Detecting Anomalous Behaviour in Robot Swarms

Ingeborg Wenger, Henrik Ebel, Peter Eberhard

Institute of Engineering and Computational Mechanics
University of Stuttgart
Pfaffenwaldring 9, 70569 Stuttgart, Germany
[ingeborg.wenger, henrik.ebel, peter.eberhard]@itm.uni-stuttgart.de

EXTENDED ABSTRACT

1 Introduction

Robotic swarms have been proposed for a variety of tasks including rescue missions, monitoring the spreading of oil leaks or contaminant clouds, mine detection, or surveillance [1]. While a lot of previous work has been done on cooperative task solving, few approaches exist for identifying and possibly correcting anomalous behavior of swarm elements such as accidental faults, computer virus infections, system failures, or malicious swarm elements [2, 3]. For instance, in [2] an intrusion detection system is proposed based on checking against pre-defined signatures of non-characteristic behavior. In [4] inverse reinforcement learning is used to detect anomalies in the movement of linear autonomous systems controlled by a linear quadratic regulator. Our work focuses on a more general approach to identifying robot behavior that is inconsistent with a known task objective. We assume that the decision-making of both benevolent and anomalous robots is determined by minimizing their respective objective function. Faulty or malicious robots either perform poorly while following the correct objective function or behave according to a different objective function than the rest of the robotic swarm. Since the successful execution of the above mentioned tasks, such as surveillance, depends on the robot swarm optimizing their coverage of an area, we identified a coverage task as a suitable model problem for this contribution.

The optimal coverage of a polygonal area W is found by minimizing the cost function

$$H(\mathbf{x}_1,\dots,\mathbf{x}_n,W)=\sum_{i=1}^{n}\int_{W_i}||\mathbf{q}-\mathbf{x}_i||^2\phi(\mathbf{q})\,\mathrm{d}W_i \tag{1}$$

with $\mathbf{x}_i\in\mathbb{R}^2$ being the position of robot i and W_i representing the Voronoi cell corresponding to $\mathbf{x}_i$ in the Voronoi decomposition of W generated by $\{\mathbf{x}_1,\dots,\mathbf{x}_n\}$. The weighting function $\phi:\mathbb{R}^2\to\mathbb{R}$ determines the importance of different locations within the area and is set to a uniform distribution by default. We consider an anomalous swarm element R_a that has been modified to assume sole coverage of a specific area W_a within W. Two different types of scenarios are inspected. In the first scenario, R_a minimizes the distance to a goal point within W_a instead of minimizing the objective function from Equation (1). For the second scenario, the objective function of R_a is modified by setting ϕ to a Gaussian distribution centered to the anomalous agent's area of interest. An example of the resulting anomalous behavior for the second scenario is shown in Figure 1.

2 Methodology

A naive way to verify the behavior of a robot R_i would involve having a subset of its neighbours minimize its objective function and checking the result against the actual behaviour of R_i. However, this is not only computationally demanding, but also error-prone due to robots being influenced by environmental noise and local information such as from obstacles. Instead, we propose a probabilistic approach in order to approximate the desired behavior of robot R_i over multiple timesteps. Identifying the objective

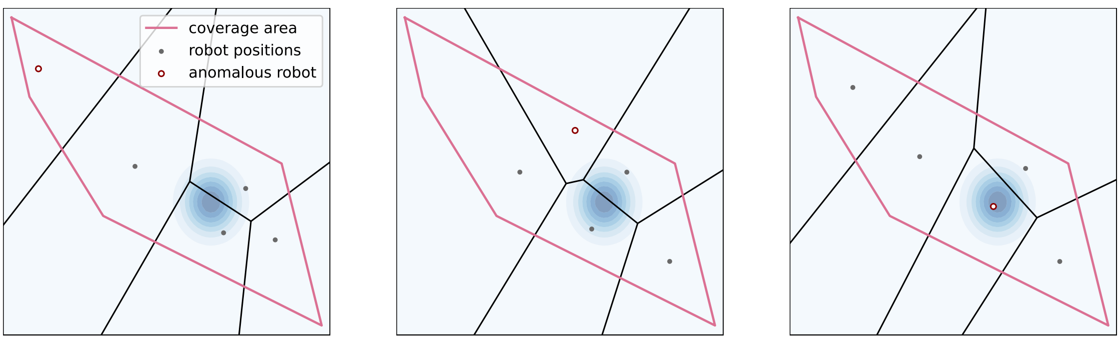

Figure 1: The anomalous robot moves towards the weighted area instead of covering the left corner of the coverage area. The blue distribution indicates the weighting function ϕ that modifies the behavior of the anomalous robot. The black lines show the Voronoi decomposition.

