



Madin, Z. R., Lawry, J., & Hunt, E. R. (2024). Collective Anomaly Perception During Multi-Robot Patrol: Constrained Interactions Can Promote Accurate Consensus. In *SAC '24: Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing* (pp. 630-637). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3605098.3635975>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1145/3605098.3635975](https://doi.org/10.1145/3605098.3635975)

[Link to publication record on the Bristol Research Portal](#)
PDF-document

University of Bristol – Bristol Research Portal

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/>



Collective Anomaly Perception During Multi-Robot Patrol: Constrained Interactions Can Promote Accurate Consensus

Zachary R. Madin
University of Bristol
Bristol, United Kingdom
zachary.madin@bristol.ac.uk

Jonathan Lawry
University of Bristol
Bristol, United Kingdom
j.lawry@bristol.ac.uk

Edmund R. Hunt
University of Bristol
Bristol, United Kingdom
edmund.hunt@bristol.ac.uk

ABSTRACT

An important real-world application of multi-robot systems is multi-robot patrolling (MRP), where robots must carry out the activity of going through an area at regular intervals. Motivations for MRP include the detection of anomalies that may represent security threats. While MRP algorithms show some maturity in development, a key potential advantage has been unexamined: the ability to exploit collective perception of detected anomalies to prioritize the location ordering of security checks. This is because noisy individual-level detection of an anomaly may be compensated for by group-level consensus formation regarding whether an anomaly is likely to be truly present. Here, we examine the performance of unmodified idleness-based patrolling algorithms when given the additional objective of reaching an environmental perception consensus via local pairwise communication and a quorum threshold. We find that generally, MRP algorithms that promote physical mixing of robots, as measured by a higher connectivity of their emergent communication network, reach consensus more quickly. However, when there is noise present in anomaly detection, a more moderate (constrained) level of connectivity is preferable because it reduces the spread of false positive detections, as measured by a group-level F-score. These findings can inform user choice of MRP algorithm and future algorithm development.

CCS CONCEPTS

• Computer systems organization → Robotic autonomy.

KEYWORDS

Multi-robot patrol, collective learning, anomaly perception, collective consensus

ACM Reference Format:

Zachary R. Madin, Jonathan Lawry, and Edmund R. Hunt. 2024. Collective Anomaly Perception During Multi-Robot Patrol: Constrained Interactions Can Promote Accurate Consensus. In *The 39th ACM/SIGAPP Symposium on Applied Computing (SAC '24)*, April 8–12, 2024, Avila, Spain. ACM, New York, NY, USA, Article 4, 8 pages. <https://doi.org/10.1145/3605098.3635975>

1 INTRODUCTION

An important application for multi-robot systems (MRS) is multi-robot patrol (MRP), where multiple robots are required to somehow

coordinate their behaviour such that all points in the environment requiring surveillance are visited regularly. Several multi-robot patrol algorithms have been developed over the past 20 years or so, which commonly aim to minimize the maximum and average ‘idleness’ [20] (return times) on each vertex in a patrol graph, using a variety of centralized and distributed robot coordination methods [13, 24]. Depending on the real-world deployment context, one aim of such patrols might be to detect subtle anomalies representing a potential security threat, for example the electromagnetic activity of an eavesdropping device planted by an adversary. Such anomalies, when detected, are quite likely to be false positives, i.e. to be false alarms caused by noise in sensors or in the environment. An important opportunity for MRS, then, is to help prioritize limited security resources, by advising a user on system-level consensus on an anomaly’s presence. In the course of MRP, one or more robots can make more than one inspection of an anomaly to obtain more information, and/or the detection system can benefit from the mechanisms of consensus formation to slow the spread of incorrect opinions (‘misinformation’). Additional anomaly inspections may be deliberate (i.e. planned by the robots following an initial detection) or occur spontaneously in the course of ongoing patrol. In an adversarial context, it would be important not to unduly compromise idleness minimization to make additional inspections, as this could be a source of vulnerability in a multi-stage, multi-location attack [30]. In the first instance, then, one may wish to retain the primary objective of idleness minimization, while allowing a consensus formation process regarding anomaly location to take place during the course of normal, longer-term operations. Maintaining typical patrolling behaviours may also be less likely to alert an adversary to their discovery.

In this work, we use a realistic ROS-based MRP simulator to examine the performance of unmodified patrolling algorithms when given the secondary objective of reaching a perceptual consensus on anomaly presence, via local pairwise communication and a quorum decision threshold. We consider the accuracy of the system-level perception in relation to the level of noise present in the anomaly detection process. Overall, MRP algorithms that promote more physical mixing of robots have faster consensus, given the assumption of local communications. However, when anomaly detection noise increases, a more moderate communication connectivity is preferable to suppress false positives. These findings could help to inform user choice of patrolling algorithm, given expectations around detection noise and attack frequency. We provide more background on relevant research in Section 2, and detail our methodology in terms of simulator, collective decision-making mechanisms, and graph metrics in Section 3. In Section 4 we present simulation results and in Section 5 we provide some conclusions and plans for future work.



