Collective Ranking of Environmental Signals through Gaussian Belief Propagation in a Patrolling Robot Swarm

Zachary R. Madin[1][0009−0005−0310−4504], Connor York[1][0009−0001−8394−2671],
Jonathan Lawry[1][0000−0002−3488−5682], and Edmund R. Hunt[1][0000−0002−9647−124X]

University of Bristol, Bristol, UK `{zachary.madin, connor.york, j.lawry, edmund.hunt}@bristol.ac.uk`

Abstract. Multi-robot patrolling (MRP) requires a team to visit all areas of an environment at regular intervals, typically minimising idleness. A practical extension, motivated by security and environmental monitoring, is to additionally form a collective ranking of all patrol locations by some measured signal, a generalisation of the best-of-n problem to the many-option, continuous-valued regime. We observe that the patrol graph admits a natural dual interpretation: it is simultaneously the topology that dictates agent movement and a factor graph over which spatial beliefs can be propagated. Exploiting this equivalence, we apply Gaussian Belief Propagation (GBP), a graph based algorithm, to collective ranking using unary measurement factors at visited nodes and pairwise smoothness factors along patrol edges. We compare GBP against simple and visit-count-weighted averaging across a range of sensor-noise conditions in simulation, and validate the approach on four Leo Rovers tracking a propagating radio signal in an office lobby. GBP outperforms both baselines on ranking accuracy, mean squared error, and time to consensus. Crucially, as noise increases and the task becomes harder, GBP degrades gracefully in simulation while both averaging methods collapse. Hardware trials reproduce the same performance ordering on a real propagating radio signal, supporting the practical relevance of the simulated results.

Keywords: Environmental monitoring · Swarms · Information fusion · Patrolling · Gaussian Belief Propagation.

1 Introduction

Surveillance and environmental monitoring applications typically require a large number of areas of interest to be covered and measured repeatedly, and the signals of interest often exhibit spatial structure in which nearby measurements are correlated. Robot swarms are well suited to such problems: they can spread over large regions, pursue multiple objectives in parallel, and cross-verify measurements between agents. A swarm that independently patrols areas of interest while sharing measurements can cover an environment faster than a single agent and be more robust to sensor noise, and by exploiting the spatial structure of the signal it can also produce better estimates at noisy or infrequently visited locations. In this work, we examine how current patrolling methods perform when tasked with monitoring environmental conditions and anomalous signals.

In previous work, we examined how collective consensus formation during patrol can suppress noisy binary anomaly detections, and found that the communication connectivity emergent from a given patrolling algorithm materially affects both the speed and the accuracy of consensus [16]. A key open question left by that work is how these findings extend when the quantity being perceived is not a binary anomaly flag but a continuous-valued environmental signal, and when the collective task is not detection but ranking, which we address in this work.

Three information fusion methods of increasing sophistication are compared: a simple mean average of measurements serves as a baseline, followed by a weighted average based on the number of visits to a node, and finally Gaussian Belief Propagation (GBP) which is a graph-based method for distributed inference. Approaches for estimating spatial measurements of environmental variables frequently rely on Gaussian Processes (GPs), which are computationally expensive when the number of observations grows large. This is because with each additional measurement, the covariance matrix expands, making GPs poorly suited for distributed monitoring in a swarm. In contrast to this we employ GBP, which exploits local message passing between agents in the swarm and supports incremental node updates to infer environmental values. The graph structure of GBP encodes the spatial relationships between patrol nodes in the graph, allowing beliefs to propagate between adjacent locations. This enables spatial correlations in the signal to be exploited during inference, an advantage that simple or weighted averaging cannot provide as these methods treat measurements at each location as independent of one another.

We examine the behaviours of patrolling algorithms when the agents are tasked with building an accurate severity rank of signals in an environment. A simulated signal environment is used to examine the performance of the swarm under increasing difficulty of problem, as well as different environments. We analyse the performance of the swarms in terms of the time taken to reach a consensus and the accuracy at that given point. To demonstrate that these results hold for a real signal, a number of trials with Leo Rover robots tracking a radio source are performed for a sub-set of the simulation conditions (Figure 1).

The key contribution of this paper is the application of GBP to the collective ranking variant of the best-of-n problem in a multi-robot patrolling context, and the novel dual-purpose use of the patrol graph. The graph structure of the problem simultaneously defines agent movement topology and serves as the factor graph structure over which spatial beliefs are propagated. This work is then validated on physical robots with real propagating electromagnetic signals. The paper is organised as follows. Section 2 discusses the relevant literature, contextualising the work within the field. Section 3 presents the methodology of our results. Section 4 and 5 are the results from simulation and real experiments respectively. Section 6 concludes the paper with a summary of our findings and directions for future work.

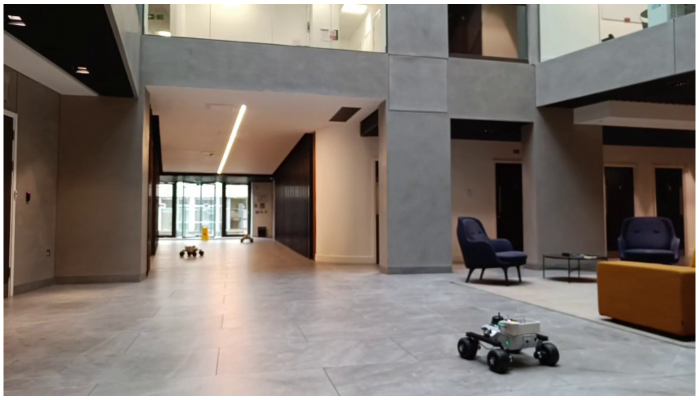

Fig. 1: Ground robots autonomously patrolling the lobby of an office building while taking radio measurements along a patrol graph during real trials.

