) and (12) take the same form as the original predictive mean and variance, cf. (5) and (6), but the MAP estimator $\hat{\mathbf{x}}$ from (9) is used.

V. DISTRIBUTED GPR ALGORITHM

In this section, we introduce a distributed algorithm for an individual agent (sensor) to estimate an environmental parameter of the surveillance region S only by exchanging local information within r -disk neighbors. Consider a mobile sensor network consisting of q mobile sensing agents distributed in S . This distributed approach can be implemented for a class of kernel functions $k(\cdot, \cdot)$ that have compact supports. The following kernel function is considered in the development of a distributed algorithm.

$$k(x, x') = \sigma_f^2 \lambda \left(\frac{\|x - x'\|}{r} \right), \quad (13)$$

where

$$\lambda(h) = \begin{cases} (1-h) \cos(\pi h) + \frac{1}{\pi} \sin(\pi h), & \text{if } h \leq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Notice that the kernel function k in (13) has a compact support, i.e., $k_{ij} = k(x^{(i)}, x^{(j)})$ is non-zero if and only if

$r_{ij} = \|x^{(i)} - x^{(j)}\|$ is less than the support r and, similarly, $k_i = K(x^{(i)}, x_*)$ is non-zero if and only if $d_{i*} = \|x^{(i)} - x_*\|$ is less than the support r . Consider a case in which each agent in a sensor network can only communicate with other agents within a limited communication range of R . In addition, we assume that there exists no central station.

The index of the robotic sensors is denoted by $I = \{1, \dots, q\}$ and the position of agent i is denoted by $x^{(i)}$. Agent i can only communicate with its neighbors in a limited range of R . The adjacency matrix Q indicates whether two agents are neighbors or not. If the element in the i -th row and j -th column of Q , i.e., $Q_{ij} = 1$, then agent i and agent j are neighbors and they have a communication link, and if $Q_{ij} = 0$, then they are not neighbors.

$$Q_{ij} = \begin{cases} 1, & \|x^{(i)} - x^{(j)}\| \leq R \text{ and } i \neq j \\ 0, & \text{otherwise,} \end{cases}$$

where R is the communication range between neighbors.

The assumptions made for the resource-constrained mobile sensor networks are listed as follows.

Assumption 1: The radius r of the support of the kernel function (13) satisfies that $0 < r < R$.

Assumption 2: Agent i can only communicate with neighbors in $N(i) = \{j \in I \mid Q_{ij} = 1\}$.

These assumptions indicate that the communication range must be longer than the measurement radius of each agent. If these assumptions are not held, the multi-agent system cannot reach a consensus due to an inability to make communication with each other.

A. Jacobi Over-Relaxation Method

Jacobi over-relaxation (JOR) is a method for solving $Ax = b$, where $A \in \mathbb{R}^{q \times q}$, and $x, b \in \mathbb{R}^q$ [16]. This method assumes that agent i knows the $[A]_{(i)} \in \mathbb{R}^q$ and b_i , and $a_{ij} = [A]_{ij} = 0$ if agent i and agent j are not neighbors which goes along with Assumption 2. The i -th element of solution $x = A^{-1}b$ can be obtained by the following iterative steps:

$$x_i^{(k+1)} = (1-h)x_i^{(k)} + \frac{h}{a_{ii}} \left(b_i - \sum_{j \in N_i} a_{ij} x_j^{(k)} \right), \quad (14)$$

where h is a constant controlling the speed of convergence.

The convergence properties of the iterative JOR algorithm with respect to eigenvalues of the matrix A is specified in [16]. Moreover, Udwardia et al. [17] proved that if $h < \frac{2}{q}$, the convergence of the JOR algorithm to the solution $x = A^{-1}b$ is guaranteed from any initial point x_{init} , where A is a symmetric, positive-definite matrix.

Remark 2: Note that, the solution of the JOR method approaches to the solution $x = A^{-1}b$ without requiring the computation of an inverse of A . As shown in Section II, a drawback of Gaussian process regression is $O(N^3)$ complexity of calculating $(K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1}$. This complexity can be relieved if the JOR method is applied in a distributed manner.

