

A Distributed Gaussian Process Model for Multi-Robot Mapping

Seth Nabarro^{1*}, Mark van der Wilk² and Andrew J. Davison¹

Abstract—We propose *DistGP*: a multi-robot learning method for collaborative learning of a global function using only local experience and computation. We utilise a sparse Gaussian process (GP) model with a factorisation that mirrors the multi-robot structure of the task, and admits distributed training via Gaussian belief propagation (GBP). Our loopy model outperforms Tree-Structured GPs [1] and can be trained online and in settings with dynamic connectivity. We show that such distributed, asynchronous training can reach the same performance as a centralised, batch-trained model, albeit with slower convergence. Last, we compare to DiNNO [2], a distributed neural network (NN) optimiser, and find DistGP achieves superior accuracy, is more robust to sparse communication and is better able to learn continually.

I. INTRODUCTION

Many mapping tasks involve large areas and dynamic maps, suggesting a divide-and-conquer approach using multiple robots. Centralised communication and computation is often not feasible in domains like environmental monitoring [3], robotic agriculture [4] or space exploration [5], which involve many robots accumulating large datasets in poorly connected regions. Further, centralised approaches can be slow where robots may need up-to-date local models to guide exploration. These considerations motivate decentralised approaches, with local computation and mesh connectivity.

In this work, we present DistGP: a multi-robot algorithm to fit an arbitrary global function using only local observation and computation; with peer-to-peer communication. We use sparse GPs whose inducing points approximate a full GP at lower cost, and act as local summaries of the function. They provide a useful abstraction for data fusion and inter-robot exchange. Further, they are assumed to be Gaussian distributed, enabling decentralised inference with GBP.

We first consider tree-structured GPs (TSGPs) [1], whose inducing points are partitioned into sparsely connected blocks. Each block covers a distinct region of the input space and is maintained by a single robot. Robots exchange GBP messages with their neighbours to ensure their inducing points are consistent. TSGP restricts connectivity to a tree to guarantee rapid convergence, however we show empirically that this constraint limits accuracy: inducing points close in Euclidean space may not be directly connected, leading to discontinuities at their boundary. Further, in dynamic multi-robot scenarios, two communicating robots are prevented from sharing information if doing so would introduce a cycle.

To address this limitation, we propose a simple extension: relax the tree constraint and allow cycles. While GBP in

our loopy model requires more iterations to converge, we demonstrate a substantial improvement in accuracy. Moreover, our approach is better suited to dynamic and asynchronous settings, where inter-robot communication patterns are continually changing – unsupported by TSGP, whose message schedule requires fixed connectivity. Our loopy model naturally exploits all robot encounters and converges with asynchronous message passing. Our proposed loopy factor graph model is illustrated in Fig. 1.

Our motivations are similar to Yu et al. [2], who use a distributed NN optimiser (DiNNO) for multi-robot learning. In contrast to DiNNO, DistGP i) enforces inter-robot consistency locally, only where both robots have inducing points, rather than in global weight-space, ii) is a probabilistic model supporting on-the-fly data fusion for online learning, and iii) is scalable as each robot stores only a local map estimate.

In short our contributions are as follows:

- 1) A new distributed GP model, trainable with GBP, which outperforms TSGPs [1].
- 2) Empirical analysis showing the model can be built on-the-fly and admits distributed, asynchronous inference via GBP. It is thus well-suited to multi-robot mapping.
- 3) Experimental evaluation of our method on environmental monitoring and occupancy mapping simulations. We show superior performance to DiNNO [2].

II. RELATED WORK

a) Distributed NN Maps: Many works build distributed implicit neural maps through inter-robot collaboration [2], [6]–[9], either using consensus alternating direction method of multipliers (CADMM) [10] or direct regularisation to constrain NN weights to be similar. DiNNO [2] is most similar to ours as it is a general method for multi-robot function learning rather than a specific method for distributed 3D reconstruction. While distributed neural maps enforce agreement in global weight space, we represent maps with factorised GPs which enable local consistency constraints, specific to regions in which both robots have inducing points.

In DiNNO, each robot maintains a NN estimate of the global map, which are brought into agreement by CADMM. In this work however, we seek to build a distributed map: each robot maintains their local segment of the map and inter-robot interactions are used to ensure each segment is consistent with its neighbours. We argue DistGP is:

- 1) **more scalable.** For problems with large areas and large numbers of robots, each robot having a global map is inefficient and reaching consensus is more challenging.
- 2) **more modular.** We divide the map into segments which can be individually copied and communicated

¹Dyson Robotics Lab, Imperial College London

²Department of Computer Science, University of Oxford

* sdn09@ic.ac.uk

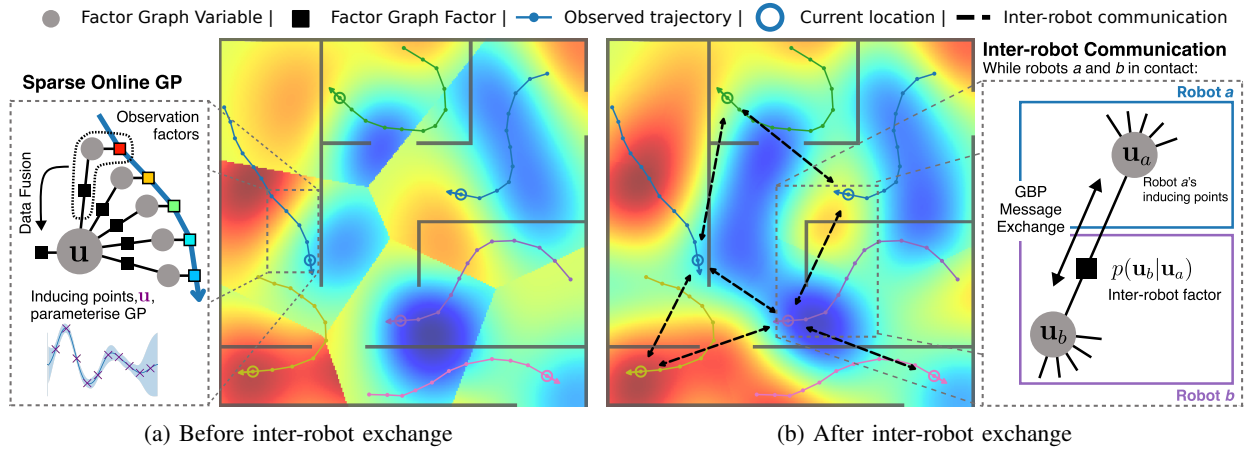


Fig. 1: In *DistGP*, each robot maintains a block of inducing points which parameterise their local GP (far-left) and into which observations can be fused. Communicating robots harmonise models by connecting inducing points via GP consistency factors and exchanging GBP messages (far-right). Here, the map at each point is predicted by the closest robot. Before communication (centre-left) the map is discontinuous and inaccurate, but this is resolved using only local inter-robot exchanges (centre-right).

as required. In contrast, it is unclear how a map stored in the weights of a NN can be subdivided.

