Decentralised self-organising maps for the online orienteering problem with neighbourhoods

Graeme Best[1] and Geoffrey A. Hollinger[1]

Abstract— **This extended abstract presents a new coordination algorithm for decentralised multi-robot information gathering. We consider planning for an online variant of the multi-agent orienteering problem with neighbourhoods. This formulation closely aligns with a number of important tasks in robotics, including inspection, surveillance, and reconnaissance. We propose a decentralised variant of the self-organising map (SOM) learning procedure, named Dec-SOM, which efficiently finds plans for a team of robots in a non-myopic manner. Decentralisation is achieved by performing a distributed allocation scheme jointly with the SOM adaptations. Preliminary simulation results indicate that Dec-SOM outperforms baseline methods and is a viable solution for decentralised, online, and non-myopic information gathering.**

I. Introduction

Mobile robots are increasingly being used to gather information about their environment, such as for infrastructure inspection [1], mine countermeasures [2], and precision agriculture [3]. Many information gathering scenarios, especially those that can be formulated as mapping, coverage, or search problems, involve observing a set of points of interest (POIs) from associated observation regions [1]–[11]. The problem is to plan action sequences for robots that maximise the number of observed POIs while satisfying time or energy budgets.

Performing information gathering with multiple robots enables scaling up the number of observations in time and space. However, to achieve desirable performance, the robots are required to effectively coordinate their actions. The coordination is preferably decentralised, especially when communication is challenging, such as in marine environments [12] and underground tunnels [13]. Decentralised coordination is difficult because the robots must plan valuable action sequences while only having partial knowledge of the plans of other robots. Ideally, this planning is performed efficiently, online, and non-myopically.

Self-organising maps (SOMs) are a special class of learning procedures that aim to find a topology-preserving dimensionality reduction of an an input space [14]. In path planning contexts, SOMs are particularly powerful at solving problems that require *jointly* optimising the selection of viewpoints and the path through these viewpoints, such as in the TSP with neighbourhoods and related variants [5], [6], [15], [16]. Multi-robot SOM variants have been proposed [5], [6], [17]; however, all of these methods are centralised. Here, we are interested in developing *decentralised* SOM algorithms for multi-robot planning.

We propose the decentralised self-organising map algorithm (Dec-SOM) for multi-robot information gathering. The task is formulated as a generalisation of the orienteering problem with neighbourhoods, where a set of continuous regions are to be discovered online and maximally visited. Dec-SOM consists of each robot optimising their path using SOM adaptations and negotiating with other robots for the allocation of goal regions. Dec-SOM is shown to (1) outperform a number of baseline comparison methods and alternative SOM approaches, (2) plan effectively in partially-known worlds, (3) efficiently adapt to changing information, and (4) experience a gradual degradation of performance as communication becomes less reliable.

II. Related Work

Informative path planning is the problem of finding paths that maximise an information gain metric, subject to budget constraints [18]. Commonly, the considered information metric is an uncertainty measure for the belief of a quantity of interest. However, such metrics are often computationally demanding or not applicable for many tasks. Another approach, which we consider in this paper, is to formulate objectives as a set of POIs to be observed. These objectives naturally align with tasks such as area coverage [1], [4], classifying physical objects [3], [6]–[8], or observing scientifically valuable regions of oceans [9]. These objectives are typically faster to compute, enabling the efficient use of non-myopic planners.

Despite the benefits of multi-robot systems, relatively little attention has been given to developing decentralised planners for information gathering. Promising recent work include the generally-applicable Dec-MCTS algorithm [11], and sequential greedy assignment for exploration [19]. Our method also shares similarities to market-based approaches [20]–[22], but our method has the benefit of optimising allocations *jointly* with path planning.

SOMs have recently emerged as a powerful method for path planning that involves range sensing, such as TSP generalisations that require selecting favourable viewpoints within continuous goal regions [15], [16]. These have been applied to problems such as sensor network data collection [5], perception of 3D objects [6], and area surveillance [10]. All of these SOM variants are for single-robot or *centralised* multi-robot systems. Little attention has been given to distributed computation, despite SOMs being parallelisable [23].

*This work was supported in part by Office of Naval Research Grant N00014-17-1-2581.

[1]The authors are with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute, Oregon State University, Corvallis OR, USA. {bestg,geoff.hollinger}@oregonstate.edu

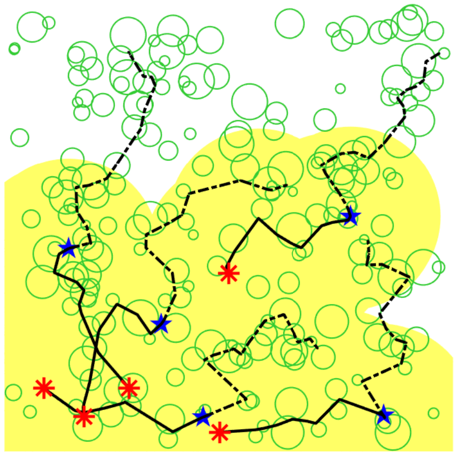

Fig. 1: Representative scenario and plans. 5 robots have naturally split up to efficiently explore different parts of the world, while focussing on regions with high goal densities. Solid lines are executed paths; dotted are plans. Explored regions are yellow; predicted regions are unshaded.