B. Discrete-Time Average Consensus Method

The discrete-time average consensus (DAC) method is used to compute the arithmetic mean of elements of a vector [18]. Let

$$\mathbf{c} = [c_1 \ c_2 \ \dots \ c_q]^T \in \mathbb{R}^q. \quad (15)$$

If the graph is connected and agent i knows the i -th element of vector \mathbf{c} , the arithmetic mean of \mathbf{c} can be computed by iterating

$$x_i^{(k+1)} = x_i^{(k)} + \epsilon \sum_{j \in N_i} a_{ij} (x_j^{(k)} - x_i^{(k)}), \quad (16)$$

with an initial condition $x_i^0 = c_i$, where $a_{ij} = 1$ if and only if agent i and agent j are connected. It is proven that if ϵ satisfies

$$0 < \epsilon < \frac{1}{\max_i (\sum_{j \neq i} a_{ij})}, \quad (17)$$

then the DAC algorithm converges to the solution [19].

After the convergence, every node in the network converges to the arithmetic mean of vector \mathbf{c} , i.e., $\frac{1}{q} \sum_{i=1}^q c_i$.

C. Distributed Maximum a Posteriori Mode Estimator

In order to find the asymptotic mode $\hat{\mathbf{x}}$ in a distributed manner, the Jacobi over-relaxation method introduced in Section V-A is used.

Agent i only knows the i -th row and the i -th column of K and $\partial K / \partial x^{(i)}$. In addition, note that for $\partial K / \partial x^{(i)}$, the elements other than the i -th row and the i -th column are zeros. Define Γ and B_l as follows:

$$\Gamma = (K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1} \bar{\mathbf{y}} \quad (18)$$

$$B_l = \frac{1}{2} (K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1} \frac{\partial K(\mathbf{x}, \mathbf{x})}{\partial x^{(l)}}. \quad (19)$$

Agent i knows the i -th row of $(K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)$ and the i -th element of $\bar{\mathbf{y}}$. So $\Gamma^{(i)}$, the i -th element of Γ , can be computed by applying the JOR method over $(K(\mathbf{x}, \mathbf{x}) + \sigma_w^2 I)^{-1} \bar{\mathbf{y}}$ based on the following recursion:

$$\Gamma^{(i)}(k+1) = (1-\alpha)\Gamma^{(i)}(k) + \frac{\alpha}{\sigma_f^2 + \sigma_w^2} \times \left(\bar{y}^{(i)} - \sum_{j \in N(i)} k(\hat{x}^{(i)}, \hat{x}^{(j)}) \Gamma^{(j)}(k) \right) \quad (20)$$

Similarly, agent i knows $[(K + \sigma_w^2 I)]_{(i)}$, the i -th row of $(K + \sigma_w^2 I)$, and the i -th element of $\bar{\mathbf{y}}$. So $\Gamma^{(i)}$, the i -th element of Γ , can be computed by applying the JOR method. However, agent i needs to receive $\frac{\partial K}{\partial x^{(l)}}$ from agent l if they are connected or use zero instead if they are not connected. In other words, agent i will

$$\begin{cases} \text{get } \left[\frac{\partial K}{\partial x^{(l)}} \right]_{(i)} & \text{from agent } l \text{ if } i \in N(l) \\ \left[\frac{\partial K}{\partial x^{(l)}} \right]_{(i)} = 0 & l \neq i \text{ and } i \notin N(l) \end{cases}$$

and

$$[B_l(k+1)]_{(i)} = (1-\alpha)[B_l(k)]_{(i)} + \frac{\alpha}{\sigma_f^2 + \sigma_w^2} \times \left(\left[\frac{\partial K}{\partial x^{(l)}} \right]_{(i)} \Big|_{\mathbf{x}=\hat{\mathbf{x}}} - \sum_{j \in N(i)} k(\hat{\mathbf{x}}^{(i)}, \hat{\mathbf{x}}^{(j)}) [B_l(k)]_{(j)} \right) \quad (21)$$

for $k \in \mathbb{Z}_{\geq 0}$ and $i \in 1, 2, \dots, q$, where $\alpha \in (0, 2\lambda_{\min}(K + \sigma_w^2 I))$. At the end of JOR iterations, agent i knows $\Gamma^{(i)}$ and $[B_l]_{(i)}$.