- 3) **also able to generate a shared global map.** In the extreme case that each robot shares its own inducing point posterior, and all others it has received thus far, we can achieve global map sharing as in DiNNO.

b) Multi-Robot GPs: As in this work, [11] use a distributed GP for multi-robot learning. Each robot first trains a local GP on its observations and they are combined as a mixture of GPs, with location-varying weights. The Expectation Maximisation algorithm optimises the local models in tandem with the mixture weights of the different observations. Unlike our method, full connectivity is needed for normalised mixture weights, and *all* robots must generate a prediction at any test point. In addition, the use of dense GPs precludes scaling to large datasets. [12] present a multi-agent inducing point GP method for human trajectory prediction, however they do not consider the problem of distributed learning of a function over the space being navigated.

c) Multi-Robot GBP: Recent work has shown GBP to be an effective distributed solver. Our work draws inspiration from [13], in using GBP exchange between robots to enable collaborative estimation. However, [13] estimate the robot locations, where we estimate a distributed map. Like us, [14] utilise GBP for multi-robot mapping. Each robot maintains a grid of the input space with one variable per cell to represent its function value. These are updated via i) observation by the robot or ii) inter-robot communication where two robots enforce consistency between their pairs of grid variables. Unlike *DistGP*, this method does not benefit from function approximation – it is unable to generalise to unobserved regions and can only be queried at fixed resolution.

III. BACKGROUND

A. Gaussian Process Regression

A GP is a generalisation of a Gaussian distribution to the infinite dimensional space of functions. It is defined by a

mean function $\mu : \mathbb{R}^D \rightarrow \mathbb{R}$, and a covariance (or kernel) function $k_\theta : \mathbb{R}^{D \times D} \rightarrow \mathbb{R}$, where D is the input dimension. $k_\theta(\mathbf{x}_a, \mathbf{x}_b)$ is the covariance between the function values at pairs of input locations $(\mathbf{x}_a, \mathbf{x}_b)$ and has hyperparameters θ .

For regression, a GP can be used as a prior $p(f(\cdot)|\theta)$ on the function to be fitted, $f : \mathbb{R}^D \rightarrow \mathbb{R}$. For an iid Gaussian likelihood on observations $\mathbf{y} = [y_1, \dots, y_N]^\top$ at inputs $X = [\mathbf{x}_1^\top, \dots, \mathbf{x}_N^\top]^\top$, the posterior of $f(\cdot)$ at test point \mathbf{x}^* is:

$$p(f(\mathbf{x}^*)|\mathbf{y}) = \mathcal{N}(\mu_*, \Sigma_*) \quad (1)$$

$$\mu_* := \mathbf{k}_{\mathbf{x}^*x} (K_x + \sigma^2 I_x)^{-1} \mathbf{y},$$

$$\Sigma_* := K_{x^*} - \mathbf{k}_{\mathbf{x}^*x} (K_x + \sigma^2 I_x)^{-1} \mathbf{k}_{x\mathbf{x}^*},$$

where $K_x \in \mathbb{R}^{N \times N}$ has ij th element $k_\theta(\mathbf{x}_i, \mathbf{x}_j)$, $K_{x^*} \in \mathbb{R}^{1 \times 1}$ has element $k_\theta(\mathbf{x}^*, \mathbf{x}^*)$, $\mathbf{k}_{\mathbf{x}^*x}$ is a length N vector with l th element $k_\theta(\mathbf{x}_l, \mathbf{x}^*)$, and σ is the observation noise. Hyperparameters θ and σ can be trained by gradient ascent of the (log-)marginal likelihood, $p(\mathbf{y}|\theta) = \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\theta)d\mathbf{f}$.

The time complexity of the matrix inverse in (1) is $O(N^3)$, which is prohibitive for large datasets, motivating approximations with cheaper training and prediction.

B. Inducing Point Approximations

Inducing point GPs seek to compress the dataset of N observations into $M < N$ pseudo-datapoints (inducing points) which capture its important characteristics at reduced computational cost. One such method is the Fully Independent Training Conditional (FITC) approximation [15], [16]. The inducing variables $\mathbf{u} \in \mathbb{R}^M$ approximate $f(\cdot)$ at the chosen input locations $Z = [\mathbf{z}_1^\top, \dots, \mathbf{z}_M^\top]^\top$. The generative model is

$$q(\mathbf{f}, \mathbf{u}) = q(\mathbf{f}|\mathbf{u})p(\mathbf{u}) = \prod_{i=1}^N p(f_i|\mathbf{u})p(\mathbf{u}), \quad (2)$$

where the conditional independence of $\mathbf{f}|\mathbf{u}$ is a model assumption which enables cheaper inference and learning, $p(\mathbf{u})$ is a dense GP prior over the inducing points. The resulting posterior predictive $p(f(\mathbf{x}^*)|\mathbf{y}, X)$ requires inversion

of dense $M \times M$ and diagonal $N \times N$ matrices, thus avoiding the $O(N^3)$ cost of inverting a dense $N \times N$ matrix.