This work is licensed under a Creative Commons Attribution International 4.0 License. <https://creativecommons.org/licenses/by/4.0/>
SAC '24, April 8–12, 2024, Avila, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0243-3/24/04.
<https://doi.org/10.1145/3605098.3635975>

2 BACKGROUND

We briefly review MRP algorithms and collective decision-making and perception. We also consider a graph-based metric to describe the emergent communication network between patrolling robots with local belief exchange.

2.1 Multi-Robot Patrolling Algorithms

MRP generally focuses on a static environment to be patrolled. This is defined as an area of interest that has nodes that are to be visited, in order to minimize idleness of each node. The term *instantaneous node idleness* was proposed by Machado et al. [20] and is used extensively in the MRP field. The application of this metric is to define the number of cycles or time since a node was last visited. In addition to this term there is also *'instantaneous graph idleness'* which is the average idleness of each node in the entire region or graph. Given a map or graph to patrol by an MRS (e.g. Figure 1), one fundamental problem to solve is the apportionment of locations (nodes) to visit by different robots, where efficient visitation will physically spread the agents between nodes. One approach is an auction-based system where each robot submits a bid either to a central auctioneer [16] or between robots in a local proximity [5]. Much work has been done to take developments from Game Theory and apply it to the robot patrol problem. Owing to the malicious nature of an attacker in a system, different patrolling strategy approaches are required [32]. With predictability comes weakness to attackers or intruders who may observe the pattern of patrol and gain access during robot downtime [3]. A more realistic algorithmic performance benchmark, beyond idleness minimization, needs to take such predictability into account [34]. Other recent approaches have taken advantage of machine learning capabilities in order to derive a 'learned' patrol strategy to offer solutions to adversarial patrolling as traditional methods are subject to exploitation by attackers [4, 40]. These reinforcement methods offer benefits when it comes to managing potential attackers, but often do not generalize well to different environments [18]. An in-depth study of the latest approaches to the problem of multi-robot patrolling and their respective benefits and shortcomings can be found in [2].

A broad selection of ten existing multi-robot patrolling algorithms from the literature, designed to minimize idleness, are used here without modification. These algorithms are structured in different ways and as such, patrolling robots exhibit a variety of behavior (Table 1). More detail on the structure and operation of the algorithms is presented in [20, 24, 25]. As we go on to discuss, the algorithms vary considerably in the extent to which robots become spatially mixed, which is significant when they have local communication.

2.2 Collective Decision-Making and Perception

One approach to managing the large quantity of information that is gathered in a multi-robot system is for each robot to contribute its 'opinions' or 'beliefs' about the state of the world, and according to some global or local mechanism, allow the group to determine which to collectively adopt, referred to as reaching a *consensus*. Often owing to communication constraints, and also to prevent rapid noise propagation, consensus formation can be performed at a local level, between individual robots. How the consensus between



Figure 1: Trajectory plot of a system of eight robots executing the SEBS patrol algorithm for over an hour on the 'Cumberland' map, with each unique color representing a single robot and its patrol route.

all robots is formed depends on the belief exchange operator that is used to respond to information. One approach is to use voting model methods, where input from each robot is submitted into a voting system in order to determine the opinion of the group [12, 29], but this approach can often suffer from noisy measurements from individual robots [8]. An approach that seeks to solve the issue caused by noisy measurements was initially proposed by Perron et al. [21], whereby a third truth state of *'uncertain'* is employed as a compromise between certain of false or certain of negative. This was applied directly to multi-agent systems in [6], which showed that the application of this Kleene uncertain ternary logic can result in better performance than a standard Boolean truth pair in noisy environments.

In the collective learning literature, a common task is for agents to perceive a feature of the environment and come to a consensus on the state of that feature. This has been abstracted as a perception of white to black floor tiles (e.g. [29]), with the ratio being the element being voted on. This scenario shares similarities with the patrolling problem in regard to communication and data exchange [33], but the information is of a single dimension or value – the ratio of black to white. In the patrolling problem there is a single value for each area of interest, with all nodes having the same importance as another.

With the collective learning and communication of information, comes the ability for erroneous information to propagate throughout the system and cause a false consensus to be reached. Often in the literature an assumption is made that the world is static and fixed – that there is no change in the real values of the measured world. In a dynamic environment – such as a realistic patrolling problem – there are changes in the world that need to be recorded, such as the activity of an attacker, or variability in the physical environment affecting sensor readings ('noise'). Some recent work has shown that robots adapt better in environments that are subject to

Short name	Full name	Patrol Strategy	Decentralized?
CBSL[27]	Concurrent Bayesian Learning Strategy	Bayesian Learning	Yes
CGG[23]	Cyclic Algorithm for Generic Graphs	Hamiltonian Path	No
CR[19]	Conscientious Reactive	Reactive	Yes
DTAG[9]	Dynamic Task Assignment Greedy	Utility Function	Yes
DTAP[9]	Dynamic Task Assignment Auction	Utility Auctioneer	Yes
GBS[26]	Greedy Bayesian Strategy	Bayesian	No
HCR[1]	Heuristic Conscientious Reactive	Heuristic Reactive	Yes
HPCC[1]	Heuristic Pathfinder Conscientious Cognitive	Heuristic Pathfinder	Yes
RAND[20]	Random	Random Selection	Yes
SEBS[26]	State Exchange Bayesian Strategy	State & Bayes Exchange	Yes

Table 1: MRP algorithms examined in this study

change when they have a constrained communication, preventing preemptive consensus formation [31].