2 Related Work

2.1 Patrolling Algorithms

The task of environmental monitoring for surveillance or security has been extensively studied in the multi-robot literature [11, 27, 3]. Most frequently the multi-agent patrolling problem is decomposed into a task allocation problem of a graph. The environment is abstracted into a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with the graph $\mathcal{G}$ consisting of $\mathcal{V}$ vertices and $\mathcal{E}$ edges. The task of monitoring the environment for a multi-agent system is to assign vertices to be visited to each agent while reducing the time since the vertex was last visited. This is defined as minimising the *"idleness"* within the patrolling literature. For the single agent instance, the task becomes a version of the Travelling Salesman Problem - the agent must visit every node at least once, minimising the time taken and overlap of distance traveled to achieve an optimal solution. This problem is known to be NP-hard in both the single and multi-agent case [23, 24], with improvements only recently being made on approximations [13].

In this work we employ two algorithms that are extensively used in the literature, chosen for their performance and comparability. Cyclic algorithm for Generic Graphs (CGG) is a method for computing Hamiltonian cycles when the patrol graph allows, and long paths to visit every node in as few steps as possible [6]. State Exchange Bayesian Strategy (SEBS) performs a form of Bayesian learning to estimate what the expected idleness of a node is, and determines which node to visit based on a utility function [29]. This algorithm enables agents to make local and independent decisions while sharing information of visit intentions with agents in their vicinity and is tolerant to communication limitations [37]. The choice of CGG and SEBS in this work is also informed by our prior comparison of ten multi-robot patrolling algorithms on the same "Cumberland" patrol graph (Figure 4b), in which SEBS emerged as a strong all-round performer in terms of both idleness minimisation and collective anomaly perception accuracy, while CGG produced a higher-connectivity emergent communication network that favoured rapid consensus at the expense of false-positive suppression [16]. Selecting these two algorithms therefore provides a principled contrast between a Bayesian utility-driven patrolling strategy and a fixed-cycle strategy whose behavioural differences are already well-characterised within our experimental framework.

In this work, the patrol graph serves a dual purpose: in addition to defining the topology of agent movement, it defines the discrete set of locations at which environmental measurements are taken, making patrolling the mechanism by which spatial observations are gathered for environmental modelling.

2.2 Environment Modeling

The task of monitoring a spatial environmental field using a multi-robot or swarm system has been applied across a wide range of application domains including ocean temperature mapping, pollution monitoring, agricultural sensing, and wildfire detection [10, 17]. A common and principled approach to modeling a continuous spatial field from discrete robot measurements is Gaussian process (GP) regression [30], which provides a probabilistic framework for interpolating between sampled locations and quantifying prediction uncertainty. GP-based methods have been applied to multi-robot systems for mapping scalar environmental fields such as sea-surface temperature, chemical gradients, and water quality parameters [14, 36, 18], and offer the advantage of naturally encoding spatial correlation structure into field estimates. However, a well-known limitation of standard GP regression is its computational complexity: inference requires the inversion of the full covariance matrix over all collected observations, an operation that scales as $\mathcal{O}(n^3)$ in the number of data points and $\mathcal{O}(n^2)$ in storage [30, 4], making it increasingly intractable as datasets grow. This cost is particularly problematic in a distributed multi-robot context, where no single agent maintains access to all observations and performing full covariance matrix inversion in a decentralised manner is a non-trivial challenge [36, 20]. Several approximation strategies have been proposed to address these limitations [18, 34], each offering trade-offs between tractability and modeling fidelity.

In settings where the primary objective is not continuous field reconstruction but rather the accumulation of point-value knowledge at specific locations in the environment as is naturally the case in the multi-robot patrolling problem a graph-structured representation of the environment offers a more appropriate and computationally lightweight abstraction. Here, the patrol graph itself defines the discrete set of nodes at which measurements are taken, and rather than fitting a global spatial model, agents can maintain and update local estimates at each node independently as it is revisited over time. A natural and tractable approach in this setting is to combine repeated measurements through a simple running average, accumulating observations across patrol cycles and refining node-level estimates as the patrol progresses [8, 2]. This approach has an analogue in the broader swarm and multi-robot exploration literature, where agents exploring an environment have been shown to build useful spatial representations through local averaging and consensus-based aggregation of sensor readings acquired during traversal [12, 2]. While such averaging methods lack the spatial interpolation properties and principled uncertainty quantification of full GP models, their simplicity and compatibility with distributed, memory-constrained robotic systems make them well-suited to the patrolling context, where the patrol graph structure already imposes a natural discretisation of the environment and the regular revisitation of nodes provides repeated opportunities to refine estimates over time.

2.3 Best-of-N and Collective Ranking

The best-of-n problem, as formalized by Valentini et al. [33], requires a swarm of agents to reach collective consensus on which of n available options best satisfies the needs of the collective. In the collective perception literature, this problem has been most commonly instantiated as a binary discrimination task, where agents must determine which of two colors (typically black or white) constitutes the majority of a tiled floor environment. In this abstraction, the quality of each option is directly represented by the proportion of the arena surface covered by a given color, with the ratio of black to white tiles representing the ground truth that the collective must perceive. While this formulation offers a tractable and well-studied benchmark for evaluating consensus formation strategies, it has limited applicability to real-world deployment scenarios. Security patrol environments, for example, are unlikely to present agents with neatly partitioned binary perceptual landscapes; instead, anomalies of interest may be spatially localized, intermittent, or confounded by environmental noise in ways that a uniform tile ratio does not capture.

