

Decentralized Learning With Limited Communications for Multi-robot Coverage of Unknown Spatial Fields

Kensuke Nakamura*, María Santos*, and Naomi Ehrich Leonard

Abstract—This paper presents an algorithm for a team of mobile robots to simultaneously learn a spatial field over a domain and spatially distribute themselves to optimally cover it. Drawing from previous approaches that estimate the spatial field through a centralized Gaussian process, this work leverages the spatial structure of the coverage problem and presents a decentralized strategy where samples are aggregated locally by establishing communications through the boundaries of a Voronoi partition. We present an algorithm whereby each robot runs a local Gaussian process calculated from its own measurements and those provided by its Voronoi neighbors, which are incorporated into the individual robot's Gaussian process only if they provide sufficiently novel information. The performance of the algorithm is evaluated in simulation and compared with centralized approaches.

I. INTRODUCTION

The problem of positioning a team of robots over a domain to optimally monitor its environmental features is referred to as the spatial coverage problem. Standard methods that tackle this problem (e.g. [1], [2]) provide control laws for robots to coordinate on their spatial distribution, concentrating around areas of high interest in the domain and spreading out over areas of low interest. Depending on the application, the multi-robot team may need to monitor different phenomena, ranging from population density for wireless coverage [3] to radiation levels in the event of a nuclear disaster [4], or weather conditions in environmental forecasting [5].

The prevalence of the environmental feature(s) to be monitored over a domain can be encoded as a spatial field. Frequently referred to as the *density function* in the context of coverage, the spatial field gives each point in the domain a scalar value that represents the relative importance of the point, with higher values of the density corresponding to regions of higher interest. Depending on the robotic application, different functions have been used in the problem of coverage including densities representing static fields [1], [2], dynamic densities that reflect changes of the field over time [6], [7], [8], [9], collections of densities that represent the presence of heterogeneous events in the environment [10], [11], or nonuniform densities that generalize the problem to domains with non-Euclidean distances [12], [13]. A common assumption in these works, however, is that the multi-robot

team knows the density function a priori, and thus is solely concerned with optimally distributing themselves to monitor the domain.

This paper addresses the problem of multi-robot spatial coverage when the density function over the domain is initially unknown. To address this problem, the robots perform two different, albeit complementary, tasks: they take samples needed to estimate the density function and they position themselves in the domain to effectively monitor it. These tasks may be completed sequentially [4] or concurrently [14], [15]. However, while the coverage task can be executed in a decentralized fashion by leveraging the well-known Lloyd's algorithm [16], learning the density function typically requires a costly central computation that processes all the samples from the team to produce an estimate [14], [17]. Instead, we propose a fully decentralized approach to simultaneously learn and cover a spatial field. Building on the Voronoi Estimation Coverage (VEC) algorithm from [15], which uses the Voronoi partition to limit the sampling area of each robot but requires a centralized computation to estimate the spatial field, we design a completely decentralized learning method where each robot aggregates its samples and its Voronoi neighbors' samples to produce its own estimate of the density function.

The paper is organized as follows. This remainder of this section includes a summary of related literature on multi-robot learning for spatial coverage. Section II introduces the problem of coverage of unknown spatial fields and gives a brief overview of the VEC algorithm in [15]. The proposed algorithm, which leverages the Voronoi partition from the coverage objective for the purposes of reducing sampling, is outlined in Section III, with Section III-B focusing on how data is filtered to reduce the amount of communications needed. Section IV presents a series of simulations that validate the approach and compare it to its centralized counterparts. Conclusions are detailed in Section V.

A. Related Work

Cooperative estimation of unknown spatial fields has been the subject of study in a variety of multi-robot system applications, e.g. [18], [19], [20], [21], [22]. In the context of spatial coverage, different strategies have been proposed to combine optimal coverage of the domain and sampling of the unknown spatial field in the control policy of the robots.

Some works optimize the robots' path only for the coverage objective, and the sampling locations occur as a direct consequence of the coverage of the current estimate of the field [17], [23], [14]. Other works opt for a sequential

This work was supported by the Office of Naval Research grant N00014-19-1-2556, Army Research Office grant W911NF-18-1-0325 and Princeton University's School of Engineering and Applied Science through the generosity of Lydia and William Addy '82.

* Equal contribution.

K. Nakamura, M. Santos, and N. E. Leonard are with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, 08544 USA; k.nakamura@princeton.edu, maria.santos@princeton.edu, naomi@princeton.edu

execution of both tasks, typically having the robots estimate the spatial field over the environment and then monitor it accordingly [4], [24].

The simultaneous estimation and coverage of an unknown spatial field has been addressed in [3], [15]. In [3], the density function is modeled using basis function approximation. However, this approach limits the types of functions that can be learned as it can lead to very large errors or instabilities when the true density function is not well represented by the chosen basis functions. While the authors resolved this instability issue in [4], the execution of the learning and coverage tasks became sequential. In [15] an algorithm was proposed to concurrently learn and cover the spatial field. However, the estimate of the function was modeled as a Gaussian process that aggregates all the samples collected by the robots in the team, and thus, the algorithm relies on a centralized computation. In addition, if the team collects P samples, the Bayesian update of the Gaussian process requires an inversion of a $P \times P$ matrix, which scales poorly with the number of measurements ($\mathcal{O}(P^3)$). In this paper, we maintain the generality afforded by Gaussian processes but leverage the spatial structure given by the Voronoi partition to decentralize the learning objective. We propose a fully decentralized algorithm for simultaneous learning and coverage of spatial fields.