2.3 Communication Graph Connectivity

Robots in a multi-robot system can be dispersed over a wide area, and the ability for robots to maintain communication with one another is dependent staying within local communication range. If a robot is isolated or weakly connected to the communication graph, it will be unable to exchange information that it gathers or receive new information from other robots. It is therefore useful to use a measure of connectivity that adequately describes the interconnections between the robots and the strength of said connections. One of the most extensively used metrics to measure this in the literature is the *algebraic connectivity* or ‘Fiedler value’ [10]. Mathematically, it is the second smallest eigenvalue of the Laplacian matrix for the graph, where the Laplacian matrix is the degree matrix subtracted from the adjacency matrix. Conceptually, it is a measure of connections between nodes as well as the strength of the edges between nodes within a graph. The algebraic connectivity has a wide range of applications in determining the strength of connection of nodes within a system, most notably in wireless communications where adequate coverage and interconnections are highly desirable properties [17, 35, 38]. The measure of algebraic connectivity can be highly correlated with performance under certain circumstances such that it can be used as an input when designing the control algorithms for mobile robots. Zavlanos and Pappas used the algebraic connectivity of a team of mobile robots’ communications as a way of controlling their spatial distribution [37]. Other work has been done to use algebraic connectivity as part of the feedback input to a system within applications such as formation control [22]. Here, we use it to assess the emergent communication networks resultant from robots moving according to various MRP algorithms.




Figure 2: Example social network of eight robots executing SEBS algorithm on ‘Cumberland’ map for 1 hour. Nodes are individual robots, with edges representing the number of communications between respective robots. Darker shade indicates more frequent communication between robots.

3 EXPERIMENTAL SET UP

We examine the performance of a multi-robot system (MRS) undertaking multi-robot patrol (MRP) when there is one, and only one, true anomaly present. The anomaly is always located in the same place on the patrol graph (‘node 30’). When there is detection noise, false positive anomaly detections occur, and we examine the ability of the collective opinion dynamics to suppress these.

3.1 Simulator

The simulation environment used in this study is ROS Patrolling Sim¹, which is described in [28]. The simulator is built using the Stage² simulation package [11] and integrated into the ROS (Robot Operating System) framework, which enables easy transfer to real robotic systems.

Each robot in the simulation is modeled as a differential drive robot, with a laser rangefinder and odometry from wheel encoders with a drift error model. As the modality of the actual sensor is not the focus of this paper, anomalies to be detected are modeled as measurements that are made with a generic sensor. It is assumed that each observation of an anomaly is instantaneous upon reaching a patrol graph node, and subject to some noise probability value according to the experiment. Upon arrival to a node on the graph, the robot searches for an anomaly and records a true or false depending on world state and noise model for that node. Given a probability value of 5%, there is a 5% chance that the sensor records the opposite value that is present in the world. Upon revisiting that same node, there is a new chance of making an incorrect measurement. Robots are modeled as having an omni-directional antenna that is capable of communicating in a defined radius with other robots in the vicinity, regardless of obstacles or walls that may be between each robot. Communication is modeled as a one-to-one (pairwise) exchange, with a minimum timeout/delay period before repeated communications between the same robots. Upon initializing communication with another robot, a pairwise comparison of belief states occurs resulting in both robots having identical beliefs.

3.2 Opinion Dynamics

In a world with n nodes on the patrol graph, robots have a set of belief values that can be either 1, 0 or $\frac{1}{2}$, meaning true, false and uncertain respectively. Each robot maintains a set of measured belief states of the world according to the number of nodes that exist in the map $M = \{m_0, \dots, m_n\}$, where m_i is the robot's value for node i . When two robots compare beliefs, they update their belief values according to the fusion truth table shown in Table 2.

		$B(p_i)$		
		0	$\frac{1}{2}$	1
$B'(p_i)$	0	0	0	$\frac{1}{2}$
	$\frac{1}{2}$	0	$\frac{1}{2}$	1
	1	$\frac{1}{2}$	1	1

Table 2: Belief fusion truth table for beliefs B and B'

Using this belief fusion provides benefits to the system as a whole, as each robot can gain information on regions of the world that they have not explored. As the simulation world is subject to anomaly detection noise at individual nodes, the ability to reject noise at the system level is highly desirable. When two robots that begin the communication process have conflicting belief states on the same

node, their resultant value is that they are both uncertain about the value of that node. This ability to reconcile conflicting information also extends to measuring a node directly. When an agent revisits a node and takes a measurement which does not match its prior belief, the belief fusion in Table 2 is performed between the prior and subsequent beliefs. By employing the use of pairwise communication, the system as a whole is able to collectively learn the state of the world faster than a single robot. The overall benefit of limiting the information breadth by using this ternary belief state model over, say, a voting model is that the three-value system is better suited to reject noisy measurements or information from faulty agents [8].