Recent work has sought to address this limitation by introducing more principled and varied approaches to the construction of environmental patterns used in collective perception tasks. Benchmarking work by [5] proposed new task difficulty metrics to better characterize the challenge posed to a collective decision-making system, moving beyond simple majority ratios to consider the spatial arrangement and clustering of features in the environment. Their approach recognizes that two environments with identical black-to-white ratios may present very different perceptual challenges depending on how features are distributed a densely clustered arrangement, for instance, may lead to strongly biased local observations for individual agents, increasing the variance of individual measurements and slowing consensus formation. This approach to the formulation of a *hardness* of the best-of-n problem still however relies on a binary option and has limited applicability to real world ranking of options. In this work, we are interested in a strictly more general task: the collective ranking of all patrol nodes by estimated signal severity, rather than the identification of a single best option.

This distinction has been formalised by Crosscombe and Lawry [7], who extend the best-of-n framework to enable agents to learn a complete preference ordering over all n options. Shan and Mostaghim [31] directly address this collective ranking task in a spatially distributed robot swarm, comparing a Borda-count ranked voting strategy against a belief-fusion baseline across sweeps of sensor noise, evidence rate, and swarm size. They find that ranked voting is more noise-resistant and scales more favourably with swarm size than belief fusion, at the cost of longer convergence and higher scatter in the swarm's final opinion distribution. They attribute this trade-off to the strong positive-feedback dynamics of fusion-based consensus, which can lock a swarm onto an incorrect option before sufficient evidence has been accumulated: an effect that is exacerbated under high noise or in large swarms. Their three-metric evaluation framework of accuracy, consensus uniformity, and convergence time motivates our use of MSE, Spearman rank correlation, and time-to-consensus as performance metrics in Section 4 and 5.

A further line of work by Shan and Mostaghim [32] considers the many-option regime in which the number of discrete options exceeds the swarm size: the relevant regime for our setting, where 40 patrol nodes are collectively ranked by a swarm of 8 agents. They compare discrete consensus strategies (ranked voting and distributed Bayesian belief sharing) against a continuous Linear Consensus Protocol (LCP) baseline, and show that the hardness of the consensus task depends on whether option qualities are correlated. When options are weakly correlated, discrete strategies are accurate but slow; when option qualities are strongly correlated or spatially concentrated, fusion mechanisms that treat options as independent become more vulnerable to premature convergence on an incorrect consensus. This distinction is directly relevant to the present work: the radially structured signal field generated by a single electromagnetic emitter produces a spatially correlated option-quality structure that a fusion method exploiting graph topology (such as GBP) should be able to leverage.

2.4 Averaging Measurements

In this work, we compare GBP against maintaining a running estimate of the environmental signal at each patrol graph node using two averaging methods. In the first, a simple mean is computed across all measurements recorded at a given node. In the second, a weighted mean is used where the contribution of each robot's local estimate is weighted by the number of visits that robot has made to that node. This results in agents with more numerous visits to a given location exerting proportionally more influence on the fused estimate. Given R robots each holding a local estimate $\hat{z}_i^r$ at node i with associated visit count v_i^r, the resultant weighted mean estimate is given by:

$$\hat{z}_i = \frac{\sum_r v_i^r \, \hat{z}_i^r}{\sum_r v_i^r} \tag{1}$$

The simple mean averages each robot's running estimate with equal weight, $\hat{z}_i = \frac{1}{R}\sum_r \hat{z}_i^r$, and so corresponds to Eq. 1 with $v_i^r = 1$ for all r.

The weighted and simple averaging baselines employed here can be viewed as discrete-graph instantiations of linear consensus [21], and serve a similar role to the LCP baseline used by Shan and Mostaghim [32] in their comparison of discrete and continuous consensus strategies. Both averaging methods aggregate observations node-by-node without reference to the graph's edge structure, and therefore provide a natural baseline against which to isolate the contribution of exploiting the patrol-graph topology during fusion.

2.5 Gaussian Belief Propagation

Gaussian Belief Propagation (GBP) is a message-passing algorithm for performing distributed inference over a factor graph, where the goal is to compute the marginal distribution of each variable given a set of observed measurements [26, 15]. In the context of robotics and multi-agent systems, GBP is particularly attractive because it operates entirely through local computation and nearest-neighbour message passing, requiring no centralised processing or global knowledge of the graph structure [9]. More recently, GBP has attracted interest in the context of distributed multi-robot systems, where its local message-passing structure aligns well with the communication constraints of robot swarms [25, 12]. Each node in the factor graph maintains a Gaussian belief over its associated variable, which is iteratively refined by receiving messages from its neighbours without requiring complete environment information. This makes GBP a natural fit for distributed robotic systems where communication is constrained to local interactions.

A factor graph encodes a joint distribution over unknown variables as a product of local factors, with *variable nodes* representing quantities to be estimated and *factor nodes* encoding probabilistic constraints or measurements [15]. Restricting to linear Gaussian factors ensures that all messages remain Gaussian throughout inference and can be computed in closed form [22]. In the patrol graph setting, variable nodes correspond to patrol graph nodes, each representing the estimated environmental quantity at that location. Unary measurement factors are added as robots visit nodes, progressively tightening each node's belief over time. The patrol graph topology is incorporated as binary pairwise difference factors between adjacent variable nodes, where edge weights are converted to factor precisions: stronger topological connections impose tighter spatial smoothness constraints between neighbouring estimates.