C. Tree-Structured Gaussian Process Approximations

While inducing point approximations such as FITC are efficient when $M \ll N$, Bui and Turner [1] argue that many datasets require $M \approx N$ for accurate prediction. For efficient inference in this setting, they propose approximating a dense prior over \mathbf{u} by dividing them into blocks $\mathbf{u} = [\mathbf{u}_{B_1}^\top, \dots, \mathbf{u}_{B_K}^\top]^\top$ and connecting them in a tree

$$q(\mathbf{u}) = q(\mathbf{u}_{B_1}) \prod_{k=2}^K p(\mathbf{u}_{B_k} | \mathbf{u}_{\text{par}(B_k)}), \quad (3)$$

where B_1 is the root node and $\text{par}(B_k)$ is the parent of B_k .

Each block generates a disjoint observation subset \mathbf{f}_{C_k}

$$q(\mathbf{f} | \mathbf{u}) = \prod_{k=1}^K q(\mathbf{f}_{C_k} | \mathbf{u}_{B_k}). \quad (4)$$

with connections between elements of \mathbf{f} and \mathbf{u} determined by proximity. For example, $f_a := f(\mathbf{x}_a)$ may be connected to the \mathbf{u}_{B_k} whose input location centroid is closest to \mathbf{x}_a . By minimising the KL divergence from the dense joint $p(\mathbf{f}, \mathbf{u})$ to approximation $q(\mathbf{f}, \mathbf{u})$, they derive the form of factors $q(\mathbf{u}_{B_k} | \mathbf{u}_{\text{par}(B_k)})$ and $q(\mathbf{f}_{C_k} | \mathbf{u}_{B_k})$, showing they are equal to standard GP conditionals $p(\mathbf{u}_{B_k} | \mathbf{u}_{\text{par}(B_k)})$ and $p(\mathbf{f}_{C_k} | \mathbf{u}_{B_k})$.

As TSGP has a sparse set of factors $\{p(\mathbf{u}_{B_k} | \mathbf{u}_{\text{par}(B_k)})\}_k$ in place of a dense prior $p(\mathbf{u})$, inference can be efficient in models where the number of inducing points is large. Bui and Turner [1] show how the posteriors $\{q(\mathbf{u}_{B_k} | \mathbf{y})\}_k$ can be inferred with GBP message passing. The tree-structured topology ensures GBP converges in one sweep from leaves to root and another from root to leaves.

After inference, prediction at a new \mathbf{x}^* can be made by finding a nearby block of inducing points $\mathbf{u}_{B_{k^*}}$ (e.g. with input locations whose centroid is closest) and estimating $f(\mathbf{x}^*)$ by marginalising over its posterior $q(\mathbf{u}_{B_{k^*}} | \mathbf{y})$:

$$q(f(\mathbf{x}^*) | \mathbf{y}) = \int p(f(\mathbf{x}^*) | \mathbf{u}_{B_{k^*}}) q(\mathbf{u}_{B_{k^*}} | \mathbf{y}) d\mathbf{u}_{B_{k^*}}. \quad (5)$$

This prediction relies only on the closest $\mathbf{u}_{B_{k^*}}$, avoiding test-time communication with other parts of the graph.

D. Belief Propagation

Belief propagation (BP) [17] is a message passing algorithm to compute the marginal beliefs of variables in a graphical model. Given the joint distribution's factorisation

$$p(\mathbf{x}_1, \dots, \mathbf{x}_N) = \frac{1}{Z} \prod_j \phi_j(\{\mathbf{x}_l\}_{l \in \text{nei}(j)}), \quad (6)$$

BP computes (or estimates) $\{p(\mathbf{x}_j)\}_{j=1}^N$. If factors ϕ_j are conditional on observations \mathbf{y} , then the marginals constitute a mean-field posterior $q(\mathbf{x}_1, \dots, \mathbf{x}_n | \mathbf{y}) = \prod_{i=1}^n p(\mathbf{x}_i | \mathbf{y})$.

Each edge in a factor graph stores both a factor-to-variable message $m_{\phi_j \rightarrow \mathbf{x}_i}(\mathbf{x}_i)$ and a variable-to-factor message $m_{\mathbf{x}_i \rightarrow \phi_j}(\mathbf{x}_i)$. BP iterates between updating these:

$$m_{\phi_j \rightarrow \mathbf{x}_i}(\mathbf{x}_i) \leftarrow \sum_{\{\mathbf{x}_r\}_{r \in \text{nei}(j) \setminus i}} \phi_j(\{\mathbf{x}_l\}_{l \in \text{nei}(j)}) \prod_{k \in \text{nei}(j) \setminus i} m_{\mathbf{x}_k \rightarrow \phi_j}(\mathbf{x}_k) \quad (7)$$

$$m_{\mathbf{x}_i \rightarrow \phi_j}(\mathbf{x}_i) \leftarrow \prod_{s \in \text{nei}(i) \setminus j} m_{\phi_s \rightarrow \mathbf{x}_i}(\mathbf{x}_i), \quad (8)$$

where $\text{nei}(r) \setminus s$ is the set of all nodes connected to r except s . After convergence, the marginals are given by the product of incoming message from all connected factors to a variable,

$$p(\mathbf{x}_i) = \prod_{s \in \text{nei}(i)} m_{\phi_s \rightarrow \mathbf{x}_i}(\mathbf{x}_i). \quad (9)$$

Message and belief updates are more efficient using the natural parameterisation of messages and factors, as the product of messages reduces to the sum of their parameters.

BP is only guaranteed to converge to exact marginals in tree-structured models, where convergence can be reached by a leaves-to-root then root-to-leaves sweep of message updates. Despite a lack of guarantees, BP often performs well on tasks involving cyclic graphs [18], [19].