3.3 Error and F-Score

In order to judge the performance of the system at arriving at an accurate consensus of the world state, a metric which will be referred to as *system error* is used, where the error is defined as:

$$E = \frac{1}{m} \frac{1}{n} \sum_{j=0}^m \sum_{i=0}^n |B_j(p_i) - S^*(p_i)| \quad (1)$$

Where m is the number of nodes on the graph, n is the number of robots in the experiment, B is the belief for a node, S^* is the true world state and p is the agent's belief at a given node [7].

Another system-level performance metric used is an F-score, which is introduced here in Equation 2. This represents the harmonic mean of the precision and recall of the system, which gives a way to compare the recorded results with the real world in order to determine the effect of noise on the performance of each algorithm. Terms in Equation 2 are t_p , the count of true positives, t_n , true negatives, f_p , false positives, f_n , false negatives and u , the number of unmeasured nodes. These are given a half weighting as per the uncertain belief in Table 2, to penalize the performance metric if nodes remain unknown. The use of this system-level F-score provides a useful insight into the performance of the system with regard to precision of the data that is collected. This scoring will be given for each algorithm at the end of the experiment when an expected consensus is to be reached.

$$F = \frac{2(t_p + t_n)}{2(t_p + t_n) + f_p + f_n + \frac{u}{2}} \quad (2)$$

3.4 Experiment Parameters

For each experiment run, each robot is initialized at the same position regardless of patrol algorithm to ensure comparability. Each set of experiments are run on the same default map ('Cumberland') that has 40 patrol graph nodes (Figure 1). This map is representative of a realistic building environment. The simulations have a system of eight robots, a group size that is both realistic and reduces the number of interference or collision events that can occur at higher robot densities [25]. There were 20 simulation runs per algorithm. Table 3 shows some further simulation parameters.

¹https://wiki.ros.org/patrolling_sim

²<https://github.com/rtv/Stage>

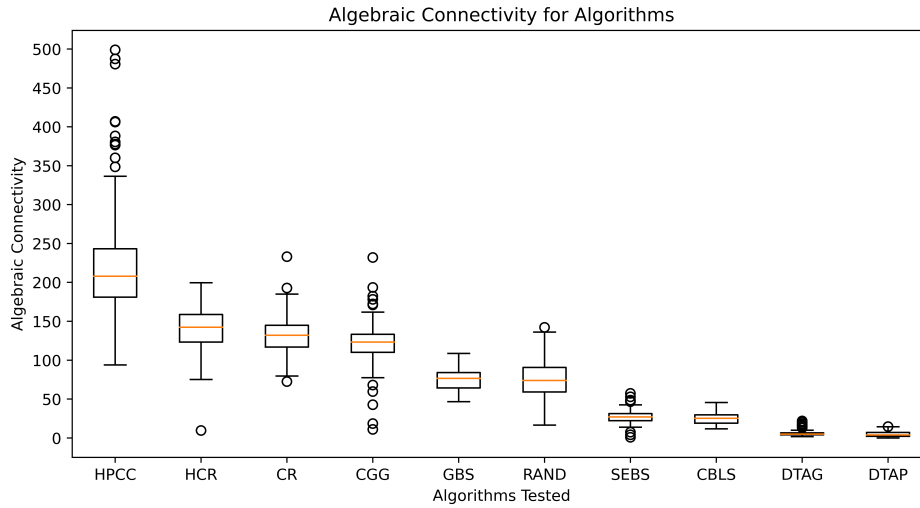


Figure 3: Algebraic Connectivity results for each algorithm tested




Figure 4: Time to reach 85% consensus on accurate view of world state against algebraic connectivity for each algorithm and noise level.

Experiment Parameter	Value
Robot communication range	5 m
Communication timeout period	30 s
Number of robots	8
Time of experiment	3600 s
Measurement noise probability	0, 5, 10 & 20%
Simulations per algorithm & noise level	20

Table 3: Simulated experiment parameter values

In order to compute the simulation runs, the software packages were ‘containerized’ and executed in parallel on a cloud-based compute service. The cloud service utilizes a 22-core Intel Xeon Gold 6238 with 16GB of RAM.

4 SIMULATION RESULTS AND DISCUSSION

A boxplot of the algebraic connectivity for each algorithm is presented in Figure 3. This shows a wide variety in communication intensity between robots, with the HPCC, HCR, CR and CGG algorithms leading to noticeably higher algebraic connectivity than the other six algorithms tested. At the other end of the scale, DTAP and DTAG have very low connectivity, while GBS, RAND, SEBS and CBLS have intermediate connectivity. Figure 3 seems to suggest, then, three approximate groupings of algorithms of high, intermediate and low connectivity.