Inference proceeds via the standard sum-product algorithm [15], with messages represented in canonical (precision) form as precision vector $\boldsymbol{\eta}$ and precision matrix $\boldsymbol{\Lambda}$, related to the mean and covariance by $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$ and $\boldsymbol{\eta} = \boldsymbol{\Lambda}\boldsymbol{\mu}$. In this form, the belief at each variable node reduces to a sum of incoming precision vectors and matrices [22], with the mean estimate recovered as $\boldsymbol{\mu}_{b_i} = \boldsymbol{\Lambda}_{b_i}^{-1}\boldsymbol{\eta}_{b_i}$.

In a tree-structured factor graph, a single forward-backward pass yields exact inference [26]. However, patrol graphs are not trees and contain cycles, giving rise to Loopy Belief Propagation (LBP), in which messages circulate iteratively and convergence is no longer guaranteed in general [35, 19]. In practice, LBP has been shown to converge and produce accurate marginals across a wide range of applications [19]. To promote stable convergence in loopy graphs, *damping* is applied to the message update rule, whereby each new message is computed as a weighted combination of the freshly computed message and the previous message at that edge:

$$\mu^{(t+1)} = (1-\alpha)\,\mu_{\text{new}}^{(t+1)} + \alpha\,\mu^{(t)}, \tag{2}$$

where $\alpha \in [0, 1)$ is the damping factor and $\mu^{(t)}$ denotes the message at iteration t [22]. Damping reduces the rate at which information propagates around cycles, preventing the runaway reinforcement that can cause divergence in graphs with strong cyclic topologies. In this work, the damping factor of $\alpha = 0.25$ was determined empirically in order to balance between oscillations around a value and a fast convergence time.

3 Methods

3.1 Simulation Environment Initialisation

This work exists in the security context of detecting concealed radio frequency emitters in a patrolled environment. We assume the presence of a sensor that returns a *severity* value representing the strength of a detected electromagnetic signal, bounded within the interval $[0, 1]$. Node severity values are drawn from a Beta distribution $\mathrm{Beta}(\alpha, \beta)$, which is well-suited to this purpose given its support over the unit interval and its flexibility in representing a wide range of unimodal and skewed distributions through the choice of shape parameters. An initial examination of multiple $\alpha = \beta$ values was performed as seen in Figure-2. Each node in the patrol graph is assigned one of these sampled severity values, providing a continuous-valued analogue to the discrete option qualities typically considered in the best-of-n literature [33], where the number of distinguishable options is small and their qualities are clearly separated. Here, we instead examine the best-of-n problem over a continuous interval, where the difficulty of correctly ordering node intensities depends on the degree of overlap between their underlying distributions.

To quantify the difficulty of the classification task, we characterise the *misordering probability* the probability that a noisy sensor reading causes an agent to incorrectly order a pair of node values as a function of the standard deviation σ of the sensing noise. Within our simulation for simplicity we assume the noise of a measured signal is Gaussian. For a given σ, we model each node's observed severity as a Gaussian distribution centred on its true Beta-sampled value with standard deviation σ, and compute the pairwise Bayes error between all node distributions. A Monte Carlo simulation is performed across a range of σ values to produce an empirical mapping from noise level to expected misordering probability, providing a principled basis for selecting experimental noise conditions that span a meaningful range of task difficulty. The noise levels used in our experiments, expressed as mis-ordering probabilities ranging from 0% to 45% in increments of 5%, are selected according to this mapping and reported in Table 1.

To mirror the realistic spatial structure of an electromagnetic emitter, the Beta-sampled severity values are assigned to patrol graph nodes according to their radial distance from a simulated signal centre, such that nodes closer to the centre receive higher severity values. This introduces a spatial correlation into the node value distribution that reflects the physically motivated assumption that signal strength decays with distance from the source, and provides a structured environmental field that can be exploited by fusion algorithms that account for spatial relationships between nodes.

3.2 Simulator

The simulator used in this paper is written in Python, and is a discrete 2D grid world where agents are represented by occupying a single cell. They can move to all eight surrounding cells, and avoid colliding with walls and other agents. Each patrolling algorithm is implemented naively as well as each of the fusion methods. At time zero, agents determine which node to visit on their patrol, and start navigating to it synchronously. Once an agent arrives at a node a measurement is taken from the node's normal distribution, which is then incorporated to the fusion method. Communication happens inter-agent without a central controller and parameters such as the communication frequency and range can be varied. The simulator uses existing maps of varying size and graph dimension that are frequently used in the literature [28].

4 Simulation Results

This section will discuss the results gained from the simulation experiments. A brief note will be made on the communication range and the influence it has on convergence and performance. The same algorithms and fusion methods were tested on two different maps with different topology: one in simulation (this section) and one real-world map (Section 5). We examined the performance of each combination of patrolling algorithm and fusion method 20 times. The parameters for the simulation are shown in Table 1.

Consensus convergence is assessed by computing the standard deviation σ_{inter} of node value estimates across all agents at each timestep. When $\sigma_{inter} < \varepsilon$ for all nodes in the patrol graph, the system is considered to have reached consensus, indicating that inter-agent estimates have stabilised and are no longer varying significantly across the robot population.

Communication range was found empirically to have no significant effect on fusion performance beyond a minimum threshold required for consensus; global communication was therefore adopted across all experiments to isolate the effects of patrolling algorithm and fusion method, as different patrolling algorithms exhibit different minimum communication requirements that would otherwise hinder direct comparison between algorithms. Studying the interaction between fusion method and emergent connectivity in the ranking setting is left for future work.