Proposition 2: Given $B_i^{(j)}$, the j -th diagonal element of B_i , $\frac{\partial h}{\partial x^{(i)}}$ calculated for agent i is

$$\frac{\partial h}{\partial x^{(i)}} = \frac{x^{(i)} - \bar{x}^{(i)}}{n\sigma_v^2} - \frac{1}{n} \sum_{j \in N(i)} \left(\Gamma^{(j)} \frac{\partial k(x^{(j)}, x^{(i)})}{\partial x^{(i)}} \Gamma^{(i)} - B_i^{(j)} \right). \quad (22)$$

Proof: h given by (10) can be expressed as

$$h = \frac{d}{2} \log(2\pi\sigma_v^2) + \frac{\sum_{i=1}^n (x^{(i)} - \bar{x}^{(i)})^2}{2m\sigma_v^2} + \frac{1}{2} \log(2\pi) + \frac{1}{2n} \log |K + \sigma_w^2 I| + \frac{1}{2m} \bar{\mathbf{y}}^T (K + \sigma_w^2 I) \bar{\mathbf{y}}.$$

Now take a derivative respect to $x^{(i)}$ to get

$$\frac{\partial h}{\partial x^{(i)}} = \frac{x^{(i)} - \bar{x}^{(i)}}{n\sigma_v^2} + \frac{1}{2n} \text{tr} \left(\frac{\partial K}{\partial x^{(i)}} (K + \sigma_w^2 I)^{-1} \right) + \frac{1}{2n} \bar{\mathbf{y}}^T (K + \sigma_w^2 I)^{-1} \frac{\partial K}{\partial x^{(i)}} (K + \sigma_w^2 I)^{-1} \bar{\mathbf{y}}.$$

Using (18) and (19), we get

$$\frac{\partial h}{\partial x^{(i)}} = \frac{x^{(i)} - \bar{x}^{(i)}}{n\sigma_v^2} + \frac{1}{n} \text{tr}(B_i^T) - \frac{1}{2n} \Gamma^T \frac{\partial K}{\partial x^{(i)}} \Gamma. \quad (23)$$

It is known that $\frac{\partial K}{\partial x^{(i)}}$ is a sparse matrix with only non-zero elements on the i -th row and column for the neighbors of i . Rewriting the above equation in the summation form leads to (22). ■

Finally using the recursive gradient method given by (24), agents can update their modes in a distributed manner, where $\bar{x}^{(i)}$ can be used as an initial condition. The overall algorithm is summarized in Algorithm 1.

Algorithm 1 Distributed Algorithm for Finding $\hat{\mathbf{x}}$ for agent i

Given Initial position $\bar{\mathbf{x}}$, corrupted measurement \bar{y} and ϵ satisfying (17).

begin $\forall_i \hat{x}_i^{(0)} = \bar{x}_i, \Gamma_i^{(0)} = 0, B_i^{(0)} = \mathbf{0} \in \mathbb{R}^{n \times n}$

repeat

repeat

1. Update Γ_i using (20).
2. Update B_i using (21).

until Γ_i and B_i converges.

1. Compute $\frac{\partial h}{\partial x^{(i)}}$ using (23).
2. Update \hat{x}_i using (24).

until \hat{x}_i converges.

$$\hat{x}^{(i)}(t+1) = \hat{x}^{(i)}(t) + \gamma \frac{\partial h}{\partial x^{(i)}} \Big|_{x^{(j)} = \hat{x}^{(j)}(t), j \in N(i)}. \quad (24)$$

D. Distributed posterior mean and variance estimators

Let us consider the first-order approximation of the PME estimator in (11). The elements of the covariance matrices $k(\hat{x}^{(i)}, x_*)$ and $k(\hat{x}^{(i)}, \hat{x}^{(j)})$ can be calculated using (13) and the converged asymptotic mode $\hat{\mathbf{x}}$ of $-h$ can be found.