There was a negative correlation between a system’s algebraic connectivity and the convergence time to 85% consensus on an accurate world view, i.e. at least 7 of 8 robots completely matching the true state of the world in their beliefs (Figure 4). There was a Pearson correlation coefficient of -0.377 ($p \ll 0.001$) for 0% noise, -0.313 ($p \ll 0.001$) for 5% noise, and -0.299 ($p \ll 0.001$) for 10% noise. On this figure 20% noise level is omitted as a majority of the simulations did not reach accurate consensus within the experiment

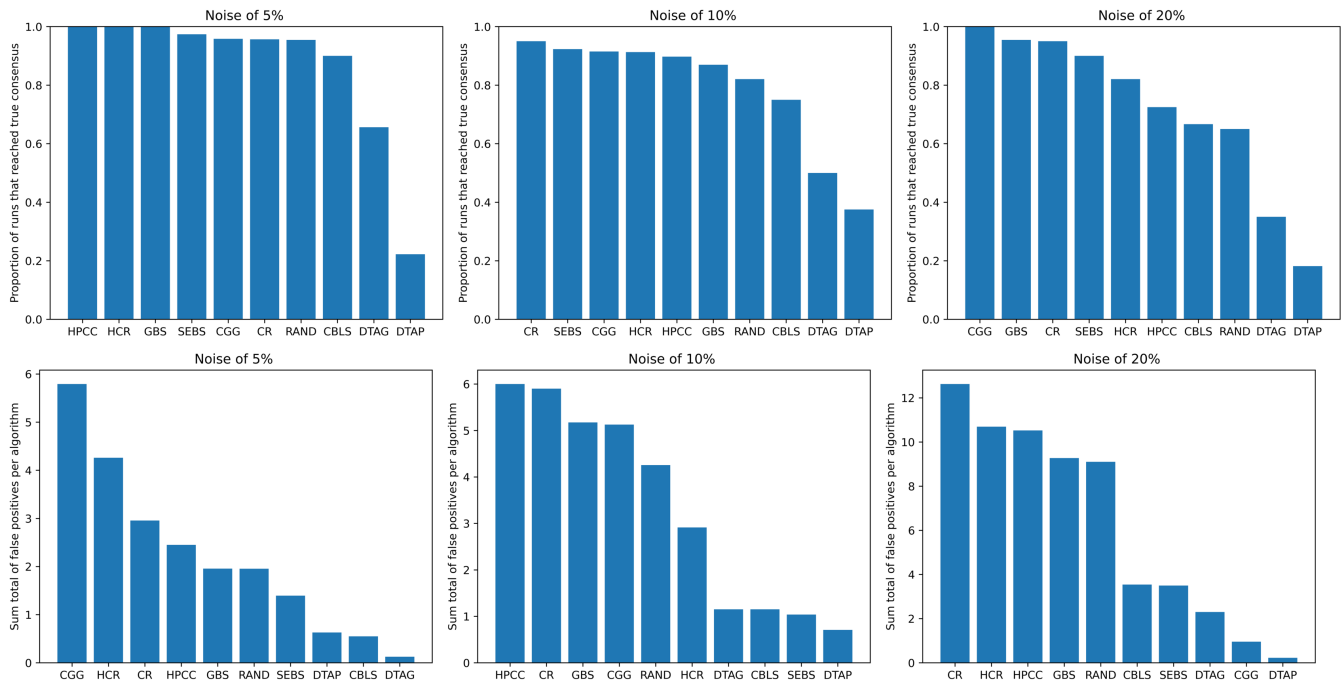


Figure 5: For 20 simulation runs, Top: Proportion of runs that reached consensus ($> 85\%$ or at least 7 of 8 robots) on particular node containing true positive (real security threat); Bottom: Average of count false positive consensus on other nodes.

run time. While faster convergence time is welcome if it is to a true positive, it may also have its disadvantages if it is at the expense of larger numbers of false positives.

Figure 5 shows that there are several algorithms that reliably give a consensus for a true positive (i.e. a correctly detected anomaly) on the correct node for different noise levels, producing a consensus 100% of the time, or at least on over 90% of occasions. This is valuable information that could help to prioritize a security professional’s inspections after time away from the building. However, in some cases there were also large numbers of false positives produced elsewhere on the patrol graph (Figure 5 bottom), averaging around 6 per run in the worst case (CGG) down much less than 1 per run with DTAG, with 5% noise. As detection noise increases, as anticipated, the number of false positives also increases, and thus it becomes apparent that it might be preferable to employ an algorithm that has good, but not the best, true positive detection performance, in return for cutting down false positives. SEBS and CBLs are notable examples of this, with much lower false positives than some alternatives that perform well in true positives, such as CGG. This trade-off is captured by the F-score.

Table 4 shows both the average idleness and F-scores for different noise levels, across the different algorithms tested. Standard deviations σ for the results are also shown. The best two algorithms are shown for each metric: very good performance of idleness minimization is found for CBLs and SEBS. For 0% noise, where false positives do not occur, HPCC and CR obtain the best F-scores, although SEBS is a very close third. As noise increases, RAND and CR lead at 5% noise, while at 10% HCR and SEBS lead, and then at 20% CGG and SEBS are the best performers on F-score. Overall, MRP

algorithms with intermediate algebraic connectivity, like SEBS and CBLs, appear to be the best all-rounders.