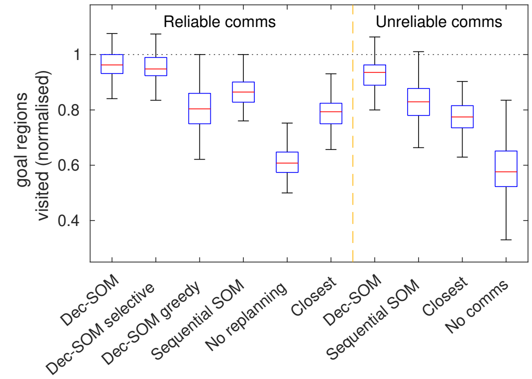

Fig. 2: Performance of Dec-SOM and various comparison methods in 100 random environments. Scores are normalised by dividing by the performance of a full-information *oracle* planner in each environment.

III. Problem Formulation

We consider the problem of planning the actions for a team of R robots, where the path of each robot r is described as a sequence of waypoint locations $\boldsymbol{x}^r = (x_0^r, x_1^r, x_2^r, ...)$. A path is feasible if the path length satisfies a given budget. The world consists of a set of circular goal regions $\mathcal{Z}$, which are to be visited by the robots. The set of goal regions $\mathcal{Z}$ is not known to the robots in advance, but they instead have an estimate $\tilde{\mathcal{Z}} \approx \mathcal{Z}$ that is refined online. We aim to address the problem of finding the set of feasible paths $\boldsymbol{x} := \{\boldsymbol{x}^1, \boldsymbol{x}^2, ..., \boldsymbol{x}^R\}$ for the team of robots that collectively maximises the number of goal regions $z_j \in \mathcal{Z}$ that are visited at least once. To facilitate decentralisation, we assume that each robot r can only modify its own plan. The robots can communicate to aid effective coordination; however, communication may be unreliable. The robots should adapt their plans in an online manner as the estimate $\tilde{\mathcal{Z}}$ is refined.

IV. Overview of Dec-SOM

We propose Dec-SOM as a solution to the above problem. The algorithm generalises the SOM learning procedure to be suitable for this decentralised setting by employing a novel distributed allocation scheme within the SOM adaptations.

An SOM provides a lower-dimensional representation of an input space, where the representation preserves a given topological structure. In our case, the input space is the goal regions that can be visited by robot r. The SOM aims to find a path for robot r that 'best fits' this input space. The algorithm jointly learns both: (1) the allocation of goal regions to individual robots and (2) the path that maximally visits the allocated goal regions.

The main loop of Dec-SOM cycles between: (1) select a goal region at random, (2) adapt the plan for robot r towards this goal region using an SOM adaptation procedure, (3) determine the value of this adaptation and, if appropriate, request the allocation of this goal from other robots, (4) retain or discard this adaptation, and (5) periodically regenerate the plan. In parallel, robots process and reply to incoming allocation requests.

One of the key components of Dec-SOM is the SOM adaptation step, which is essential for generating desirable paths. When a goal region z_j is presented, the closest waypoint or edge to any point on the disk z_j is selected as the *winner* and moved to the closest point in z_j. All other waypoints in $\boldsymbol{x}^r$ are moved towards z_j by a fraction determined by the topological distance to the winner.

The other key component is the negotiation between robots. If robot r desires to retain an adaptation towards z_j, but has not been allocated z_j then it requests this allocation from other robots. It does so by broadcasting z_j and a score s. We define s as the increase in path length as a result of the allocation, divided by the fraction of budget used. Another robot will reply if it holds the allocation. The winning robot retains the allocation.

V. Experiments

Preliminary experimental results compare the performance of Dec-SOM to various baseline comparison methods. These experiments involve planning for a team of 5 robots in random worlds of 200 goal regions, as illustrated in Fig. 1.

Results are shown in Fig. 2. We observe that Dec-SOM is competitive with an oracle (full information case), meaning it is able to effectively address the challenge of having a noisy world estimate. Sequential planning (plan for one robot at a time) performed relatively poorly, which demonstrates the benefits of iterative negotiations. Greedy planning performed remarkably poorly since it lacks foresight. The no replanning case performed the worst, which demonstrates the need to adapt plans. When communication was unreliable, Dec-SOM exhibited a small reduction in performance, however was still able to perform reasonably well despite this difficulty.

VI. Future Work

Future work includes performing more detailed experiments and exploring generalisations of Dec-SOM. Straightforward extensions would be to incorporate polygonal goals, non-uniform rewards, heterogeneous sensing, and 3D environments. It would be worthwhile to investigate approaches for reducing the communication requirements [24]. It would also be interesting to exploit probabilistic estimates for $\tilde{\mathcal{Z}}$ and address cases where robots have inconsistent beliefs.