II. COVERAGE OF UNKNOWN SPATIAL FIELDS

The multi-robot coverage problem deals with how to position a team of mobile robots such that they can optimally monitor the environmental features in a domain [1], [2]. In this context, a resource being more prevalent in an area or an event being more likely may impact the relative importance of points in the domain to be covered. Differences in the relative importance of points can be modeled with a density function over the domain that assigns higher values to points that are to be monitored more closely.

Consider a team of M mobile robots that are to cover a closed and convex domain $D \subset \mathbb{R}^d$ (e.g., $d = 2$ for ground robots, $d = 3$ for aerial robots) with respect to an initially unknown spatial field $f^* : D \rightarrow \mathbb{R}$. For the coverage task, a typical approach is to partition the domain into regions of dominance such that Robot i , with position $x_i \in D$, is in charge of monitoring the set V_i of points in the domain that are closest to it:

$$V_i(\mathbf{x}) = \{q \in D \mid \|q - x_i\| \leq \|q - x_j\|, \forall i \neq j\}, \quad (1)$$

with $\mathbf{x} = [x_1^\top, \dots, x_M^\top]^\top$ the stacked positions of the robots and $i, j \in \{1, \dots, M\}$. Note that Robot i can compute V_i from the intersection of the bisectors of each segment $x_i x_j$, j a Delaunay neighbor of i . Assuming that the sensing radius of each robot is big enough, each robot can compute its Delaunay neighbors locally, e.g. as in [25].

With this partition, the quality of coverage of Robot i with respect to the density function f^* is inversely related to cost

$$h_i(\mathbf{x}, f^*(\cdot)) = \int_{V_i(\mathbf{x})} \|q - x_i\|^2 f^*(q) dq.$$

The squared distance models the decaying quality of sensor data at point q with increasing distance from the robot at x_i . The performance of the multi-robot team can be evaluated as the sum of the individual robot costs:

$$H(\mathbf{x}, f^*(\cdot)) = \sum_{i=1}^M h_i(\mathbf{x}, f^*(\cdot)) = \sum_{i=1}^M \int_{V_i(\mathbf{x})} \|q - x_i\|^2 f^*(q) dq, \quad (2)$$

with a better coverage corresponding to lower values of H .

Finding a global minimum to (2) is an NP-hard problem [23]. However, one configuration that is known to locally minimize the cost (2) is the Centroidal Voronoi Tessellation (CVT) [26]. In a CVT, the position of each robot corresponds to the center of mass of its Voronoi cell, given by

$$c_i(\mathbf{x}, f^*(\cdot)) = \frac{\int_{V_i(\mathbf{x})} q f^*(q) dq}{\int_{V_i(\mathbf{x})} f^*(q) dq}. \quad (3)$$

The multi-robot team can thus minimize (2) by making each robot move towards its center of mass [1], which corresponds to a continuous version of Lloyd's algorithm [16],

$$\dot{x}_i = \kappa(c_i(\mathbf{x}, f^*(\cdot)) - x_i), \quad \kappa > 0, \quad \forall i \in \{1, \dots, M\}. \quad (4)$$

A strong assumption behind the control law in (4) is that each robot knows the spatial field $f^*(\cdot)$ and, thus, is able to compute the center of mass in (3). The recently developed Voronoi Estimation Coverage (VEC) algorithm [15] leverages Gaussian process regression to learn the estimate of the true spatial field. The estimate $f(\cdot)$ of the true function $f^*(\cdot)$ is modeled as a Gaussian process

$$f(\cdot) \sim GP(\mu(\cdot; \rho), k(\cdot, \cdot; \tau)),$$

with $\mu : D \rightarrow \mathbb{R}$ and $k : D \times D \rightarrow \mathbb{R}$ the mean and covariance functions, respectively, and ρ and τ , the hyperparameters. Adopting a linear mean and a squared-exponential covariance function as in [15], we have

$$\mu(x; \rho) = \rho^\top x, \quad k(x, x'; \tau) = \exp\left(-\frac{\|x - x'\|^2}{2\tau^2}\right). \quad (5)$$

The mean and covariance in (5), characterized by ρ and τ , can be approximated by maximum-likelihood estimation from noisy measurements of the density function [27]. Robot i takes a noisy scalar measurement $y_i^{(n)}$ at iteration n :

$$y_i^{(n)} = f^*(x_i) + \epsilon, \quad (6)$$

with $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$, where σ_n^2 models the variance in the measurement. Denote by N the total number of iterations, with $n \in \{1, \dots, N\}$ the time iteration. The set of points sampled in the domain D at time n is denoted by $\mathbf{x}^{(n)} = [x_1^{(n)\top}, \dots, x_M^{(n)\top}]^\top \in \mathbb{R}^{Md}$ and the corresponding measurements are $\mathbf{y}^{(n)} = [y_1^{(n)}, \dots, y_N^{(n)}]^\top \in \mathbb{R}^M$. We let $\mathbf{x}^{(1:n)}$ and $\mathbf{y}^{(1:n)}$ refer to all the sampled locations and measurements taken from iteration 1 to n .