5 CONCLUSION AND FUTURE WORK

Multi-robot patrolling (MRP) is a well-studied problem with several effective control algorithms for a user to choose from. However, when there is anomaly detection noise, there is considerable untapped potential for a multi-robot system to help a user to prioritize their security inspections. This is because robots can share their detection information and form a collective consensus about which locations are most likely to be worth checking – and such locations could be the first stops on a regular human security patrol, for example. Here, we found that MRP algorithms controlling robots with local, pairwise belief exchange, resulted in systems that tended to converge to consensus more quickly when there was higher connectivity in their emergent communications network – that is, when they made robots physically pass by each other more often. However, for some algorithms, easy consensus formation also led to high numbers of false positive results. We developed a multi-robot F-score to obtain an overall view of system accuracy for each algorithm and noise level. We found that some of the leading MRP algorithms with respect to idleness minimization [27], which also had intermediate connectivity, also had good F-scores, with SEBS (State Exchange Bayesian Strategy, [26]) and CBLs (Concurrent Bayesian Learning Strategy, [27]) having good all-round performance. Their coordination of robots to enhance efficient patrolling also has the advantage of reducing the mixing of robots, and hence constrains somewhat their local communication.

Algorithm	Avg. Graph Idleness (s)	σ	F-Score							
			0% Noise		5% Noise		10% Noise		20% Noise	
			Avg.	σ	Avg.	σ	Avg.	σ	Avg.	σ
CBLS	59.2	7.7	0.962	0.0145	0.949	0.0124	0.916	0.0265	0.871	0.0210
CGG	100.4	17.7	0.954	0.0323	0.928	0.0462	0.931	0.0217	0.962	0.0257
CR	100.3	18.0	0.974	0.0008	0.952	0.0167	0.923	0.0181	0.860	0.0314
DTAG	68.8	6.2	0.914	0.0646	0.910	0.0828	0.883	0.0595	0.818	0.0633
DTAP	81.1	21.4	0.867	0.1291	0.839	0.1312	0.880	0.0764	0.852	0.1057
GBS	81.4	16.3	0.962	0.0000	0.951	0.0122	0.924	0.0251	0.875	0.0391
HCR	115.3	23.5	0.956	0.0200	0.933	0.0275	0.944	0.0335	0.872	0.0372
HPCC	169.2	59.7	0.974	0.0000	0.951	0.0165	0.922	0.0361	0.881	0.0339
RAND	99.3	20.9	0.954	0.0025	0.953	0.0129	0.918	0.0260	0.868	0.0314
SEBS	60.3	7.2	0.970	0.0142	0.935	0.0377	0.936	0.0242	0.902	0.0547

Table 4: Performance (average node idleness, F-score) for each algorithm tested at 0, 5, 10 and 20% noise in measurement, averaged across 20 simulation runs. Best two performing algorithms in each category are highlighted in bold text.

The algorithms and their emergent consensus formation behaviors have been studied here on a moderately-sized map (40 nodes, average degree 2.2). This investigation will be extended further for different maps. Given a patrol graph (map) with a different level of connectivity or larger number of nodes for the same number of robots to patrol, different interaction levels may result. This would require further investigation into the relationship between the number of robots to the number of nodes that each robot is expected to visit. Given a larger map for the same number of robots it could be anticipated that there will be fewer interactions due to the larger spatial separation between the robots (i.e. lower robot density). This may result in a slower consensus, or no consensus at all. Patrol graph connectivity (e.g. average degree) is also likely to be a relevant factor in emergent communication network connectivity.

We have examined the performance of unmodified MRP algorithms, where robot behavior does not change upon initial anomaly detection. A natural next step would be to add in a behavioral response such as a re-inspection or communication to a nearest neighbor to invite them to also visit that location. With such interaction, the spatial distribution of robots could itself indicate anomaly distribution [14, 15]. Given the good performance of algorithms resulting in *intermediate* network connectivity [31], this could also be used as a control input [22, 37], with global measurement or local connectivity estimation in a decentralized implementation [36, 38, 39]. We will also explore further scenarios in terms of anomaly regularity, given that real-world adversary activity is likely to be transient or intermittent. In such a case, it may well be preferable to err on the side of false positives, and prefer higher communication connectivity. Our results should help inform understanding of the trade-offs involved in using multi-robot systems to aid the perception of difficult-to-detect environmental features.