1) *Gaussian Belief Propagation*: For Gaussian models, the natural parameters are the precision (inverse covariance) $\Lambda := \Sigma^{-1}$ and information vector $\eta := \Sigma^{-1}\mu$. Factor-to-variable updates (7) take the form (see [20])

$$\eta_{\phi_j \rightarrow \mathbf{x}_i} \leftarrow \eta_i^{(\phi_j)} - \Lambda_{i, \setminus i}^{(\phi_j)} \left(\Lambda_{\setminus i, i}^{(\phi_j)} + \Lambda_{\mathbf{x}_i \rightarrow \phi_j} \right)^{-1} \left(\eta_{\setminus i}^{(\phi_j)} + \eta_{\mathbf{x}_i \rightarrow \phi_j} \right)$$

$$\Lambda_{\phi_j \rightarrow \mathbf{x}_i} \leftarrow \Lambda_{i, i}^{(\phi_j)} - \Lambda_{i, \setminus i}^{(\phi_j)} \left(\Lambda_{\setminus i, i}^{(\phi_j)} + \Lambda_{\mathbf{x}_i \rightarrow \phi_j} \right)^{-1} \Lambda_{\setminus i, i}^{(\phi_j)}$$

where $(\eta^{(\phi_j)}, \Lambda^{(\phi_j)})$ are parameters of ϕ_j and subscript $\mathbf{x}_i \rightarrow \phi_j$ denotes concatenation of all messages incoming to ϕ_j except that from \mathbf{x}_i . Variable-to-factor messages (8) are

$$\eta_{\mathbf{x}_i \rightarrow \phi_j} \leftarrow \sum_{k \in \text{nei}(i) \setminus j} \eta_{\phi_k \rightarrow \mathbf{x}_i}, \quad \Lambda_{\mathbf{x}_i \rightarrow \phi_j} \leftarrow \sum_{k \in \text{nei}(i) \setminus j} \Lambda_{\phi_k \rightarrow \mathbf{x}_i}.$$

As in the general case, convergence in loopy Gaussian models is not guaranteed, though if GBP *does* converge in such cases, marginal means are guaranteed to be correct [21].

IV. METHOD

A. Problem Statement

We aim to learn a scalar map $f : \mathbb{R}^D \rightarrow \mathbb{R}$ distributed between K robots, $f_{\text{dist}} = \{f_k\}_{k=1}^K$, $f_{\text{dist}}(\mathbf{x}^*) = f_i(\mathbf{x}^*) \iff \mathbf{x}^* \in \mathcal{R}_i$, where partial map $f_i(\cdot)$ is maintained by robot i which covers region \mathcal{R}_i . For static maps, $D = D_{\text{space}}$, where dynamic maps are also a function of time: $D = D_{\text{space}} + 1$.

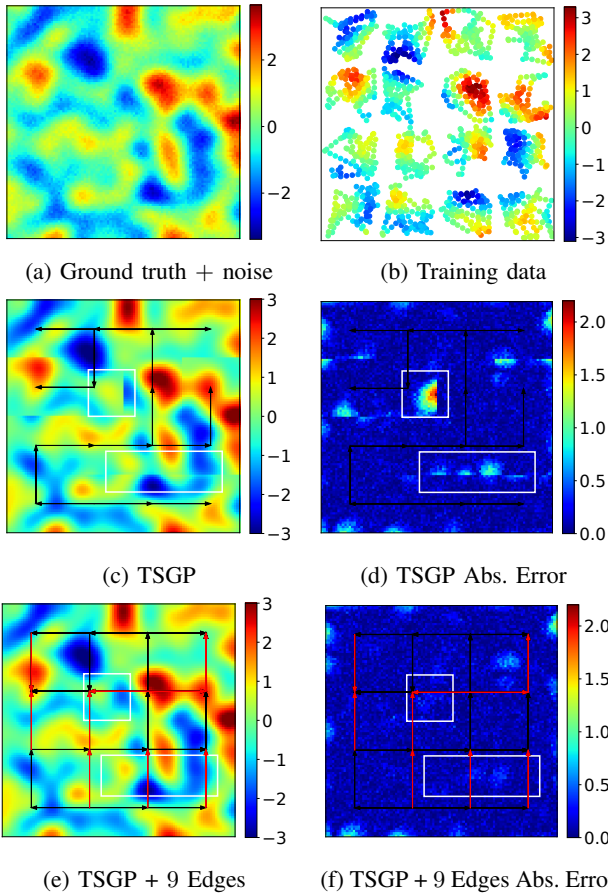


Fig. 2: **Limitations of TSGP.** Some robots remain unconnected to ensure tree connectivity (as in c). This leads to discontinuities in the prediction at the boundary between unconnected robots (d), which can be reduced significantly by adding loop-forming edges (f) (see white rectangles).

B. Improving Over Tree-Structured GPs

The following features TSGP make it a good candidate for distributed multi-robot mapping:

- 1) Each inducing block \mathbf{u}_{B_k} models a different subset of observations (4), so robots can update their local posteriors $q(\mathbf{u}_{B_k}|y_{C_k})$ with new experience without needing to communicate with others.
- 2) Each \mathbf{u}_{B_k} only connects to its neighbours (3). With GBP inference, this ensures inter-robot messages are local exchanges – no global communication is needed.

To evaluate TSGP, we train it on simulated multi-robot data. We sample a static, ground truth function from a GP with $D = 2$, which the robots observe (with additive Gaussian noise) at each timestep. The input space covering $[-1, 1] \times [-1, 1]$ is divided into a 4×4 grid, with a robot occupying each cell. We generate robot trajectories by randomly sampling a target within each cell, which the robot heads towards at a constant speed. Once within $\epsilon = 0.02$, the target is resampled. The resulting observations are presented in Fig. 2b. Each robot’s inducing points are placed on a regular grid spanning its cell. Blocks of inducing points between robots in adjacent cells are connected, such that the graph




Fig. 3: **Benefit of additional edges.** The prediction error (vertical axis) reduces as more edges are added to the TSGP model (horizontal axis, $x = 0$ is TSGP), but plateaus after ~ 9 extra edges, where accuracy is equal to a dense GP prior on all inducing points (green). Results averaged over 10 seeds, errorbars cover ± 1 standard error (SE) around mean.

of all inducing point blocks forms a tree (black arrows in Fig. 2c). The model is batch-trained on the full dataset.

The resulting predictions and error (Figs. 2c and 2d) show discontinuities in the learnt function and large errors at the boundaries between unconnected cells. This pathology results from the model’s tree-structure: a tree connecting nodes on a grid with $D \geq 2$ will necessarily include unconnected neighbours. The learnt function is discontinuous at their boundary due to an absence of inter-robot factors.