Proposition 3: If a sensor network is connected, every agent can compute $\hat{\mathbb{E}}(y_*|D)$ with the DAC method by exchanging only local information using:

$$\theta_i(k+1) = \theta_i(k) + \epsilon \sum_{j \in N(i)} (\theta_j(k) - \theta_i(k)), \quad (25)$$

where $\theta_i(0) = k^{(i)} \Gamma^{(i)}$.

Proof: By Proposition 1 and (18), PME in (11) can be represented as follows:

$$\begin{aligned} \hat{\mathbb{E}}(y_*|D) &= k^T(x_*, \hat{\mathbf{x}}) (K(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + \sigma_w^2 I)^{-1} \bar{\mathbf{y}} \\ &= k^T(x_*, \hat{\mathbf{x}}) \Gamma. \end{aligned} \quad (26)$$

The solution of (26) can be solved using DAC since (26) can be matched to (15). Recall that (26) is a scalar value and can be represented as $\sum_{i=1}^q k^T(x_*, \hat{\mathbf{x}}^{(i)}) \Gamma^{(i)}$ and the i -th agent knows both $k^T(x_*, \hat{\mathbf{x}}^{(i)})$ and $\Gamma^{(i)}$.

From Section V-B, we know that θ_i converges to $\frac{1}{q} \hat{\mathbb{E}}[y_*|D]$ when $0 < \epsilon < \frac{1}{\max_i \sum_{i \neq j} Q_{ij}}$ and each agent can easily compute $\hat{\mathbb{E}}[y_*|D]$ by $q \times \theta_i$. ■

The predictive variance can be approximated similarly with

$$\widehat{\text{var}}(y_*|D) = \sigma_f^2 - k^T (K + \sigma_w^2 I)^{-1} k. \quad (27)$$

First, suppose that

$$\Phi = (K(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + \sigma_w^2 I)^{-1} k(x_*, \hat{\mathbf{x}}). \quad (28)$$

Agent i can calculate the i -th element of (28) by the following recursion (29).

$$\begin{aligned} \Phi^{(i)}(k+1) &= (1-\alpha)\Phi^{(i)}(k) + \frac{\alpha}{\sigma_f^2 + \sigma_w^2} \\ &\times \left(\bar{y}^{(i)} - \sum_{j \in N(i)} k(\hat{x}^{(i)}, \hat{x}^{(j)}) \Phi^{(j)}(k) \right). \end{aligned} \quad (29)$$

Algorithm 2 Distributed Algorithm for Mean Field Prediction for agent i

Given Mode $\hat{\mathbf{x}}$, corrupted measurement \bar{y} , unmeasured position x_* and converged Γ .

begin $\theta_i^{(0)} = k(x_*, \hat{x}_i) \Gamma^{(i)}$

repeat

1. Update θ_i using (25).

until θ_i converges.

output $\hat{\mathbb{E}}[y_*|D] = q \times \theta_i$

Algorithm 3 Distributed Algorithm for Variance Field Prediction for agent i

Given Mode $\hat{\mathbf{x}}$, unmeasured position x_* and converged Γ .
begin Set arbitrary $\Phi^{(i)}$.
repeat
 1. Update $\Phi^{(i)}$ using (29).
until $\Phi^{(i)}$ converges.
begin $\theta_i^{(0)} = k(x_*, \hat{x}_i)\Phi^{(i)}$
repeat
 1. Update θ_i using (30).
until θ_i converges.
output $\widehat{\text{var}}(y_*|D) = q \times \theta_i$

Once JOR converges, the error variance can be computed using following recursion based on the DAC method

$$\theta_i(k+1) = \theta_i(k) + \epsilon \sum_{j \in N(i)} (\theta_j(k) - \theta_i(k)), \quad (30)$$

where

$$\theta_i(0) = k^T(x_*, \hat{\mathbf{x}}^{(i)})\Phi^{(i)}. \quad (31)$$

Without loss of generality, the iterative DAC solution in (25) and (30) can be extended to vector formulation, where $\theta_i \in \mathbb{R}^k$. The overall algorithms are summarized in Algorithm 2 and 3.