ACKNOWLEDGEMENTS

ZRM is supported by a University of Bristol PhD Scholarship. 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] Alessandro Almeida. 2003. *Patrulhamento Multiagente em Grafos com Pesos*. M.Sc. Thesis. Univ. Federal de Pernambuco, Brasil.
- [2] Nicola Basilico. 2022. Recent Trends in Robotic Patrolling. *Current Robotics Reports* 3, 2 (2022), 65–76. <https://doi.org/10.1007/s43154-022-00078-5>
- [3] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. 2012. Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence* 184-185 (2012), 78–123. <https://doi.org/10.1016/j.artint.2012.03.003>
- [4] Kenneth Bogert and Prashant Doshi. 2018. Multi-robot inverse reinforcement learning under occlusion with estimation of state transitions. *Artificial Intelligence* 263 (Oct. 2018), 46–73. <https://doi.org/10.1016/j.artint.2018.07.002>
- [5] Han-Lim Choi, Luc Brunet, and Jonathan P. How. 2009. Consensus-Based Decentralized Auctions for Robust Task Allocation. *IEEE Transactions on Robotics* 25, 4 (2009), 912–926. <https://doi.org/10.1109/TRO.2009.2022423> Conference Name: IEEE Transactions on Robotics.
- [6] Michael Crosscombe and Jonathan Lawry. 2017. Exploiting Vagueness for Multi-agent Consensus. In *Multi-agent and Complex Systems*, Quan Bai, Fenghui Ren, Katsuhide Fujita, Minjie Zhang, and Takayuki Ito (Eds.). Springer Singapore, Singapore, 67–78. https://doi.org/10.1007/978-981-10-2564-8_5
- [7] Michael Crosscombe and Jonathan Lawry. 2022. The Impact of Network Connectivity on Collective Learning. In *Distributed Autonomous Robotic Systems (Springer Proceedings in Advanced Robotics)*, Fumitoshi Matsuno, Shun-ichi Azuma, and Masahito Yamamoto (Eds.). Springer International Publishing, Cham, 82–94. https://doi.org/10.1007/978-3-030-92790-5_7
- [8] Michael Crosscombe, Jonathan Lawry, Sabine Hauert, and Martin Homer. 2017. Robust distributed decision-making in robot swarms: Exploiting a third truth state. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4326–4332. <https://doi.org/10.1109/IROS.2017.8206297>
- [9] Alessando Farinelli, Luca Iocchi, and Daniele Nardi. 2017. Distributed on-line dynamic task assignment for multi-robot patrolling. *Autonomous Robots* 41 (2017). <https://doi.org/10.1007/s10514-016-9579-8>
- [10] Miroslav Fiedler. 1973. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal* 23, 2 (1973), 298–305. <https://doi.org/10.21136/CMJ.1973.101168>