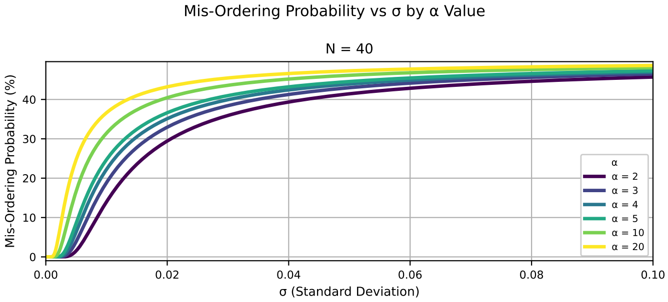

Fig. 2: Monte Carlo simulation of mis-ordering probability using Bayes Error for different standard deviations of normal distributions used in node means

Table 1: Simulation parameter values used across all experiments.

Parameter	**Value**
Number of robots	8
Communication range	Global
Communication timeout method	Upon new Data
Duration of experiment	10,000 steps or consensus
Patrol graph nodes	40
Mis-ordering probability	0 - 45%
Simulations per algorithm	20

4.1 Performance Metrics

Performance is evaluated across three metrics. Convergence time is recorded as the number of simulation timesteps until the variance of estimates across all agents falls below an empirically determined threshold ε for every node in the patrol graph, indicating that the swarm has reached a stable consensus. Once consensus has been reached, the following metrics determine performance of the swarm. Summed mean squared error is computed by summing the squared difference between each agent's estimate and the ground truth value across all agents and nodes. Spearman's rank correlation coefficient ρ measures the distance between the swarm's collective node ordering and the ground truth ranking.

4.2 Results

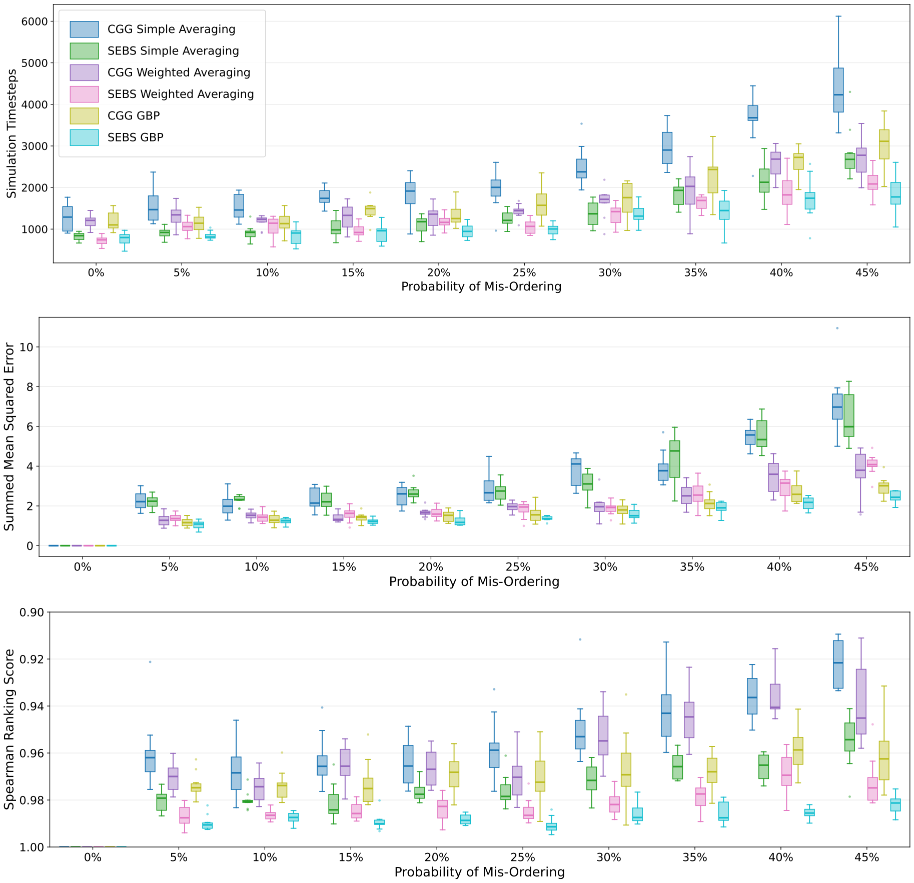

Fig. 3: Performance metrics by patrol algorithm and fusion method, for varying mis-ordering probabilities with an 8 robot swarm. Top: Consensus time. Middle: Mean Squared Error summed across the simulated swarm. Bottom: Spearman Ranking Score (axis inverted so that lower indicates better performance).

Across all conditions, the three fusion methods exhibit a clear performance hierarchy (Figure 3). Simple averaging performs worst across all three metrics, with both MSE and Spearman rank distance degrading substantially as mis-ordering probability increases beyond 20%. Weighted averaging consistently outperforms simple averaging, achieving lower MSE and better ranking accuracy with a reduced time to consensus, reflecting the benefit of weighting agent estimates by visit count. GBP outperforms both averaging methods across all metrics at all noise levels. Of the two patrolling algorithms, SEBS outperforms CGG across both performance metrics and reaches consensus in fewer timesteps, suggesting that its Bayesian utility-driven node selection produces more informative measurement sequences than the fixed cyclic traversal of CGG. Notably, GBP demonstrates a robustness to increasing noise that the averaging methods do not. As mis-ordering probability increases beyond 25%, both averaging methods show significant degradation in MSE and Spearman ranking score, while GBP maintains accurate performance with minimal degradation even at the highest noise conditions tested. This is attributable to GBP's exploitation of spatial correlations encoded in the patrol graph topology, which allows beliefs at noisy or infrequently visited nodes to be informed by estimates at neighbouring nodes, an advantage unavailable to methods that treat each node's measurements independently