VI. SIMULATION RESULTS

In this section, we perform a number of numerical simulations to validate the prediction performance of the proposed distributed GPR algorithm. For simulation, we have randomly generated ten reference fields from a Gaussian process with a covariance function given in (13). We then compare the predicted fields using three different algorithms against the reference field. Three algorithms used in simulations are the quick and dirty solution (QDS) given in (5), the centralized solution using the simple Laplace approximation using the MAP estimator (MAP-GP) in (11), and the proposed distributed Gaussian process regression (D-GP).

We assume that there are ten sensing agents ($q = 10$). For each reference field, sensing agents are randomly located and each agent only knows a noisy position of itself. The variance of the position uncertainty is set to $\sigma_v^2 = 1$. Each agent then makes a noisy measurement from the reference field. The collection of all measurements by all agents is denoted by $D' = \{\bar{\mathbf{x}}, \bar{\mathbf{y}}\}$. The hyper-parameter σ_f^2 of the kernel function in (13) and the variance of the measurement noise σ_w^2 are estimated by maximizing the likelihood [3], i.e.,

$$\{\widehat{\sigma}_f^2, \widehat{\sigma}_w^2\} = \arg \max \log P(\bar{\mathbf{y}}|\bar{\mathbf{x}}, \sigma_f^2, \sigma_w^2),$$

given $D' = \{\bar{\mathbf{x}}, \bar{\mathbf{y}}\}$. Then we have performed the reconstruction of the entire field from D' using three algorithms. For more accurate validation, we have repeated the above procedure ten times for each reference field, resulting ten independent runs for each reference field. For each reference field, we have computed the mean and variance of the reconstruction error from ten independent runs.

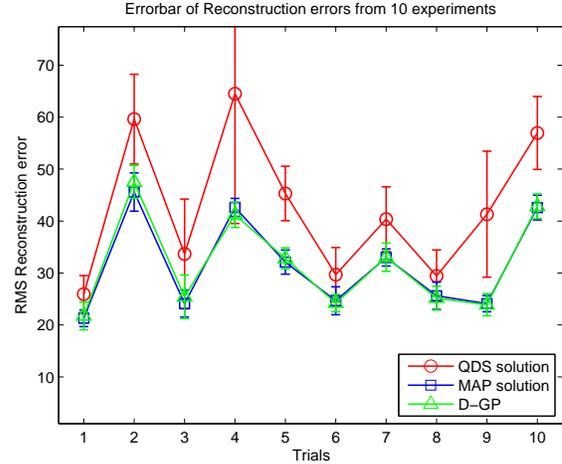


Fig. 1. Reconstruction errors of three algorithms (QDS, MAP-GP, and D-GP) for ten different scenarios. Error bars indicate one standard deviation from ten independent runs for each scenario.

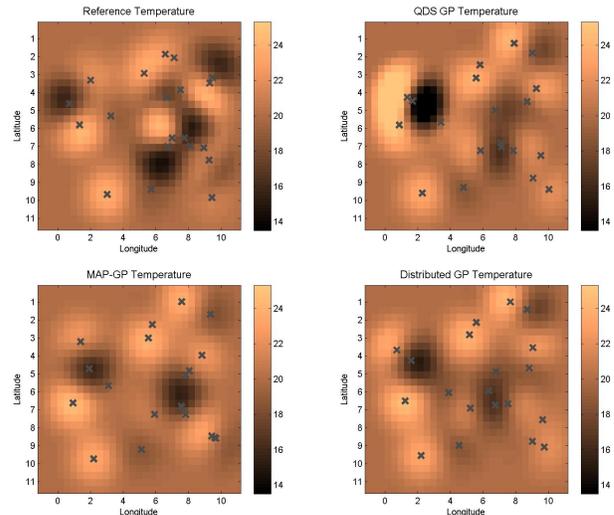


Fig. 2. An example of a reference field and fields reconstructed by three algorithms (QDS, MAP-GP, and D-GP). The reference field is shown in the upper left corner and, clockwise from the top, fields reconstructed using QDS, D-GP, and MAP-GP. The MAP estimates of sensor positions are marked by gray cross marks.