The centralized maximum likelihood estimate for a Gaussian process is known in closed form (see, e.g. [27]) as

$$\begin{aligned} (\rho^{(n)}, \tau^{(n)}) &= \arg \min_{\rho, \tau} \{ \log |\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau) + \sigma_m^2 \mathbf{I}| + \\ &(\boldsymbol{\mu}^{(1:n)}(\rho) - \mathbf{y}^{(1:n)})^\top (\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau) + \sigma_m^2 \mathbf{I})^{-1} (\boldsymbol{\mu}^{(1:n)}(\rho) - \mathbf{y}^{(1:n)}) \}, \end{aligned} \quad (7)$$

with \mathbf{I} the $Mn \times Mn$ identity matrix, and $\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau) \in \mathbb{R}^{Mn \times Mn}$ the covariance matrix with entries $k(x_i^{(j)}, x_{i'}^{(j')}; \tau)$ for $i, i' \in \{1, \dots, M\}$ and $j, j' \in \{1, \dots, n\}$. Finally, $\boldsymbol{\mu}^{(1:n)}(\rho)$ represents the mean output of the GP.

With the estimate of the hyperparameters, the robot team reconstructs the mean and covariance functions as

$$\begin{aligned} \mu^{(n)}(x) &= \mu(x; \rho^{(n)}) + \\ \mathbf{k}(x; \mathbf{x}^{(1:n)})^\top (\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau^{(n)}) + \sigma_m^2 \mathbf{I})^{-1} (\mathbf{y}^{(1:n)} - \boldsymbol{\mu}^{(1:n)}(\rho^{(n)})), \\ k^{(n)}(x, x') &= k(x, x'; \tau^{(n)}) - \\ \mathbf{k}(x; \mathbf{x}^{(1:n)})^\top (\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau^{(n)}) + \sigma_m^2 \mathbf{I})^{-1} \mathbf{k}(x'; \mathbf{x}^{(1:n)}), \end{aligned}$$

where $\mathbf{k}(x; \mathbf{x}^{(1:n)}) \in \mathbb{R}^{Mn}$ is the vector with entries $k(x_i^{(j)}, x; \tau^{(n)})$ for $i \in \{1, \dots, M\}$ and $j \in \{1, \dots, n\}$. The reconstructed covariance function is then used to estimate the standard deviation

$$\sigma^{(n)}(x) = \sqrt{k^{(n)}(x, x)}. \quad (8)$$

Note the inversion of the matrix $(\mathbf{K}_{\mathbf{x}^{(1:n)}}(\tau) + \sigma^2 \mathbf{I})$ in (7). The matrix $\mathbf{K}_{\mathbf{x}^{(1:n)}}$ contains the kernel evaluation (5) between every pair $(x_i^j, x_{i'}^{j'})$, elements of $\mathbf{x}^{(1:n)}$ and is size $Mn \times Mn$. This computation requires each agent to send information to a centralized computer at each time step to perform this evaluation and scales $\mathcal{O}(M^3 n^3)$.

At each time step $n \in \{1, \dots, N\}$, the VEC algorithm updates the maximum likelihood estimate of $\rho^{(n)}$ and $\tau^{(n)}$. These are used to create an estimate of the density function as in [14], [15],

$$f^{(n)}(q) = \mu^{(n-1)}(q) - \sqrt{\beta^{(n)}} \sigma^{(n-1)}(q), \quad q \in D. \quad (9)$$

Here, $\{\beta^{(n)}\}_{n=1}^N$ is a non-negative non-decreasing sequence that leads to a small underestimate of the prediction. This encourages robots to be “pushed” to regions of higher importance, i.e. where the mean of estimate is higher [14]. Using this estimate, each robot selects a goal point which is a linear interpolation between the center of mass $c_i^{(n)}$ and a point $e_i^{(n)}$ that maximizes the uncertainty in (8) within its respective Voronoi cell. Regarding the center of mass, $c_i^{(n)}$, note: (i) $c_i^{(n)}$ is computed using (3) with the estimate $f^{(n)}(\cdot)$, and, thus, minimizes $H(\cdot, f^{(n)})$; and (ii), during iteration n , $c_i^{(n)}$ is not a static point—the center of mass evolves as a consequence of the movement the robots, which impacts the Voronoi partition of the domain. The controller thus sends each robot to their goal point:

$$\dot{x}_i = \kappa \left((1 - \gamma^{(n)}) c_i^{(n)} + \gamma^{(n)} e_i^{(n)} - x_i \right), \quad \kappa > 0. \quad (10)$$

The weighting of the linear interpolation is determined by $\{\gamma^{(n)}\}_{n=1}^N$, a nonnegative decreasing sequence less than

1. The weighting between the two points varies between timesteps, such that at later times more preference is given to points that minimize the cost (2). This linear interpolation is guaranteed to be within the domain and remain within each robot’s respective Voronoi cell due to the assumed convexity of the domain. When the equilibrium of (10) is achieved throughout the team, each robot collects its noisy sample of the environment, $y_i^{(n)}$. These are centrally aggregated to update the maximum likelihood estimate of the hyperparameters ρ and τ as in (7), which leads to a new iteration of the algorithm.