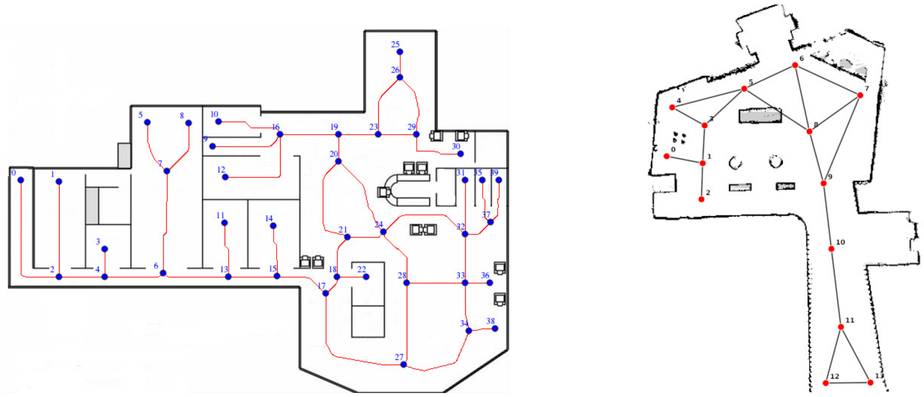

(a)

(b)

Fig. 4: Patrol graph environments used in experiments: (a) "Cumberland" (40 nodes, simulation trials) and (b) Bristol Lobby (14 nodes, real-robot trials).

5 Real Robot Experiments

5.1 Radio Signal Used

The conclusions from the simulated results were examined in a real world environment, with the agents sampling an environment with a propagating radio signal (Figure 1). In these experiments, a simple measure of signal intensity (RSSI) is used to represent the power of the signal received by the agents. Each agent performs their role of patrolling, while taking measurements when arriving at a node. The measurement values are fed into the three fusion methods examined in the simulation results, and compared for the Mean Squared Error, Spearman Rank distance and time to convergence. As the behaviour of the swarm is not dictated by the measurements but by the patrolling algorithm, all three methods are computed for the same trial data and then plotted for comparison. Ten trials of the patrolling algorithm SEBS were performed in the environment.

For the signal to be measured, an XBee S2C Zigbee [1] is used with a single static device acting as the transmitter, and each agent having a device as a receiver. The XBee Zigbee device operates on the 2.4GHz frequency, and transmits with a power of 3.1mW (+5 dBm). Signal packets are transmitted at 20Hz, and the receiver sensitivity is −100dBm. This device and method were selected to satisfy the size and power constraints of the real robotic platform, while the moderately low transmit power ensures a detectable decay in signal strength across the environment, producing the desired signal propagation pattern. The ground truth of the signal prior to experiment for ranking purposes was determined by taking repeated recordings over time and averaging across values.

The environment in which the robots patrol and sample the radio signal is an office lobby covering an area of 20 × 14 m. Nodes on the graph structure are distributed such that the distance between neighbouring nodes is no greater than 5 meters to ensure fair coverage of the environment. A representation of the navigation graph and mapped environment can be seen in Figure 4(b).

A small swarm of homogeneous robotic agents is used in this patrolling trial, each consisting of a non-holonomic rover-style platform. Four Leo Rovers are used in this experiment, each equipped with a battery pack, LiDAR, XBee S2C Zigbee transceiver, Latte Panda SBC and Raspberry Pi, performing combined high-level SLAM and navigation.

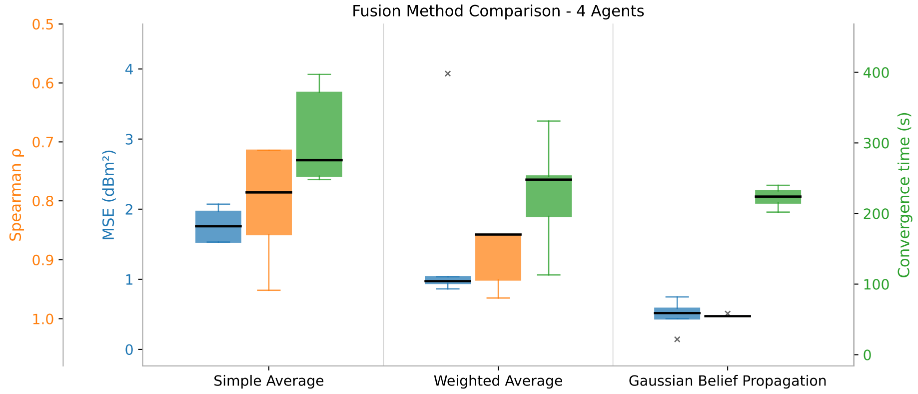

Fig. 5: Comparison between fusion methods used in real trial data with only SEBS patrolling algorithm. Mean Squared Error, Spearman Ranking ρ on left vertical axis (inverted), and convergence time on right vertical axis

Simulation operates on normalised severity values in [0,1] while real trials measure RSSI in dBm, so absolute MSE magnitudes are not directly comparable across the two settings; only the relative ordering of fusion methods is. As shown in Figure 5, the hierarchy observed in simulation is preserved on hardware: simple averaging performs worst, weighted averaging improves on it, and GBP achieves the lowest MSE, highest Spearman ranking score, and shortest median convergence time. This holds despite real signal propagation being subject to multipath interference, environmental occlusion, and hardware noise not captured in simulation.