References

[1] M. Hassan and D. Liu, "Simultaneous area partitioning and allocation for complete coverage by multiple autonomous industrial robots," *Autonomous Robots*, vol. 41, no. 8, pp. 1609–1628, 2017.

[2] S. Sariel, T. Balch, and N. Erdogan, "Naval mine countermeasure missions," *IEEE Robotics Automation Mag.*, vol. 15, no. 1, pp. 45–52, 2008.

[3] F. Sukkar, G. Best, C. Yoo, and R. Fitch, "Multi-robot region-of-interest reconstruction with Dec-MCTS," in *Proc. of IEEE ICRA*, 2019.

[4] C. Dornhege, A. Kleiner, A. Hertle, and A. Kolling, "Multirobot coverage search in three dimensions," *J. Field Robotics*, vol. 33, no. 4, pp. 537–558, 2016.

[5] J. Faigl and G. A. Hollinger, "Autonomous data collection using a self-organizing map," *IEEE Trans. Neural Networks and Learing Systems*, vol. 29, no. 5, pp. 1703–1715, 2018.

[6] G. Best, J. Faigl, and R. Fitch, "Online planning for multi-robot active perception with self-organising maps," *Autonomous Robots*, vol. 42, no. 4, pp. 715–738, 2018.

[7] P. Tokekar, J. V. Hook, D. Mulla, and V. Isler, "Sensor planning for a symbiotic UAV and UGV system for precision agriculture," *IEEE Trans. Robotics*, vol. 32, no. 6, pp. 1498–1511, 2016.

[8] P. A. Plonski and V. Isler, "Approximation algorithms for tours of height-varying view cones," *Int. J. Robotics Research*, vol. 38, no. 2-3, pp. 224–235, 2019.

[9] S. McCammon and G. A. Hollinger, "Topological hotspot identification for informative path planning with a marine robot," in *Proc. of IEEE ICRA*, 2018.

[10] J. Faigl, P. Váňa, R. Pěnička, and M. Saska, "Unsupervised learning-based flexible framework for surveillance planning with aerial vehicles," *J. Field Robotics*, vol. 36, no. 1, pp. 270–301, 2019.

[11] G. Best, O. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "Dec-MCTS: Decentralized planning for multi-robot active perception," *Int. J. Robotics Research*, vol. 38, no. 2-3, pp. 316–337, 2019.

[12] G. A. Hollinger, S. Yerramalli, S. Singh, U. Mitra, and G. Sukhatme, "Distributed coordination and data fusion for underwater search," in *Proc. of IEEE ICRA*, 2011, pp. 349–355.

[13] D. Tardioli, L. Riazuelo, D. Sicignano, C. Rizzo, F. Lera, J. L. Villarroel, and L. Montano, "Ground robotics in tunnels: Keys and lessons learned after 10 years of research and experiments," *J. Field Robotics*, 2019, doi:10.1002/rob.21871.

[14] T. Kohonen, "Essentials of the self-organizing map," *Neural Networks*, vol. 37, pp. 52–65, 2013.

[15] J. Faigl, "Approximate solution of the multiple watchman routes problem with restricted visibility range," *IEEE Trans. Neural Networks*, vol. 21, no. 10, pp. 1668–1679, 2010.

[16] J. Faigl, R. Pěnička, and G. Best, "Self-organizing map-based solution for the orienteering problem with neighborhoods," in *Proc. of IEEE SMC*, 2016, pp. 1315–1321.

[17] S. Somhom, A. Modares, and T. Enkawa, "Competition-based neural network for the multiple travelling salesmen problem with minmax objective," *Computers & Operations Research*, vol. 26, no. 4, pp. 395–407, 1999.

[18] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *Int. J. Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.

[19] M. Corah and N. Michael, "Distributed matroid-constrained submodular maximization for multi-robot exploration: Theory and practice," *Autonomous Robots*, vol. 43, no. 2, pp. 485–501, 2019.

[20] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer, "Multi-robot exploration controlled by a market economy," in *Proc. of IEEE ICRA*, vol. 3, 2002, pp. 3016–3023.

[21] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Trans. Robotics*, vol. 25, no. 4, pp. 912–926, 2009.

[22] L. Liu and D. A. Shell, "Large-scale multi-robot task allocation via dynamic partitioning and distribution," *Autonomous Robots*, vol. 33, no. 3, pp. 291–307, 2012.

[23] T. D. Hämäläinen, "Parallel implementations of self-organizing maps," in *Self-Organizing Neural Networks*. Springer, 2002, pp. 245–278.

[24] G. Best, M. Forrai, R. R. Mettu, and R. Fitch, "Planning-aware communication for decentralised multi-robot coordination," in *Proc. of IEEE ICRA*, 2018, pp. 1050–1057.