To quantify this effect, we train a sequence of models, starting at TSGP and adding edges one-by-one, prioritising those between robots close in Euclidean space but distant in the graph. We also train two baselines: i) with no inter-robot connections (each robot fits an independent FITC GP), and ii) with a full GP prior over *all* inducing points. We train each model 10 times on different datasets, each randomly sampled as in Fig. 2b. We evaluate on a dense grid, but remove test points less than $d = 0.06$ from any training observations.

The results (Fig. 3) show that models with more edges have higher accuracy, until ~ 9 additional edges after which performance matches the full GP. Examining predictions and errors for a model with 9 extra edges on the example dataset (Figs. 2e and f) we see discontinuities are much reduced (e.g. in white rectangles, Figs. 2d and 2f). We believe performance plateaus at 9 extra edges (red arrows, Fig. 2e) because here each node connects to its immediate NESW neighbours. Further longer-range edges have little impact. This is consistent with repeated analyses on 3×3 and 5×5 grids, which plateau at 4 and 16 extra edges respectively.

C. Asynchronous Model Construction and Training

We have shown that TSGP’s mapping performance can be improved with additional connections. However, we have thus far assumed that all data are available upfront, and any robot could be connected to any other. In multi-robot mapping these assumptions are not realistic, so we now design a model which can be trained incrementally, with dynamic connectivity as per the robots’ relative locations. We will now describe (1) how observations are fused into local models, (2) where new inducing points are added, (3) how inter-robot connections are made, (4) how kernel hyperparameters are learnt and (5) how predictions are generated.

1) *Adding New Observations:* Desirable features of a robot’s local model include i) scalability w.r.t the number of observations, ii) ability to efficiently incorporate new observations, iii) ability to discard old observations after model update. To this end, we employ the FITC GP approximation [15], [16] (see factor graph in Fig. 1a). For a robot with inducing points \mathbf{u}_{B_i} and observations $\mathbf{y}^{(C_i)} = [y_1^{(C_i)}, \dots, y_{N_i}^{(C_i)}]^\top$ at $X^{(C_i)} = [(\mathbf{x}_1^{(C_i)})^\top, \dots, (\mathbf{x}_{N_i}^{(C_i)})^\top]^\top$, this approximation can be written

$$q(\mathbf{y}_i, \mathbf{f}_{C_i} | \mathbf{u}_{B_i}) = \prod_{n=1}^{N_i} p(y_n^{(C_i)} | f_n^{(C_i)}) p(f_n^{(C_i)} | \mathbf{u}_{B_i}). \quad (10)$$

For $N_i < M_i$, this approximation has inference complexity $O(M_i^2 N_i)$, and is therefore scalable to large N_i . Further, a new observation $y_{N_i+1}^{(C_i)}$ can be easily added online by creating a $f_{N_i+1}^{(C_i)}$ variable, new $p(y_{N_i+1}^{(C_i)} | f_{N_i+1}^{(C_i)})$, $p(f_{N_i+1}^{(C_i)} | \mathbf{u}_{B_i})$ factors, and calculating the resulting message to variable \mathbf{u}_{B_i} . Last, individual observations can be discarded after their message from factor $p(f_n^{(C_i)} | \mathbf{u}_{B_i})$ to \mathbf{u}_{B_i} has been computed. We maintain a unary factor on \mathbf{u}_{B_i} to accumulate messages from discarded observations. After this, factors $p(y_n^{(C_i)} | f_n^{(C_i)})$, $p(f_n^{(C_i)} | \mathbf{u}_{B_i})$ and variable $f_{i,m}$ are removed.

Discarding observation factors online introduces two sources of approximation relative to batch fitting. First, in most cases we add inducing points incrementally, as the robot gathers experience. Thus inducing points which have been added after an observation factor has been discarded will not benefit from the corresponding message. This does not affect inference when the inducing point locations are fixed throughout. Second, for non-linear observation factors $p(y_n^{(C_i)} | f_n^{(C_i)})$, the message to \mathbf{u}_{B_i} depends on the estimate of $f_n^{(C_i)}$. Observations added after $y_n^{(C_i)}$ is removed could have affected $\mathbb{E}[f_n^{(C_i)}]$ and therefore the message to \mathbf{u}_{B_i} .

2) *Adding New Inducing Points:* We efficiently add inducing points to the model using the *greedy variance selection* method of [22]. This method i) chooses an inducing point at the observation location with the highest marginal output variance under $p(f_n^{(C_i)} | \mathbf{u}_{B_i})$, ii) adds it to Z_{B_i} , iii) updates the model and iv) repeats until M_i is reached.

In our setting, the set of observations grows with time; we do not have the static dataset assumed in [22]. Nonetheless we employ greedy variance selection in an online fashion, adding a new inducing point to the model every k timesteps. We restrict our choice of new locations to observations which have not yet been fused and discarded (see Sec. IV-C.1). This ensures that the new inducing point will benefit from a strong message from the co-located observation.

In addition to selecting inducing points from observations, we find adding inducing points near the boundary between connecting robots provides the model more flexibility to make their representations consistent. To this end, we use two strategies: i) each robot adds a line of inducing points along the boundary; or ii) each robot adds inducing points at the same location as a subset of the other’s near the boundary.

3) *Adding New Connections:* We next consider how two communicating robots should exchange information to syn-

chronise their representations. We connect their inducing points via a shared factor, but what form should this factor take? We first observe that the optimal prior over inducing points, i.e. that which most closely approximates a full GP over observations, is itself a full GP over the inducing points:

$$p(\mathbf{u}_{B_1}, \dots, \mathbf{u}_{B_b}) = \mathcal{N}(\mathbf{0}, K_{Z_{\text{all}}}) \quad (11)$$

where $K_{Z_{\text{all}}} \in \mathbb{R}^{M \times M}$ is the kernel matrix over all inducing points. This joint prior can be factorised sequentially:

$$p(\mathbf{u}_{B_1}, \dots, \mathbf{u}_{B_b}) = p(\mathbf{u}_{B_1}) \prod_{i=2}^b p(\mathbf{u}_{B_i} | \{\mathbf{u}_{B_j}\}_{j < i}) \quad (12)$$

meaning it can be distributed s.t. each robot “owns” one factor. However, this fully connected structure is computationally infeasible: inference scales as $O(M^3)$, where $M := \sum_i \dim(\mathbf{u}_{B_i})$. Moreover, in realistic multi-robot settings, each robot can only communicate with nearby peers, where this factorisation implies all-to-all communication.