The VEC algorithm has shown remarkable performance for the problem of simultaneously learning and covering a spatial field [15]. One large drawback of the method, however, is the centralized computation of the hyperparameters ρ and τ . For a large team of robots collecting samples over a long time horizon, this computation can become prohibitively expensive as samples are continuously collected and also leaves the system vulnerable should anything happen to the centralized computer.

III. DECENTRALIZED LEARNING WITH VORONOI COMMUNICATIONS

The main contribution of this paper is an extension to the VEC algorithm that decentralizes the computation of the hyperparameters ρ and τ . Our proposed solution is motivated by the fact that the decision making of Robot i when computing $c_i^{(n)}$ and $e_i^{(n)}$ for the controller in (10) depends only on the estimate $f_i^{(n)}(\cdot)$ and uncertainty $\sigma^{(n)}(\cdot)$ within its cell V_i and its vicinity. The information far from the robot has no impact. By focusing only on the samples collected by a robot and those of its adjacent neighbors, a robot will still be able to recover an accurate estimate of the density function in its vicinity without incurring prohibitive computational costs.

We propose the Decentralized Voronoi Estimation Coverage (DVEC) algorithm (see Algorithm 1), where each robot has an individual estimate of the density function, $f_i(\cdot)$. This estimate is modeled as a Gaussian process regression,

$$f_i(\cdot) \sim GP(\mu_i(\cdot; \rho_i), k_i(\cdot, \cdot; \tau_i)),$$

with the mean μ_i and covariance k_i functions having the same structure as in (5). The novelty of this approach relies on the fact that each robot calculates this regression only with its own samples and samples received from its neighbors.

In this section we introduce two communication protocols that regulate the information exchange between neighboring robots, which we refer to as *naive* and *constrained communication*. In the first case, robots always exchange samples with their Delaunay neighbors. In the second, we introduce an energy-saving criterion that allows robots to discard samples that do not add substantial information to the density model, thus reducing the computational burden on the regression. In both cases there is an implicit assumption that communication for this is synchronous. This is due to our algorithm not allowing an individual robot to take a measurement until the entire team has converged to the equilibrium of (10).

A. Naive Communication Protocol

The DVEC algorithm takes advantage of the Voronoi tessellation (1) to create a communication network between the robots where they can exchange information about their samples of the environment. Using the spatial structure given by the partition, a naive way of aggregating samples for the Gaussian process is to let Robot i exchange information about its sample $\{x_i^{(n)}, y_i^{(n)}\}$ with its Delaunay neighbors. Let \mathcal{N}_i denote the set of neighbors of Robot i in the Voronoi partition, i.e. all robots j whose Voronoi cell V_j shares a boundary with V_i . If all robots share their samples with their Voronoi neighbors, the information available to Robot i at iteration n can be written as the sets,

$$\begin{aligned} \mathbf{x}_i^{(n)} &= \{x_i^{(n)}\} \cup \left(\bigcup_{j \in \mathcal{N}_i} \{x_j^{(n)}\} \right), \\ \mathbf{y}_i^{(n)} &= \{y_i^{(n)}\} \cup \left(\bigcup_{j \in \mathcal{N}_i} \{y_j^{(n)}\} \right). \end{aligned} \quad (11)$$

In a slight abuse of notation, we use the symbols $\mathbf{x}_i^{(n)}$ and $\mathbf{y}_i^{(n)}$ to indistinctly denote the sets in (11) and the column vectors whose entries correspond to the stacked elements of those sets (analogously to the definition of $\mathbf{x}^{(n)}$). The ordering of the elements in \mathbf{x}_i and \mathbf{y}_i does not affect the computation of the Gaussian process regression. Furthermore, we denote as $\mathbf{x}_i^{(1:n)}$ and $\mathbf{y}_i^{(1:n)}$ all the samples aggregated up to iteration n . Note that the neighbors of Robot i at each iteration may vary: in general, $\mathcal{N}_i^{(l)} \neq \mathcal{N}_i^{(l')}$ for two different iterations l, l' .

After communicating with its neighbors, Robot i computes its own estimate of the hyperparameters ρ_i and τ_i as in (7) using in this case $\mathbf{x}_i^{(1:n)}$ and $\mathbf{y}_i^{(1:n)}$. This allows Robot i to maintain its own estimate of the spatial field $f_i(\cdot)$ and the corresponding uncertainty $\sigma_i(\cdot)$, calculated as in (9) and (8), respectively. The controller then interpolates its goal point between the center of mass calculated with respect to its individual estimate, $c_i^{(n)}$, and the point that maximizes the uncertainty in the cell $e_i^{(n)}$. The pseudocode of the algorithm for an individual robot is outlined in Algorithm 1. The algorithm follows closely the steps of VEC, but DVEC considers all the learning computations at the individual level.