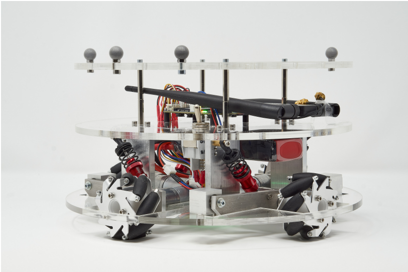

Figure 2: Non-holonomic robot used for experimental validation.

function of a robot based on its behavior allows to detect anomalous behavior while possibly enabling the swarm to prevent the anomalous element from completing its goal. To that end, we investigate an inverse reinforcement learning approach based on a nonlinear reward function r_i determining the behavior of robot i. Similar to the approach proposed in [5], we consider a Gaussian process regression for r_i. The reward function is found by maximizing the log-likelihood of the robot behavior under r_i,

$$\log P(\mathbf{S}_i|r_i) = \sum_t \log P(\mathbf{s}_{i,t+1}|\mathbf{s}_{i,t}, r_i) \tag{2}$$

with $\mathbf{S}_i \in \mathbb{R}^{T\times k}$ corresponding to a trajectory of length T of states $\mathbf{s}_{i,t} \in \mathbb{R}^k$ that has been observed for robot i. Special attention is paid to the feature selection of the states such that the learned reward function has a meaningful structure and is able to generalize to unknown states. The swarm behavior is simulated by modeling a swarm of non-holonomic robots, see Figure 2, with the equation of motion

$$\mathbf{M}\ddot{\mathbf{y}} + \mathbf{D}\dot{\mathbf{y}} = \mathbf{u}, \tag{3}$$

where $\mathbf{M} \in \mathbb{R}^{2\times 2}$ and $\mathbf{D} \in \mathbb{R}^{2\times 2}$ describe the mass and damping matrices, $\mathbf{y} \in \mathbb{R}^2$ corresponds to the robot positions and $\mathbf{u} \in \mathbb{R}^2$ determines the control input. We incorporate prior information about the robot model into the reward function r, for instance via limiting the maximum distance per timestep that a robot will be allowed to cover. Additionally, we take information about the cost function $H(\mathbf{x}_1,\ldots,\mathbf{x}_n,W)$ into account, with benevolently behaving robots R_b approximately minimizing the cost function. Our approach is developed using simulations and will be verified in experiments on a swarm of non-holonomic robots as the one shown in Figure 2.

Our contribution shows a first approach for anomaly detection in the objective function of swarm robots whose behavior might not be conforming to the task objective. Future work will investigate different methods for learning a reward function. Furthermore, we aim at finding suitable methods for counteracting anomalous behavior that leverage the learned objective function of an anomalous robot.

Acknowledgments

This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grant 433183605 (Research on Multibody Dynamics and Control for Collaborative Elastic Object Transportation by a Heterogeneous Swarm with Aerial and Land-Based Mobile Robots) and through Germany's Excellence Strategy (Project PN4-4 Theoretical Guarantees for Predictive Control in Adaptive Multi-Agent Scenarios) under Grant EXC 2075-390740016 (SimTech).

References

[1] S. Subchan, B. A. White, A. Tsourdos, M. Shanmugavel and R. Zbikowski. Pythagorean Hodograph (PH) Path Planning for Tracking Airborne Contaminant using Sensor Swarm. Proceedings of the IEEE Instrumentation and Measurement Technology Conference, pages 501-506, 2008.

[2] I. Sargeant and A. Tomlinson. Intrusion Detection in Robotic Swarms. In Y. Bi, R. Bhatia, S. Kapoor, editors, Intelligent Systems and Applications. IntelliSys 2019. Advances in Intelligent Systems and Computing, 1038:968-980, Springer, Cham, 2019.

[3] Y.-S. Dai, M. Hinchey, M. Madhusoodan, J. L. Rash and X. Zou. A Prototype Model for Self-Healing and Self- Reproduction In Swarm Robotics System. Proceedings of the 2nd IEEE International Symposium on Dependable, Autonomic and Secure Computing, pages 3-10, 2006.

[4] B. Lian, Y. Kartal, F.L. Lewis, D.G. Mikulski, G.R. Hudas, Y. Wan and A. Davoudi. Anomaly Detection and Correction of Optimizing Autonomous Systems with Inverse Reinforcement Learning. IEEE Transactions on Cybernetics, 2022. doi: 10.1109/TCYB.2022.3213526.

[5] S. Levine, Z. Popovic and V. Koltun. Nonlinear Inverse Reinforcement Learning with Gaussian Processes. Proceedings of the 24th International Conference on Neural Information Processing Systems. Advances in Neural Information Processing Systems, pages 19-27, Curran Associates Inc., 2011.