The mean and variance of the root mean squared (RMS) errors between the predicted field and the reference field are shown in Figure 1. As expected, the centralized algorithm MAP-GP shows the best performance in terms of the reconstruction error, followed by D-GP and QDS. However, the mean of the reconstruction error of D-GP is comparable to that of MAP-GP, demonstrating the performance of the proposed distributed algorithm. One example of the reference field and the predicted fields using three algorithms are shown in Figure 2. Figure 3 demonstrates the convergence of JOR used in our algorithm.

Since our objective function (10) is non-convex, we used a gradient-based nonlinear optimization method for both MAP-GP and D-GP and this can explain the difference between

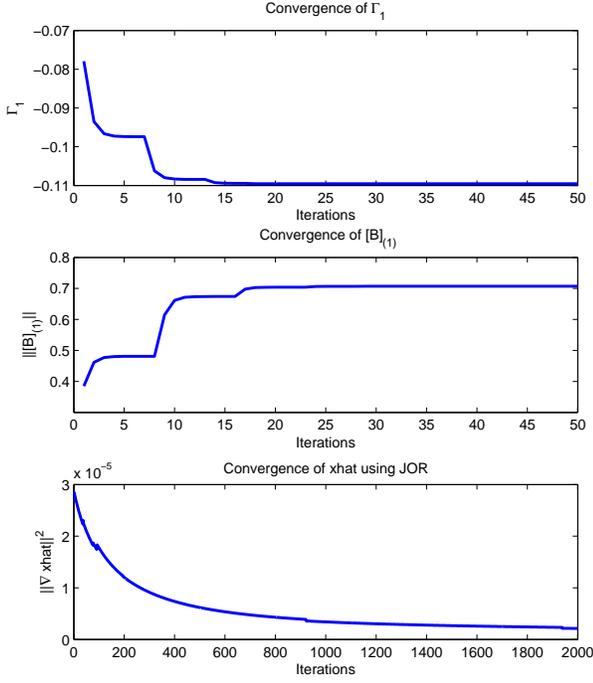


Fig. 3. Convergence of the JOR method. The upper and middle figure indicate Γ_1 and $[B]_{(1)}$ of agent 1, respectively. The bottom figure shows the norm of the gradient of \hat{x} , i.e., the solution of the MAP estimator in (9).

| Method | | | | | | |
|---------|-------|----------|--------------|----------|--------------|----------|
| trial | QDS | | MAP-GP | | D-GP | |
| | mean | variance | mean | variance | mean | variance |
| 1 | 25.93 | 3.57 | 21.28 | 1.60 | 21.73 | 2.67 |
| 2 | 59.63 | 8.64 | 45.59 | 3.68 | 47.54 | 3.16 |
| 3 | 33.67 | 10.56 | 24.11 | 2.62 | 25.41 | 4.18 |
| 4 | 64.52 | 25.00 | 42.63 | 1.74 | 41.21 | 2.44 |
| 5 | 45.31 | 5.25 | 32.09 | 2.32 | 32.80 | 4.03 |
| 6 | 29.67 | 5.24 | 24.65 | 2.68 | 24.18 | 1.66 |
| 7 | 40.37 | 6.20 | 32.98 | 1.65 | 33.03 | 2.72 |
| 8 | 29.44 | 5.00 | 25.62 | 2.63 | 25.15 | 2.30 |
| 9 | 41.31 | 12.14 | 24.09 | 1.55 | 23.88 | 2.12 |
| 10 | 56.96 | 7.03 | 42.60 | 2.38 | 42.85 | 2.36 |
| average | 42.68 | 8.86 | 31.56 | 2.28 | 31.78 | 2.76 |

TABLE I

MEAN AND VARIANCE OF RECONSTRUCTION ERRORS OF THREE ALGORITHMS: QDS: QUICK AND DIRTY SOLUTION, MAP-GP: CENTRALIZED SOLUTION, AND D-GP: PROPOSED APPROACH.

the reconstruction errors of these two algorithms as shown in Figure 1. However, as shown in Table I, D-GP has outperformed QDS in terms of the reconstruction error and its performance is comparable to MAP-GP in most cases.