6 Conclusion & Future Work

We examined a dual-objective robot swarm performing patrolling and environmental monitoring in a security context, framing the task as a collective ranking problem over continuous-valued signals at a large number of patrol nodes. Using two established idleness-minimisation algorithms, we compared Gaussian Belief Propagation against simple and weighted averaging as fusion methods, and found that GBP yields more accurate ranking and faster convergence at the cost of additional computation. The same performance hierarchy was observed when deploying the system on real robots tracking a propagating radio signal, establishing GBP as a reliable distributed method for collective environmental monitoring. These results are directly actionable for security applications in which a human operator, with limited time to verify a swarm's output, must prioritise the most severe readings.

Our findings align with and extend observations from the discrete collective decision-making literature. Shan and Mostaghim [31, 32] argue that fusion strategies which enforce consensus through positive feedback can lock a swarm onto an incorrect collective decision when measurements are noisy or option qualities are spatially correlated. This is due to treating each option independently and reinforcing shared beliefs among nearby agents. The averaging methods considered here exhibit this failure mode as measurement noise increases: they treat each patrol node as an independent estimation problem and collapse onto locally dominant but globally incorrect beliefs. GBP, by contrast, replaces pure consensus enforcement with a structural constraint: the patrol graph acts as a factor graph that propagates beliefs between adjacent variable nodes, so a noisy or infrequently visited node inherits information from its better-sampled neighbours, allowing the swarm to converge on a globally consistent estimate.

A natural direction for future work is to incorporate agent pose uncertainty into the fusion framework. This problem formulation lends itself well to GBP, where the uncertainty of measurements are added as an additional two dimensions to the inference framework. This addition leaves the distributed message passing benefits unchanged, while naturally propagating pose uncertainty into node estimates. This would enable equivalent environmental monitoring on platforms with less accurate localisation, in turn making larger swarm sizes tractable.

Acknowledgments. ZRM is supported by a University of Bristol PhD Scholarship. CY is supported by UK FCDO Services. ERH is supported by the Royal Academy of Engineering under the Research Fellowship programme. Simulations were performed on University of Bristol Self-Service Cloud.

References

1. XBee/XBee-PRO® S2C Zigbee® RF Module User Guide, https://docs.digi.com/resources/documentation/digidocs/90002002/default.htm
2. Albani, D., IJsselmuiden, J., Haken, R., Trianni, V.: Monitoring and Mapping with Robot Swarms for Agricultural Applications. In: 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS). IEEE (2017)
3. Almeida, A., Ramalho, G., Santana, H., Tedesco, P., Menezes, T., Corruble, V., Chevaleyre, Y.: Recent Advances on Multi-agent Patrolling. In: Bazzan, A.L.C., Labidi, S. (eds.) Advances in Artificial Intelligence – SBIA 2004. pp. 474–483. Lecture Notes in Computer Science, Springer, Berlin, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28645-5_48
4. Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D.W., O'Neil, M.: Fast Direct Methods for Gaussian Processes. IEEE Transactions on Pattern Analysis and Machine Intelligence **38**(2), 252–265 (2015), publisher: IEEE
5. Bartashevich, P., Mostaghim, S.: Benchmarking Collective Perception: New Task Difficulty Metrics for Collective Decision-Making. In: Moura Oliveira, P., Novais, P., Reis, L.P. (eds.) Progress in Artificial Intelligence, vol. 11804, pp. 699–711. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-30241-2_58
6. Chevaleyre, Y.: Theoretical analysis of the multi-agent patrolling problem. In: Proceedings. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2004. (IAT 2004). pp. 302–308 (Sep 2004). https://doi.org/10.1109/IAT.2004.1342959
7. Crosscombe, M., Lawry, J.: Collective preference learning in the best-of-n problem: From best-of-n to ranking n. Swarm Intelligence **15**(1-2), 145–170 (Jun 2021). https://doi.org/10.1007/s11721-021-00191-9
8. Crosscombe, M., Lawry, J.: The Impact of Network Connectivity on Collective Learning. In: Distributed Autonomous Robotic Systems. pp. 82–94. Springer Proceedings in Advanced Robotics, Springer (2022)
9. Davison, A.J., Ortiz, J.: FutureMapping 2: Gaussian Belief Propagation for Spatial AI. arXiv preprint arXiv:1910.14139 (2019)
10. Dunbabin, M., Marques, L.: Robots for Environmental Monitoring: Significant Advancements and Applications. IEEE Robotics & Automation Magazine **19**(1), 24–39 (2012), publisher: IEEE
11. Huang, L., Zhou, M., Hao, K., Hou, E.: A survey of multi-robot regular and adversarial patrolling. IEEE/CAA Journal of Automatica Sinica **6**(4), 894–903 (Jul 2019). https://doi.org/10.1109/JAS.2019.1911537, conference Name: IEEE/CAA Journal of Automatica Sinica
12. Jones, S., Hauert, S.: Distributed Spatial Awareness for Robot Swarms. Autonomous Robots (2025). https://doi.org/10.1007/s10514-025-10228-1
13. Karlin, A.R., Klein, N., Gharan, S.O.: A (slightly) improved approximation algorithm for metric TSP. In: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. pp. 32–45. STOC 2021, Association for Computing Machinery, New York, NY, USA (Jun 2021). https://doi.org/10.1145/3406325.3451009
14. Lin, T.X., Al-Abri, S., Coogan, S., Zhang, F.: A Distributed Scalar Field Mapping Strategy for Mobile Robots. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 11581–11586. IEEE, Las Vegas, NV, USA (Oct 2020). https://doi.org/10.1109/IROS45743.2020.9340836