The naive communication strategy adopted in DVEC results in each robot having an estimate calculated based on the information gathered in its vicinity, but that has less data and, as a result, more uncertainty in regions of the domain that are far away from the robot. However, in most cases, the most relevant information for a robot to maximize its contribution to the coverage objective relates to the estimate of the density within the robot's Voronoi cell and its immediate surroundings. Thus, we do not anticipate severe performance losses by using a local estimate of the spatial field for the coverage task. In Section IV we verify this assumption in a series of simulations.

Algorithm 1: Decentralized Voronoi Estimation Coverage (DVEC)

```

/* Algorithm for Robot  $i$  */
Input: Parameter  $\kappa$ , Domain  $D$ , sequences
 $\{\beta^{(n)}\}_{n=1}^N$  and  $\{\gamma^{(n)}\}_{n=1}^N$ 
1 Initialization;
2 Random initial position  $x_i^{(0)}$ 
3 Sample initial position  $y_i^{(0)} = f^*(x_i^{(0)}) + \epsilon_i^{(0)}$ 
4 Exchange data  $\{\mathbf{x}_i^{(0)}, \mathbf{y}_i^{(0)}\}$ 
5 Calculate parameters  $\rho_i^{(0)}$  and  $\tau_i^{(0)}$ 
6 Calculate mean  $\mu_i^{(0)}(\cdot)$  and standard deviation  $\sigma_i^{(0)}(\cdot)$ 
7 for each iteration  $n \in \{1, \dots, N\}$  do
8   Calculate the estimate of  $f^*$  for all  $q \in D$  as
      $f_i^{(n)}(q) = \mu_i^{(n-1)}(q) - \sqrt{\beta^{(n)}} \sigma_i^{(n-1)}(q)$ 
     Choose estimation position as
      $e_i^{(n)} = \arg \max_{x \in V_i^{(n-1)}} \sigma^{(n-1)}(x)$ 
9   while  $x_i \neq (1 - \gamma)c_i^{(n)} + \gamma e_i^{(n)}$  do
10    Compute  $c_i^{(n)}(\cdot, f_i^{(n)}(\cdot))$  as in (3)
11    Move according to
      $\dot{x}_i = \kappa \left( (1 - \gamma^{(n)})c_i^{(n)} + \gamma^{(n)}e_i^{(n)} - x_i \right)$ 
12  end
13  Sample  $y_i^{(n)} = f^*(x_i^{(n)}) + \epsilon_i^{(n)}$ 
14  Exchange data  $\{\mathbf{x}_i^{(n)}, \mathbf{y}_i^{(n)}\}$ 
15  Update  $\rho_i^{(n)}$  and  $\tau_i^{(n)}$ 
16  Update  $\mu_i^{(n)}(\cdot)$  and  $\sigma_i^{(n)}(\cdot)$ 
17 end

```

B. Constrained Communication Protocol

The communication protocol in Section III-A considers a scenario when agents always incorporate their neighbors' samples to their model, irrespectively of the value added by such information. In this section, we present an alternative protocol that limits the amount of information shared between the agents in an effort to reduce redundant or uninformative transmissions that overload the communication system and increase the number of computations in the regression process.

Consider the collection of samples of Robot i , $\mathbf{y}_i^{(1:n)}$ at locations $\mathbf{x}_i^{(1:n)}$. We would like to know if a new sample $y_j^{(n)}$ obtained by Robot j , $j \in \mathcal{N}_i$ should be added to $\mathbf{x}_i^{(n)}$ by analyzing the amount of potential information that $\{x_j^{(n)}, y_j^{(n)}\}$ would add to the estimate f_i . To this end, the communication constrained estimation algorithm in [22] provides an upper bound on the difference between the prediction of the spatial field $f_i(\bar{x})$ at a location $\bar{x} \in D$ with and without incorporating a new sample at a location $\chi_j \in D$ to the Gaussian process regression. This is formally notated as

$$|f_i(\bar{x}; \mathbf{x}_i^{(1:n)}) - f_i(\bar{x}; \tilde{\mathbf{x}}_i^{(1:n)})|,$$

where $\tilde{\mathbf{x}}_i^{(1:n)} = \mathbf{x}_i^{(1:n)} \cup \chi_j$. We analogously use the notation

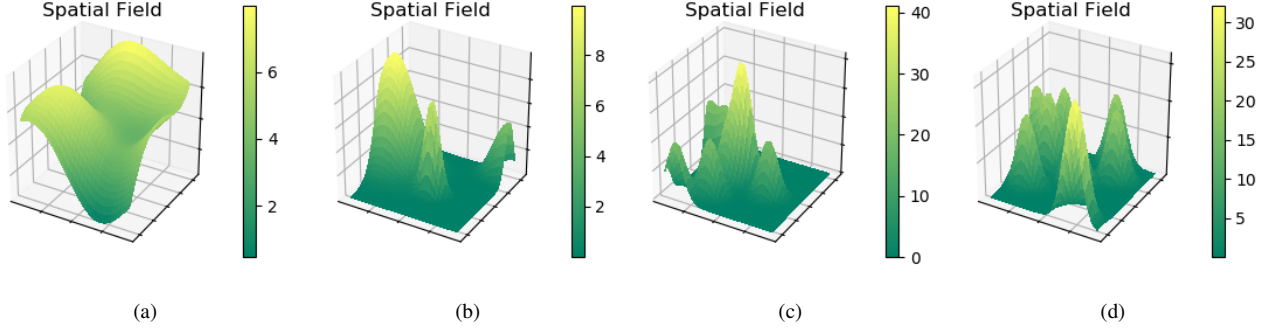


Fig. 1. Density functions used in the simulations. (a) shows the density function from [14]. The densities in (b), (c), and (d) are from [15]: (b) shows a sum of 3 Gaussian bivariate distributions with different means and variances, and the densities in (c) and (d) are sums of 9 Gaussian bivariate distributions with equal variances generated with random means.