Despite these constraints, (12) provides a theoretical target: it implies an inter-robot communication protocol which would converge to the ideal $p(\mathbf{u}_{B_1}, \dots, \mathbf{u}_{B_b})$ in the limit of full connectivity. The protocol is as follows:

- 1) **Initialisation:** To begin, each robot has a unary GP prior on its inducing variables $p(\mathbf{u}_{B_i}) = \mathcal{N}(\mathbf{0}, K_{Z_{B_i}})$, equal to the message that \mathbf{u}_{B_i} would receive from a global dense prior over inducing points (11), due to the marginal consistency of GPs. This ensures message updates are well-conditioned, enabling inference in each robot’s local model before any inter-robot contact.
- 2) **Inter-robot Exchange:** Suppose robot i now encounters j . The two compare their indices to determine which will act as the parent in the connecting factor. If $i > j$, for example, robot j will be the parent whose inducing points will condition i ’s prior. Three updates are now executed: i) the unary prior is amended to a conditional: $p(\mathbf{u}_{B_i}) \rightarrow p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j})$ and connected to \mathbf{u}_{B_j} , ii) variable-to-factor messages $m_{\mathbf{u}_{B_i} \rightarrow p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j})}$, $m_{\mathbf{u}_{B_j} \rightarrow p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j})}$ are updated, and iii) factor-to-variable messages $m_{p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j}) \rightarrow \mathbf{u}_{B_i}}$, $m_{p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j}) \rightarrow \mathbf{u}_{B_j}}$ are updated.
- 3) **Decoupling:** The message exchanges in 2 are executed every timestep the two robots are within range. After separation, the child i stores the previous state of $m_{\mathbf{u}_{B_j} \rightarrow p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j})}$ and $p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j})$. The factor may be further conditioned on another robot’s inducing variables later. A unary factor with value $m_{p(\mathbf{u}_{B_i} | \mathbf{u}_{B_j}) \rightarrow \mathbf{u}_{B_j}}$ is added to \mathbf{u}_{B_j} .

In the limit of all-to-all communication, we highlight that this routine constitutes inference in a model with a full GP prior (12). However, there are two sources of approximation in our approach. First, robots which are always far away and never make contact may never be connected via a shared edge. These robots are treated as independent. Second, robots which connect and then move out of range will be propagating stale messages. Thus, the asynchronous nature

of the connections may limit performance. We empirically evaluate the impact of these approximations in Sec. IV-D.

4) *Distributed Hyperparameter Learning*: Kernel hyperparameters θ are typically tuned by gradient ascent of the log-marginal likelihood (see Sec. III-A). Bui and Turner [1] propose a message-passing routine to update kernel hyperparameters in TSGP. However, their method requires a full up-down sweeps of the tree to compute the gradient before it is applied. This is not applicable in our case, with cyclic models and sporadic connectivity. Instead, we compute stochastic, approximate marginal likelihood gradients based on each robot’s individual experience, and use this to update their local copy of the kernel hyperparameters. We interleave hyperparameter tuning and GBP, updating the GP factors using the latest hyperparameters at each stage.

5) *Prediction*: Prediction at a test point \mathbf{x}^* can be made directly from the posterior of a chosen block of inducing points as in (5). The choice of which block $q(\mathbf{u}_{B_i}|\mathbf{y})$ for a query \mathbf{x}^* should ideally be consistent with how observations are connected to inducing points during training. However, this will depend on the application and where $q(\mathbf{u}_{B_i}|\mathbf{y})$ are stored at prediction time. For DistGP, we explore two methods for mapping test queries to robots: i) predict at \mathbf{x}^* with the robot whose current location $\mathbf{x}_{r_i,t}$ is closest, $r^* = \operatorname{argmin}_{r_i \in \{r_1, \dots, r_R\}} \|\mathbf{x}^* - \mathbf{x}_{r_i,t}\|^2$ and ii) pre-allocate the input space regions $\{\mathcal{R}_i\}$ in which each robot is responsible for mapping and predicting: $r^* = r_i \iff \mathbf{x}^* \in \mathcal{R}_i$.

D. Evaluation of Asynchronous Model Construction

To quantify the impact of asynchronous model construction and GBP (as described in Sec. IV-C), we compare its performance to a batch-fitted model with a full GP prior over all inducing points. We generate the trajectories and observations as in Sec. IV-B, and again place inducing points on a regular grid within each robot’s cell. However, in this case the robots learn online, incorporating new observations, connecting and exchanging messages as per Sec. IV-C. Robots can only communicate with others in a $d_{\text{comm}} = 0.15$ radius of their current location. This distance is a small fraction of the overall input space covering $[-1, 1] \times [-1, 1]$. For comparison, we fit the batch models every 400 steps, using all observations made so far. We evaluate predictions in six randomly chosen 0.3×0.3 subregions which are “held out”: robots may move through these regions but do not make observations when inside them.

The results in Fig. 4 show that while the distributed implementation (blue) is slower to converge, it achieves the same result as the batch-fitted dense GP prior over all inducing points (green). Further, the distributed implementation is able to rapidly achieve performance superior to a batch-fitted TSGP (orange) and a batch fitted model without connections (red). These results suggest we can maintain accuracy while learning a distributed map asynchronously.

V. EXPERIMENTS

We now evaluate DistGP, testing it on two distributed mapping simulations: sea-surface temperature (SST) prediction and occupancy mapping. We compare against DiNNO [2].




Fig. 4: **Asynchronous + distributed vs batch + centralised model.** We compare i) online learning, ad-hoc inter-robot connections, ad-hoc GBP (blue) against ii) a centralised implementation with synchronous GBP sweeps on batch data (red, orange, green). The former converges to the same performance as the centralised, dense prior (green). Averages over 10 seeds, error bars cover ± 1 SE around the mean.