15. Loeliger, H.A.: An Introduction to Factor Graphs. IEEE Signal Processing Magazine **21**(1), 28–41 (2004), publisher: IEEE
16. Madin, Z.R., Lawry, J., Hunt, E.R.: Collective Anomaly Perception During Multi-Robot Patrol: Constrained Interactions Can Promote Accurate Consensus. In: Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing (SAC '24). pp. 1–8. ACM (2024). https://doi.org/10.1145/3605098.3635975
17. Mansfield, J., Montazeri, A.: A Survey on Autonomous Environmental Monitoring Approaches: Towards Unifying Active Sensing and Reinforcement Learning. Frontiers in Robotics and AI **11** (2024). https://doi.org/10.3389/frobt.2024.1336612
18. Mishra, R., Koay, T.B., Chitre, M., Swarup, S.: Multi-USV Adaptive Exploration Using Kernel Information and Residual Variance. Frontiers in Robotics and AI **8** (May 2021). https://doi.org/10.3389/frobt.2021.572243, publisher: Frontiers
19. Murphy, K.P., Weiss, Y., Jordan, M.I.: Loopy Belief Propagation for Approximate Inference: An Empirical Study. In: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence. pp. 467–475. Morgan Kaufmann (1999)
20. Nabarro, S., Wilk, M.v.d., Davison, A.J.: A Distributed Gaussian Process Model for Multi-Robot Mapping (Mar 2026). https://doi.org/10.48550/arXiv.2603.07351, arXiv:2603.07351 [cs]
21. Olfati-Saber, R., Fax, J.A., Murray, R.M.: Consensus and Cooperation in Networked Multi-Agent Systems. Proceedings of the IEEE **95**(1), 215–233 (Jan 2007). https://doi.org/10.1109/JPROC.2006.887293
22. Ortiz, J., Evans, T., Davison, A.J.: A visual introduction to Gaussian Belief Propagation (Jul 2021). https://doi.org/10.48550/arXiv.2107.02308, arXiv:2107.02308 [cs]
23. Papadimitriou, C.H.: The Euclidean travelling salesman problem is NP-complete. Theoretical Computer Science **4**(3), 237–244 (Jun 1977). https://doi.org/10.1016/0304-3975(77)90012-3
24. Pasqualetti, F., Franchi, A., Bullo, F.: On optimal cooperative patrolling. In: 49th IEEE Conference on Decision and Control (CDC). pp. 7153–7158 (Dec 2010). https://doi.org/10.1109/CDC.2010.5717873, iSSN: 0191-2216
25. Patwardhan, A., Murai, R., Davison, A.J.: Distributing Collaborative Multi-Robot Planning with Gaussian Belief Propagation. IEEE Robotics and Automation Letters **8**(2), 552–559 (2023), publisher: IEEE
26. Pearl, J.: Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. In: Proceedings of the Second AAAI Conference on Artificial Intelligence. pp. 133–136. AAAI Press (1982)
27. Portugal, D., Rocha, R.: A Survey on Multi-robot Patrolling Algorithms. In: Camarinha-Matos, L.M. (ed.) Technological Innovation for Sustainability, vol. 349, pp. 139–146. Springer Berlin Heidelberg, Berlin, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19170-1_15
28. Portugal, D., Rocha, R.P.: On the performance and scalability of multi-robot patrolling algorithms. In: 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics. pp. 50–55 (Nov 2011). https://doi.org/10.1109/SSRR.2011.6106761, iSSN: 2374-3247
29. Portugal, D., Rocha, R.P.: Distributed multi-robot patrol: A scalable and fault-tolerant framework. Robotics and Autonomous Systems **61**(12), 1572–1587 (Dec 2013). https://doi.org/10.1016/j.robot.2013.06.011
30. Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. MIT Press (2006)

31. Shan, Q., Mostaghim, S.: Noise-resistant and scalable collective preference learning via ranked voting in swarm robotics. Swarm Intelligence (Sep 2022). https://doi.org/10.1007/s11721-022-00214-z
32. Shan, Q., Mostaghim, S.: Many-option collective decision making: discrete collective estimation in large decision spaces. Swarm Intelligence **18**(2), 215–241 (Sep 2024). https://doi.org/10.1007/s11721-024-00239-6
33. Valentini, G., Ferrante, E., Dorigo, M.: The Best-of-n Problem in Robot Swarms: Formalization, State of the Art, and Novel Perspectives. Frontiers in Robotics and AI **4** (2017). https://doi.org/10.3389/frobt.2017.00009
34. Viseras, A., Wiedemann, T., Manss, C., Magel, L., Mueller, J., Shutin, D., Merino, L.: Decentralized multi-agent exploration with online-learning of Gaussian processes. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). pp. 4222–4229 (May 2016). https://doi.org/10.1109/ICRA.2016.7487617
35. Weiss, Y., Freeman, W.T.: Correctness of Belief Propagation in Gaussian Graphical Models of Arbitrary Topology. Neural Computation **13**(10), 2173–2200 (2000), publisher: MIT Press
36. Xu, Y., Choi, J.: Spatial prediction with mobile sensor networks using Gaussian processes with built-in Gaussian Markov random fields. Automatica **48**(8), 1735–1740 (Aug 2012). https://doi.org/10.1016/j.automatica.2012.05.029
37. Yan, C., Zhang, T.: Multi-robot patrol: A distributed algorithm based on expected idleness. International Journal of Advanced Robotic Systems **13**(6), 1729881416663666 (Dec 2016). https://doi.org/10.1177/1729881416663666, publisher: SAGE Publications