$\tilde{\mathbf{y}}_i^{(1:n)} = \mathbf{y}_i^{(1:n)} \cup \psi_j$, with ψ_j is the noisy sample taken at χ_j according to (6). The notation $f_i(\bar{x}; \mathbf{x}_i^{(1:n)})$ is used here to indicate that the estimate of the density function at \bar{x} is obtained by using the collection of samples associated with $\mathbf{x}_i^{(1:n)}$ in the Gaussian process regression of f_i .

The bound is

$$|f_i(\bar{x}; \mathbf{x}_i^{(1:n)}) - f_i(\bar{x}; \tilde{\mathbf{x}}_i^{(1:n)})| \leq \|\mathbf{k}(\bar{x}; \tilde{\mathbf{x}}_i^{(1:n)})^\top\| \|\tilde{\mathbf{y}}_i^{(1:n)}\| \Delta K,$$

where

$$\Delta K \leq \max\{\|\mathcal{X}\|, \|\mathcal{Z}\|\} + \|\mathcal{Y}\|,$$

and

$$\begin{aligned} \mathcal{Z} &= (k(\chi_j, \chi_j; \tau_j) - \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})^{-1} \mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1} \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)}))^{-1} \\ \mathcal{Y} &= -\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1} \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)}) \mathcal{Z} \\ \mathcal{X} &= -\mathcal{Y} \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})^\top \mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1} \end{aligned} \quad (12)$$

Bounds on these terms can be evaluated using the Cauchy-Schwarz inequality and the feature-space view of the kernel function as in [22],

$$\begin{aligned} \|\mathcal{X}\| &\leq \|\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1} \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})\|^2 \|\mathcal{Z}\|, \\ \|\mathcal{Y}\| &\leq \|\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1} \mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})\| \|\mathcal{Z}\|, \\ \|\mathcal{Z}\| &\leq \frac{\frac{1}{k(\chi_j, \chi_j; \tau_j)} + \left(\frac{\|\mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})\|}{k(\chi_j, \chi_j; \tau_j)}\right)^2 \|\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1}\|}{1 - \left(\frac{\|\mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})\|}{k(\chi_j, \chi_j; \tau_j)}\right) \|\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)^{-1}\|^2}. \end{aligned} \quad (13)$$

The notation above is in principle the same as from Section II; however, the dimensions of the matrix $\mathbf{K}_{\mathbf{x}_i^{(1:n)}}(\tau_i)$ and vector $\mathbf{k}(\chi_j; \mathbf{x}_i^{(1:n)})$ cannot be stated *a priori*. Their dimensions depend on the communications between agent i and its neighbors up until iteration n . In (12), \mathcal{Z} represents the covariance in the prediction at the test point given our training points in the first n iterations. The other factors contributing to \mathcal{X} and \mathcal{Y} are related to the expression for the mean predictions at a test point [27]. A more detailed treatment can be found in [22].

Using the method outlined in [22], one can quantify the impact a new data point will have on the estimate based on the value of ΔK , which is computed using (13). [22] used this value to let a robot rank the impact of every candidate data point to be received and accepted only the most impactful points from other robots within a positive radius R from itself. In contrast, our implementation sets a threshold value for ΔK and determines if a sample from a neighbor at iteration n , $\{x_j^{(n)}, y_j^{(n)}\}$, $j \in \mathcal{N}_i^{(n)}$ should be added or not to the Gaussian process of $f_i^{(n)}$ based on the value of ΔK associated with the sample.

IV. SIMULATION RESULTS

The algorithms DVEC with naive and constrained communications (henceforth referred to as DVEC-nc and DVEC-cc, respectively) were compared to the baseline, VEC, in a series of simulations to determine the impact of considering local estimates of the density function. All algorithms were tested in a collection of four different spatial fields, shown in Fig. 1. Each trial ran for 15 iterations with a team of $M = 7$ robots, each one moving according to single integrator dynamics, i.e., $\dot{x}_i = u$, $u \in \mathbb{R}$ given by the VEC and DVEC algorithms.