A. Dynamic Sea Surface Temperature Tracking

Many environmental monitoring applications require dynamic maps over large areas, in regions with poor connectivity. These challenges motivate a decentralised multi-robot mapping approach like ours. To evaluate our method in this setting, we simulate a multi-robot SST monitoring problem. The Optimally Interpolated SST (OISST) dataset [23], [24] combines satellite, buoy and ship measurements, and interpolates them on a dense grid of $1/4$ lat/lon resolution at daily frequency. We select 5 disjoint ocean regions, each covering a $4400 \times 4400 \text{ km}^2$ area, and divided into a 5×5 grid. We simulate 25 aquatic robots spread over the region, one per cell, which make pointwise SST measurements as they move around. We assume the robots move at a low speed of 7 km/h , can communicate with others within $d_{\text{comm}} = 1500 \text{ km}$ and observe the temperature every 3 hours. As we only have access to daily mean temperatures, we assume the field remains constant throughout each day.

The aim is for the combined robot predictions over the area to accurately match the ground truth map at any given point in time. As the robots move slowly and the map evolves with time, their observation coverage is limited. To do well, therefore, requires robots are able to exchange useful information with their peers to improve their predictions in less recently observed regions. Combined predictions are generated on a dense grid, with each robot predicting the SST at all test points inside their cell.

We train our method by adding observations online and incorporating their messages into the inducing point posterior as per Sec. IV-C. Each robot retains 5 observations in their model after which we fuse them into the unary factor on their inducing variables. Every other step, one of these observations is selected as a new inducing point via greedy variance selection. While this amounts to a relatively dense set of inducing points, we only retain the 200 most recent inducing points, as older inducing points become irrelevant to current predictions as the field evolves. When two robots in adjacent cells connect, they both add a line of inducing points close to their shared border to provide extra flexibility

Region	DistGP (Ours)	DiNNO
Central Atlantic	0.134 ± 0.001	0.175 ± 0.001
North Indian Ocean	0.160 ± 0.001	0.222 ± 0.001
South Indian Ocean	0.230 ± 0.001	0.316 ± 0.002
West Pacific	0.108 ± 0.001	0.166 ± 0.001
East Pacific	0.240 ± 0.002	0.303 ± 0.001

TABLE I: **SST prediction MSE.** Error in predicted temperatures over the map, averaged over all timesteps after a 25 day “burn-in”. Intervals cover \pm SE over 4 seeds.

Comm. Setting	DistGP (Ours)	DiNNO
Default	0.134 ± 0.001	0.175 ± 0.001
Comm. Interval 1 \rightarrow 5	0.138 ± 0.001	0.249 ± 0.001
Comm. Interval 1 \rightarrow 10	0.143 ± 0.001	0.294 ± 0.001
d_{comm} 1500km \rightarrow 1100km	0.143 ± 0.001	0.420 ± 0.005
d_{comm} 1500km \rightarrow 700km	0.169 ± 0.001	– [†]

TABLE II: **Varying communication.** MSE with different communication settings for the Central Atlantic region. [†]optimisation diverged. Intervals cover \pm SE over 4 seeds.

to make their fitted functions consistent. We use the product of two Matérn-1/2 kernels as our covariance function, one over 2D space and one over time; and we tune kernel θ as well as the observation noise σ with the distributed method described in Sec. IV-C.4.

The base model for DiNNO is a 6-layer MLP with width 64 and leaky-ReLU activations. At each timestep we train these networks using minibatch samples from the previous two weeks observations. This window length was found to perform best during validation. Hyperparameters for both methods were chosen based on simulations in the Central Atlantic region from 2022/01/01 to 2022/07/01. We evaluate global map accuracy on a dense grid at every 3 hour timestep between 2022/07/01 and 2023/03/15, following a 25 day burn-in period.

The results (Tab. I), show that DistGP significantly outperforms DiNNO for all regions, suggesting more effective information sharing. We further examine the robustness of both methods under different communication settings: i) less frequent (every 5 or 10 steps instead of every step), and ii) shorter range ($d_{\text{comm}} \in \{700\text{km}, 1100\text{km}\}$ instead of 1500km). In both cases, we find DiNNO suffers considerable degradation in accuracy where DistGP performs similarly to the more frequent, longer-range communication.

B. Online Occupancy Mapping

We now evaluate on the 2D online mapping simulation from Yu et al. [2]. The trajectories of 7 robots around a 2D environment (Fig. 5a) are predefined¹, and each robot makes a sweep of simulated LIDAR measurements from their position at each timestep. The aim is to use these observations, and inter-robot communication, to learn a map online. To simulate streaming data, Yu et al. [2] give each robot a buffer of the 400 most recent scans to train its model (total scans per trajectory $\sim 4 \times 10^3$). However, each robot

¹See Fig. 3 of [2] for trajectories

has to repeat its trajectory ~ 100 s of times to achieve a good map. We believe this is due to catastrophic forgetting in NNs, and ask whether our probabilistic representation enables greater efficiency.

We treat each LIDAR scan of ~ 500 observations as a batch, use it to update the inducing point posterior, then discard it as per Sec. IV-C.1 (in contrast to DiNNO’s 400 scan buffer). As the occupancy observations are discrete $y_n \in \{0, 1\}$, we use binary observation factors with energy equal to binary cross-entropy (BCE), and $p_n = \sigma(f(\mathbf{x}_n))$ as the predicted probability of occupancy. We linearise these factors to approximate outgoing messages as Gaussian.

Observing that the variation in the map is concentrated at the walls of the environment, we choose inducing points to be close to the walls. From each LIDAR scan we extract a set of candidate inducing points by concatenating: i) the input locations for which density is 1 (positive examples), and ii) positive examples perturbed a small distance back towards the robot (negative examples). We choose new inducing points from this shortlist using the greedy variance method.

Upon making contact, two robots connect via a shared factor (Sec. IV-C). Each robot also subsamples a set of the other’s nearby inducing point locations and creates new inducing points in their own model at these locations. Further, as the trajectories are partially overlapping, the robots exchange inducing point posteriors, allowing them to predict in regions observed by the other without having to incorporate all the other’s inducing points into their own model. All robots use a Matérn-1/2 kernel with fixed lengthscale.