- [11] Brian P Gerkey, Richard T Vaughan, and Andrew Howard. 2003. The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems. In *Proceedings of the International Conference on Advanced Robotics (ICAR 2003)*.
- [12] Farzad Farnoud Hassanzadeh, Eitan Yaakobi, Behrouz Touri, Olga Milenkovic, and Jehoshua Bruck. 2013. Building consensus via iterative voting. In *2013 IEEE International Symposium on Information Theory*. 1082–1086. <https://doi.org/10.1109/ISIT.2013.6620393>
- [13] Li Huang, MengChu Zhou, Kuangrong Hao, and Edwin Hou. 2019. A survey of multi-robot regular and adversarial patrolling. *IEEE/CAA Journal of Automatica Sinica* 6, 4 (2019), 894–903. <https://doi.org/10.1109/JAS.2019.1911537>
- [14] Edmund R. Hunt. 2022. What Could Models of Superorganismal Cognition Offer to Embodied AI? *IOP Conference Series: Materials Science and Engineering* 1261, 1 (2022), 012023. <https://doi.org/10.1088/1757-899X/1261/1/012023>
- [15] Edmund R. Hunt, Nigel R. Franks, and Roland J. Baddeley. 2020. The Bayesian Superorganism: Collective Probability Estimation in Swarm Systems (*ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*). 315–323. https://doi.org/10.1162/isal_a_00247
- [16] Kao-Shing Hwang, Jin-Ling Lin, and Hui-Ling Huang. 2009. Cooperative patrol planning of multi-robot systems by a competitive auction system. In *2009 ICCAS-SCIE*. 4359–4363.
- [17] Huijia Lin, Maohua Lu, Nikola Milosavljevic, Jie Gao, and Leonidas J. Guibas. 2008. Composable Information Gradients in Wireless Sensor Networks. In *2008 International Conference on Information Processing in Sensor Networks (IPSN 2008)*. 121–132. <https://doi.org/10.1109/IPSN.2008.21>
- [18] Najun Liu, Yinghao Cai, Tao Lu, Rui Wang, and Shuo Wang. 2020. Real–Sim–Real Transfer for Real-World Robot Control Policy Learning with Deep Reinforcement Learning. *Applied Sciences* 10, 5 (2020), 1555. <https://doi.org/10.3390/app10051555>
- [19] Aydano Machado. 2002. *Patrulha Multiagente: Uma Análise Empírica e Sistemática*. M.Sc. Thesis. Univ. Federal de Pernambuco, Brasil.
- [20] Aydano Machado, Geber Ramalho, Jean-Daniel Zucker, and Alexis Drogoul. 2003. Multi-agent Patrolling: An Empirical Analysis of Alternative Architectures. In *Multi-Agent-Based Simulation II (Lecture Notes in Computer Science)*, Jaime Simão Sichman, François Bousquet, and Paul Davidsson (Eds.). Springer, Berlin, Heidelberg, 155–170. https://doi.org/10.1007/3-540-36483-8_11
- [21] E. Perron, D. Vasudevan, and M. Vojnovic. 2009. Using Three States for Binary Consensus on Complete Graphs. In *IEEE INFOCOM 2009*. IEEE, Rio de Janeiro, Brazil, 2527–2535. <https://doi.org/10.1109/INFOCOM.2009.5062181>
- [22] Hasan A. Poonawala, Aykut C. Satici, Hazen Eckert, and Mark W. Spong. 2015. Collision-Free Formation Control with Decentralized Connectivity Preservation for Nonholonomic-Wheeled Mobile Robots. *IEEE Transactions on Control of Network Systems* 2, 2 (2015), 122–130. <https://doi.org/10.1109/TNS.2014.2378876>
- [23] David Portugal and Rui Rocha. 2010. MSP algorithm: multi-robot patrolling based on territory allocation using balanced graph partitioning. In *Proceedings of the 2010 ACM Symposium on Applied Computing*. ACM, Sierre Switzerland, 1271–1276. <https://doi.org/10.1145/1774088.1774360>
- [24] David Portugal and Rui Rocha. 2011. A Survey on Multi-robot Patrolling Algorithms. In *Technological Innovation for Sustainability*, Luis M. Camarinha-Matos (Ed.). Vol. 349. Springer Berlin Heidelberg, Berlin, Heidelberg, 139–146. https://doi.org/10.1007/978-3-642-19170-1_15
- [25] David Portugal and Rui P. Rocha. 2012. Decision methods for distributed multi-robot patrol. In *2012 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 1–6. <https://doi.org/10.1109/SSRR.2012.6523869>
- [26] David Portugal and Rui P. Rocha. 2013. Distributed multi-robot patrol: A scalable and fault-tolerant framework. *Robotics and Autonomous Systems* 61, 12 (2013), 1572–1587. <https://doi.org/10.1016/j.robot.2013.06.011>
- [27] David Portugal and Rui P. Rocha. 2016. Cooperative multi-robot patrol with Bayesian learning. *Autonomous Robots* 40, 5 (2016), 929–953. <https://doi.org/10.1007/s10514-015-9503-7>
- [28] David Bina Siassipour Portugal. 2013. *Effective Cooperation and Scalability in Multi-Robot Teams for Automatic Patrolling of Infrastructures*. Ph. D. Dissertation. Universidade de Coimbra, Portugal. https://ap.isr.uc.pt/archive/DPortugal_PhDThesis_2014.pdf
- [29] Qihao Shan, Alexander Heck, and Sanaz Mostaghim. 2021. Discrete Collective Estimation in Swarm Robotics with Ranked Voting Systems. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*. 1–8. <https://doi.org/10.1109/SSCI50451.2021.9659868>
- [30] Efrat Sless Lin, Noa Agmon, and Sarit Kraus. 2019. Multi-robot adversarial patrolling: Handling sequential attacks. *Artificial Intelligence* 274 (2019), 1–25. <https://doi.org/10.1016/j.artint.2019.02.004>
- [31] Mohamed S. Talamali, Arindam Saha, James A. R. Marshall, and Andreagiovanni Reina. 2021. When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics* 6, 56 (2021), eabf1416. <https://doi.org/10.1126/scirobotics.abf1416>
- [32] Kristal K. Trejo, Julio B. Clempner, and Alexander S. Poznyak. 2015. A Stackelberg security game with random strategies based on the extraproximal theoretic approach. *Engineering Applications of Artificial Intelligence* 37 (2015), 145–153. <https://doi.org/10.1016/j.engappai.2014.09.002>
- [33] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. 2017. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics & AI* 4 (2017). <https://doi.org/10.3389/frobt.2017.00009>
- [34] James C. Ward and Edmund R. Hunt. 2023. An Empirical Method for Benchmarking Multi-Robot Patrol Strategies in Adversarial Environments. In *Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing (Tallinn, Estonia) (SAC '23)*. Association for Computing Machinery, New York, NY, USA, 787–790. <https://doi.org/10.1145/3555776.3577802>
- [35] Feng Xue and P.R. Kumar. 2004. The Number of Neighbors Needed for Connectivity of Wireless Networks. *Wireless Networks* 10, 2 (2004), 169–181. <https://doi.org/10.1023/B:WINE.0000013081.09837.c0>
- [36] Peng Yang, Randy A. Freeman, Geoffrey J. Gordon, Kevin M. Lynch, Siddhartha S. Srinivasa, and Rahul Sukthankar. 2008. Decentralized estimation and control of graph connectivity in mobile sensor networks. In *2008 American Control Conference*. IEEE, Seattle, WA, 2678–2683. <https://doi.org/10.1109/ACC.2008.4586897>
- [37] Michael M. Zavlanos, Magnus B. Egerstedt, and George J. Pappas. 2011. Graph-theoretic connectivity control of mobile robot networks. *Proc. IEEE* 99, 9 (2011), 1525–1540. <https://doi.org/10.1109/JPROC.2011.2157884>
- [38] Michael M. Zavlanos and G.J. Pappas. 2008. Distributed connectivity control of mobile networks. *IEEE Transactions on Robotics* 24, 6 (2008), 1416–1428. <https://doi.org/10.1109/TRO.2008.2006233>
- [39] Michael M. Zavlanos, Herbert G. Tanner, Ali Jadbabaie, and George J. Pappas. 2009. Hybrid Control for Connectivity Preserving Flocking. *IEEE Trans. Automat. Control* 54, 12 (2009), 2869–2875. <https://doi.org/10.1109/TAC.2009.2033750>
- [40] Xin Zhou, Weiping Wang, Tao Wang, Yonglin Lei, and Fangcheng Zhong. 2019. Bayesian Reinforcement Learning for Multi-Robot Decentralized Patrolling in Uncertain Environments. *IEEE Transactions on Vehicular Technology* 68, 12 (2019), 11691–11703. <https://doi.org/10.1109/TVT.2019.2948953>