In order to compare the performance of VEC, DVEC-nc and DVEC-cc, we use the error in the estimate of the density function as defined in [15],

$$\Delta^{(n)} = \int_D |f^{(n)}(q) - f^*(q)| dq. \quad (14)$$

However, while in the case of VEC the estimate of the function is uniquely defined at iteration n by $f^{(n)}$, for DVEC we have a collection of estimates $\{f_i^{(n)}\}_{i=1}^M$. In order to directly compare against the centralized estimation of VEC, a composition method from [10] was used to combine the estimates for analysis at each iteration. For this comparison, we use the composition defined by taking the spatial field estimate that gives the maximum value of the density $\max_i f_i^{(n)}(q)$, $\forall q \in D$. Since the estimate of the spatial field given by (9) underestimates the mean estimate of the field proportionally with the uncertainty of the measurement at that point, taking the maximum was shown

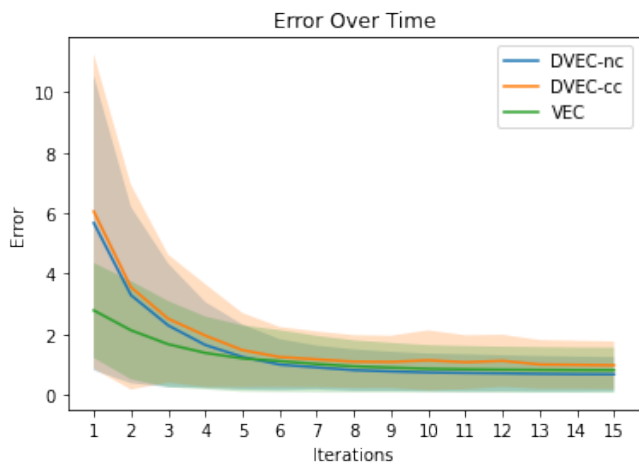


Fig. 2. Error in estimate as calculated by (14) for VEC (green), DVEC-nc (blue) and DVEC-cc (orange) over 15 iterations from 5 different initial conditions of the multi-robot team. The shaded regions represent standard deviation across the trials over the four tested densities in Fig. 1.

in testing to lead to lower errors on the composed density estimate as compared to using other composition operators. The composition can be parallelized as each point in the discretization of the domain is independent of the others.

The comparison on the error in the density estimate is shown in Fig. 2. The error in the DVEC algorithms is initially higher than VEC but quickly reaches a similar level by iteration 6 for both communication methods. Afterwards, the estimate ends up with near identical mean error (0.815 for VEC, 0.678 for DVEC-nc, and 0.888 for DVEC-cc at iteration 15). The larger standard deviation for the DVEC experiments stem from the initial communications between agents; since communication is only allowed between Voronoi neighbors, the initial configuration impacts the initial estimate much more than if the estimates were computed centrally. The standard deviations at the final iteration were similar: 0.743 for VEC, 0.568 for DVEC-nc, and 0.681 for DVEC-cc.

The algorithms are also compared in terms of regret at each iteration, defined as

$$r^{(n)} = H(f^*(\cdot), \mathbf{x}^{(n)}) - \min_{\mathbf{x}} H(f^*(\cdot), \mathbf{x}). \quad (15)$$

The comparison between the regret for VEC, DVEC-nc and DVEC-cc is shown in Fig. 3. Here, the regret starts at a similar value for all algorithms since it is a function of the position and true density function, and not the density estimates. At later iterations, the regret between DVEC-nc and DVEC-cc is similar, but larger than VEC. This is likely due to the CVT being only a *local* minimum of the cost function (2). Although the estimate $f^{(n)}(\cdot)$ is shown in Fig. 2 to converge to the true density $f^*(\cdot)$ for all methods, a better spatial configuration appears to be established earlier for VEC, around iterations 2 and 3. Since both methods prioritize exploration in early iterations and focus more on coverage later, the configurations do not change much after the first few iterations.

It is important to note that Fig. 2 represents composed




Fig. 3. Regret as calculated by (15) for VEC (green), DVEC-nc (blue) and DVEC-cc (orange) over 15 iterations from 5 initial conditions. The shaded regions represent standard deviation across the trials over the four tested densities.

estimates for DVEC-nc and DVEC-cc, which in this case is $\max_i f_i^{(n)}(q), \forall q \in D$. Even though the composed estimate of the robot team may have low error (Fig. 2), the composed estimate does not impact the decisions for individual agents. The larger standard deviation and slight increase in regret for the DVEC methods is hypothesized to be due to the incomplete information at the individual level, which may lead to suboptimal coverage positions.

The average accumulated number of information transfers per robot is compared between DVEC-nc and DVEC-cc in Fig. 4. The equivalent number of information transfers to recover the performance from VEC is also shown in green. The plot shows that as iterations increase, the communication constrained protocol DVEC-cc results in fewer transmissions of data across the entire robot team. Both versions of DVEC transfer significantly less information than VEC. The transfer of data for DVEC-cc stops almost entirely around iteration 7, consistent with Fig. 2 when the error in density estimate has converged. This is also determined by the thresholding parameter for ΔK . In spite of this reduced communication for DVEC-cc, the regret and error for DVEC-nc and DVEC-cc have experimentally been shown to be nearly identical across a range of spatial fields and initial configurations.

V. CONCLUSIONS

In this paper, we introduced a decentralized algorithm for a team of robots to simultaneously learn and cover an unknown spatial field. Leveraging the spatial structure of the coverage problem where individual robots monitor events within their Voronoi cell, we propose a decentralized approach to learning and covering the spatial field whereby each robot runs a Gaussian process that aggregates the samples of itself and its Voronoi neighbors.