VII. CONCLUSION

We have presented a distributed Gaussian process regression algorithm that can incorporate both localization and measurement noises in a Bayesian framework. The proposed distributed algorithm is constructed by combining Jacobi

over-relaxation (JOR) and discrete time average consensus (DAC) in a distributed manner and suitable for mobile sensor networks with limited communication ranges. At the same time, the proposed algorithm relieves the heavy computation burden of computing the inverse of a large covariance matrix required for Gaussian process regression. The performance of our proposed algorithm is evaluated extensively using a number of simulations. It has been demonstrated that the proposed distributed Gaussian process regression (D-GP) outperforms the quick-and-dirty solution (QDS) in most cases and its performance is comparable to the centralized solution.

REFERENCES

- [1] D. Ehrlich, J. Estes, and A. Singh, "Applications of NOAA-AVHRR 1 km data for environmental monitoring," *Remote Sensing*, vol. 15, no. 1, pp. 145–161, 1994.
- [2] D. Culler, D. Estrin, and M. Srivastava, "Guest editors' introduction: overview of sensor networks," *Computer*, pp. 41–49, 2004.
- [3] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 1.
- [4] J. Choi, J. Lee, and S. Oh, "Swarm intelligence for achieving the global maximum using spatio-temporal Gaussian processes," in *Proc. of the American Control Conference*, 2008, pp. 135–140.
- [5] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg, "Near-optimal sensor placements: Maximizing information while minimizing communication cost," in *Proc. of the International Conference on Information Processing in Sensor Networks*, 2006.
- [6] C. Williams and M. Seeger, "Using the Nystrom method to speed up kernel machines," in *Proc. of the Advances in Neural Information Processing Systems*, 2001.
- [7] N. Lawrence, M. Seeger, and R. Herbrich, "Fast sparse Gaussian process methods: The informative vector machine," in *Proc. of the Advances in Neural Information Processing Systems*, 2002.
- [8] Y. Xu, J. Choi, and S. Oh, "Mobile sensor network navigation using Gaussian processes with truncated observations," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1118–1131, 2011.
- [9] P. Oguz-Ekim, J. Gomes, J. Xavier, and P. Oliveira, "Robust localization of nodes and time-recursive tracking in sensor networks using noisy range measurements," *IEEE Transactions on Signal Processing*, vol. 59, no. 8, pp. 3930–3942, 2011.
- [10] R. Karlsson and F. Gustafsson, "Bayesian surface and underwater navigation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4204–4213, 2006.
- [11] M. Jadaliha, Y. Xu, J. Choi, N. Johnson, and W. Li, "Gaussian process regression for sensor networks under localization uncertainty," *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 223–237, 2013.
- [12] L. Tierney and J. Kadane, "Accurate approximations for posterior moments and marginal densities," *Journal of the American Statistical Association*, vol. 81, no. 393, pp. 82–86, 1986.
- [13] L. Tierney, R. Kass, and J. B. Kadane, "Fully exponential Laplace approximations to expectations and variances of nonpositive functions," *Journal of the American Statistical Association*, vol. 84, no. 407, 1989.
- [14] Y. Miyata, "Fully exponential Laplace approximations using asymptotic modes," *Journal of the American Statistical Association*, vol. 99, no. 468, pp. 1037–1049, 2004.
- [15] —, "Laplace approximations to means and variances with asymptotic modes," *Journal of Statistical Planning and Inference*, vol. 140, no. 2, pp. 382–392, 2010.
- [16] D. Bertsekas and J. Tsitsiklis, *Parallel and distributed computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [17] F. Udawadia, "Some convergence results related to the JOR iterative method for symmetric, positive-definite matrices," *Applied Mathematics and Computation*, vol. 47, no. 1, pp. 37–45, 1992.
- [18] J. Cortés, "Distributed kriged Kalman filter for spatial estimation," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2816–2827, 2009.
- [19] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," in *Proc. of the IEEE*, vol. 95, no. 1, 2007, pp. 215–233.