While our method generates a map distributed between robots, DiNNO produces an estimate of the global map on each robot. We therefore use two modes of evaluation:

- 1) **Global map accuracy.** As in [2], we assess DiNNO by the average accuracy of all robots’ global maps. For our method, we construct the global map which would be achieved by the robots exchanging both i) their own inducing point posterior and ii) posteriors received from others, when connecting.
- 2) **Distributed map accuracy.** In each evaluation round, we assess the accuracy when predicting each test point using the closest robot. As the locations change over time, the same test point may be predicted by different robots over the course of the experiment.

The maps (Figs 5b, c and d) show that DistGP is able to converge to an accurate map in a single pass, i.e. with each robot traversing their trajectory once. In contrast, DiNNO must repeat the trajectories hundreds of times to get close to the ground truth. We further note that our final map more closely matches the ground truth under both metrics (Fig. 6), and the distributed and centralised predictions perform almost identically for both methods.

VI. CONCLUSION

We have presented DistGP: a method for multi-robot mapping in which each robot learns an sparse GP and communicates with others via GBP to ensure consistency. It extends TSGP [1] to include loops (which boost accuracy)




Fig. 5: **Distributed occupancy mapping.** Our method converges to an accurate map in a single pass through the robots’ trajectories, as observations are fully fused into the model. In contrast, DiNNO trains NNs iteratively and therefore requires many passes to overcome catastrophic forgetting. We note that our map at convergence is also more accurate than DiNNO’s.




Fig. 6: Test error during mapping (log-scale horizontal axis).

and support asynchronous GBP. Our experiments demonstrate the inducing point map representation enables effective online data fusion and robust inter-robot communication relative to NN-based DiNNO [2].

There are many possibilities to build on this work. First, GPs are known to provide accurate uncertainty estimates in many settings, making them effective surrogate models for Bayesian optimisation, and useful for active learning. Our method could be an exciting opportunity for distributed Bayesian optimisation or active learning, in which multiple robots explore different subregions of the input space in parallel. Second, we note that each robot learns kernel hyperparameters in isolation. A message passing procedure to facilitate collaborative estimation in loopy graphs could improve performance considerably.

VII. ACKNOWLEDGEMENTS

We are grateful to Javier Yu for his advice on tuning DiNNO. SN and AJD are funded by EPSRC Prosperity Partnerships (EP/S036636/1) and Dyson Technology Ltd.

REFERENCES

- [1] T. D. Bui and R. E. Turner, “Tree-structured gaussian process approximations,” *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [2] J. Yu, J. A. Vincent, and M. Schwager, “Dinno: Distributed neural network optimization for multi-robot collaborative learning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1896–1903, 2022.
- [3] 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.
- [4] D. Albani, J. IJsselmuiden, R. Haken, and V. Trianni, “Monitoring and mapping with robot swarms for agricultural applications,” in *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, 2017, pp. 1–6.
- [5] L. Yliniemi, A. K. Agogino, and K. Tumer, “Multirobot coordination for space exploration,” *AI Magazine*, vol. 35, no. 4, pp. 61–74, 2014.
- [6] Y. Deng, Y. Tang, Y. Yang, D. Wang, and Y. Yue, “Macim: Multi-agent collaborative implicit mapping,” *IEEE Robotics and Automation Letters*, vol. 9, no. 5, pp. 4369–4376, 2024.
- [7] M. Asadi, K. Zareinia, and S. Saeedi, “Di-nerf: Distributed nerf for collaborative learning with relative pose refinement,” *IEEE Robotics and Automation Letters*, 2024.
- [8] H. Zhao, B. Ivanovic, and N. Mehr, “Distributed nerf learning for collaborative multi-robot perception,” *arXiv preprint arXiv:2409.20289*, 2024.
- [9] —, “Ramen: Real-time asynchronous multi-agent neural implicit mapping,” *arXiv preprint arXiv:2502.19592*, 2025.
- [10] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, *et al.*, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [11] 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)*. IEEE, 2018, pp. 6359–6364.
- [12] G. Habibi and J. P. How, “Human trajectory prediction using similarity-based multi-model fusion,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 715–722, 2021.
- [13] R. Murai, J. Ortiz, S. Saeedi, P. H. Kelly, and A. J. Davison, “A robot web for distributed many-device localization,” *IEEE Transactions on Robotics*, vol. 40, pp. 121–138, 2023.
- [14] A. Patwardhan and A. J. Davison, “A distributed multi-robot framework for exploration, information acquisition and consensus,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 12 062–12 068.
- [15] E. Snelson and Z. Ghahramani, “Sparse gaussian processes using pseudo-inputs,” *Advances in neural information processing systems*, vol. 18, 2005.
- [16] L. Csató and M. Opper, “Sparse on-line gaussian processes,” *Neural computation*, vol. 14, no. 3, pp. 641–668, 2002.
- [17] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.
- [18] R. Gallager, “Low-density parity-check codes,” *IRE Transactions on information theory*, vol. 8, no. 1, pp. 21–28, 1962.
- [19] K. Murphy, Y. Weiss, and M. I. Jordan, “Loopy belief propagation for approximate inference: An empirical study,” *arXiv preprint arXiv:1301.6725*, 2013.
- [20] A. J. Davison and J. Ortiz, “Futuremapping 2: Gaussian belief propagation for spatial ai,” *arXiv preprint arXiv:1910.14139*, 2019.
- [21] Y. Weiss and W. Freeman, “Correctness of belief propagation in gaussian graphical models of arbitrary topology,” *Advances in neural information processing systems*, vol. 12, 1999.
- [22] D. R. Burt, C. E. Rasmussen, and M. Van Der Wilk, “Convergence of sparse variational inference in gaussian processes regression,” *Journal of Machine Learning Research*, vol. 21, no. 131, pp. 1–63, 2020.
- [23] R. W. Reynolds, T. M. Smith, C. Liu, D. B. Chelton, K. S. Casey, and M. G. Schlax, “Daily high-resolution-blended analyses for sea surface temperature,” *Journal of climate*, vol. 20, no. 22, pp. 5473–5496, 2007.
- [24] B. Huang, C. Liu, V. Banzon, E. Freeman, G. Graham, B. Hankins, T. Smith, and H.-M. Zhang, “Improvements of the daily optimum interpolation sea surface temperature (doisst) version 2.1,” *Journal of Climate*, vol. 34, no. 8, pp. 2923–2939, 2021.