Simulations demonstrated the effectiveness of Decentralized Voronoi Estimation Coverage (DVEC) with both naive and constrained communication exchanges between Voronoi neighbors. Communications and computation were reduced




Fig. 4. Average accumulated number of information transfers for an individual robot in the 7 robot team over each iteration for VEC (green), DVEC-nc (blue) and DVEC-cc (orange). An information transfer is defined as an incident where a robot transmits a location and sample to another robot.

by taking into account the novelty of the information being received, rejecting data points that do not greatly improve the estimate of the spatial field. Despite the constrained communication drastically lowering the amount of information transferred between agents, the estimation error and coverage regret was shown to be comparable to DVEC under a naive communication scheme that exchanged information at every step. Both implementations of DVEC presented minor drawbacks in terms of performance when compared to a centralized computer.

VI. ACKNOWLEDGEMENTS

The authors thank Professor Gennaro Notomista for his invaluable explanations and insights about communication-constrained coverage.

REFERENCES

- [1] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, Apr. 2004.
- [2] M. Schwager, J. McLurkin, and D. Rus, "Distributed coverage control with sensory feedback for networked robots," in *Proceedings of Robotics: Science and Systems*, Aug. 2006.
- [3] M. Schwager, D. Rus, and J.-J. Slotine, "Decentralized, adaptive coverage control for networked robots," *The International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.
- [4] M. Schwager, M. P. Vitus, D. Rus, and C. J. Tomlin, "Robust adaptive coverage for robotic sensor networks," in *Robotics Research*. Springer, 2017, pp. 437–454.
- [5] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robotics Automation Magazine*, vol. 19, no. 1, pp. 24–39, 2012.
- [6] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, "Collective motion, sensor networks, and ocean sampling," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, 2007.
- [7] L. C. A. Pimenta, M. Schwager, Q. Lindsey, V. Kumar, D. Rus, R. C. Mesquita, and G. A. S. Pereira, *Simultaneous Coverage and Tracking (SCAT) of Moving Targets with Robot Networks*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 85–99.
- [8] S. G. Lee, Y. Diaz-Mercado, and M. Egerstedt, "Multirobot control using time-varying density functions," *IEEE Transactions on Robotics*, vol. 31, no. 2, pp. 489–493, 2015.
- [9] M. Santos, S. Mayya, G. Notomista, and M. Egerstedt, "Decentralized minimum-energy coverage control for time-varying density functions," in *International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, 2019, pp. 155–161.
- [10] M. Santos, Y. Diaz-Mercado, and M. Egerstedt, "Coverage control for multirobot teams with heterogeneous sensing capabilities," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 919–925, 2018.
- [11] M. Santos and M. Egerstedt, "Coverage control for multi-robot teams with heterogeneous sensing capabilities using limited communications," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 5313–5319.
- [12] F. Lekien and N. E. Leonard, "Nonuniform coverage and cartograms," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 351–372, 2009.
- [13] S. Bhattacharya, R. Ghrist, and V. Kumar, "Multi-robot coverage and exploration on Riemannian manifolds with boundaries," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 113–137, 2014.
- [14] A. Benevento, M. Santos, G. Notarstefano, K. Paynabar, M. Bloch, and M. Egerstedt, "Multi-robot coordination for estimation and coverage of unknown spatial fields," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7740–7746.
- [15] M. Santos, U. Madhushani, A. Benevento, and N. E. Leonard, "Multi-robot learning and coverage of unknown spatial fields," in *International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, 2021, pp. 137–145.
- [16] S. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.
- [17] W. Luo and K. Sycara, "Adaptive sampling and online learning in multi-robot sensor coverage with mixture of Gaussian processes," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6359–6364.
- [18] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," in *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, ser. IJCAI'09. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2009, p. 1843–1850.
- [19] F. Zhang and N. E. Leonard, "Cooperative filters and control for cooperative exploration," *IEEE Transactions on Automatic Control*, vol. 55, no. 3, pp. 650–663, 2010.
- [20] C. Peterson and D. A. Paley, "Distributed estimation for motion coordination in an unknown spatially varying flowfield," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 3, p. 894–898, 2013.
- [21] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [22] G. Notomista and M. Egerstedt, "Communication constrained distributed spatial field estimation using mobile sensor networks," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 9642–9649, 2020.
- [23] W. Luo, C. Nam, G. Kantor, and K. Sycara, "Distributed environmental modeling and adaptive sampling for multi-robot sensor coverage," in *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2019, p. 1488–1496.
- [24] A. Carron, M. Todescato, R. Carli, L. Schenato, and G. Pillonetto, "Multi-agents adaptive estimation and coverage control using Gaussian regression," in *European Control Conference (ECC)*, 2015, pp. 2490–2495.
- [25] L. J. Guibas, D. E. Knuth, and M. Sharir, "Randomized incremental construction of Delaunay and Voronoi diagrams," *Algorithmica*, vol. 7, no. 1, pp. 381–413, 1992.
- [26] Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi tessellations: Applications and algorithms," *SIAM Rev.*, vol. 41, no. 4, pp. 637–676, Dec. 1999.
- [27] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*. Springer, 2003, pp. 63–71.