Identifying Anomalous and Antagonistic Behavior in Networks of Multibody Systems

Henrik Ebel, Ingeborg Wenger and Peter Eberhard

May 17, 2024

Identifying Anomalous and Antagonistic Behavior in Networks of Multibody Systems

Henrik Ebel*, **Ingeborg Wenger#**, **Peter Eberhard#**

* Department of Mechanical Engineering, LUT University
P.O. Box 20, 53851 Lappeenranta, Finland
henrik.ebel@lut.fi

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

Abstract

Networks of systems are of increasing relevance, in particular in robotics, where a cooperating and communicating network of robots can achieve things unachievable by a single robot. In such applications, the robots themselves as well as the overall system consisting of the robots, their task space and the objects therein can be modeled as multibody systems. Employing a network of systems with underlying distributed decision-making algorithms is in some ways more robust than centralized or single-system approaches since the network can potentially reorganize to compensate for breakdowns of individual systems. However, networked, interconnected behavior is usually rather intricate to understand, design, and analyze, which also makes it vulnerable. In particular, in uncontrolled environments, it may happen that malevolent agents enter the network, derogating the performance of the whole network. Due to interdependencies and dynamic couplings, it may not be immediately clear to spectators which system is the culprit. This work proposes and discusses two distinct approaches to automatically detect anomalous behavior in networks of multibody systems, one of which is based on inverse optimal-control, whereas the other is based on machine learning. Anomalous behavior in this work reaches from merely passively erroneous behavior, e.g., as it may appear due to defects, to actively antagonistic behavior that may strike a balance between causing maximum damage and remaining undetected.

This work considers two prototypical problems from networked robotics that serve as building blocks for many practical applications. Firstly, formation control is considered, where the systems need to coordinate their states relatively to one another, e.g., a group of mobile robots attaining a certain formation, see the left-hand side of Figure 1 for an illustrative simulation example. Therein, four robots use a formation controller to track an ∞-shaped path with the geometric center of the formation while keeping a square formation that is rotated at a constant rate around the formation center. The trajectories of the robots are plotted in the same colors as the circles representing the robots. Formation control can, e.g., be useful for cooperative manipulation tasks [1]. Secondly, coverage problems are considered, where the robots need to optimally cover an area of interest, which can be useful for surveillance and monitoring purposes, e.g., in emergency scenarios. A simulated coverage scenario is illustrated on the right-hand side of Figure 1.

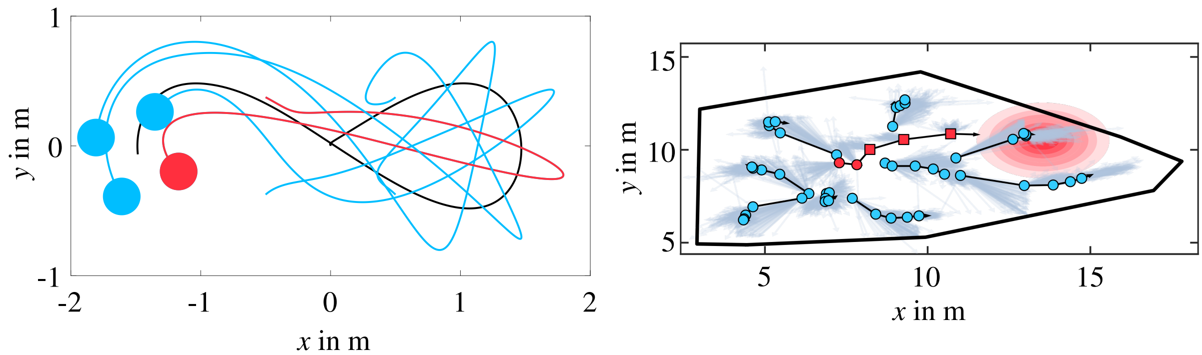

Figure 1: On the left, the red-colored robot's behavior is correctly detected as anomalous by the inverse optimal control approach, revealing that, goal-wise, it puts more emphasis on low control effort than on high accuracy. On the right, in those time instances marked by a square, the red-colored robot is correctly detected as anomalous by the machine-learning-based approach. The antagonistic robot aims to get closer to its red-shaded area of interest instead of following the overall goal of the swarm.

The two proposed identification approaches are furnished as follows. The first approach makes use of the fact that, for technical systems such as robots, the dynamics and physical functioning of the system are often known. In such a case, an antagonistic robot may be subject to the same physical dynamics as a benevolent robot but it may use a controller designed to achieve a different goal. Then, a deterministic behavior model is identified for each robot in the form of a model predictive controller, where the cost function is identified using available observation data from the network to fit the behavior of each robot. This approach fits into the general framework of inverse optimal control [2], which has mainly seen applications in biomechanics as a way to analyze motion in a non-networked setting [3]. Having identified the goals of each robot in the form of an individualized cost function, this contribution examines whether and how this allows to deduce which robot may behave anomalously.

In other applications, an a-priori model of the dynamics might not be available, system behavior may not be fully deterministic, or formulating an inverse optimal control problem may require too much information on the real decision-making structure employed. In such instances, it can be more fitting to directly learn a stochastic behavior model of each system in the network from observed data using tools from machine learning. Hence, this work's second approach builds upon a so-called normalizing flow [4], which is a type of neural network that can be used to approximate an unknown probability distribution from data. The network learns an invertible mapping from a predefined, simple base distribution (e.g., Gaussian or uniform) to the target distribution. This way, based on observation data from benevolent robots, a conditional probability distribution is learned for the actions of a benevolent robot given the current state of the coverage task. It is then expected that an anomalous robot will more often perform actions that are in areas of low probability density, permitting to categorize the robots' dynamic behavior. Still, a direct usage of a normalizing flow can be problematic in coverage tasks since the network's inputs, i.e., its features, which are of fixed dimensionality, need to describe the state of the coverage task. However, depending on the setting and situation, describing the task's state requires a varying number of pieces of information, such as relative positions of a varying number of neighboring robots. To that end, we employ two long-short-term-memory networks (LSTMs) to embed the information on the coverage problem's state in a latent space of fixed dimensionality. One LSTM is used to deal with position information of the robots, whereas the other deals with information on the boundaries of the coverage area. Concretely, the LSTMs are fed the necessary information as a sequence of inputs of variable lengths, whereas their outputs form the fixed-dimension state description employed by the normalizing flow. The whole network structure, consisting of LSTMs and normalizing flow, is trained end to end.

This contribution shows that, when using the two approaches in settings fitting their underlying assumptions, very satisfactory detection rates can be achieved. Furthermore, the limits of the methods are discussed explicitly, also in respect to one another. In addition, it is examined how the picture changes when, instead of operating on data from multibody simulations, data from hardware experiments is employed.

Acknowledgments

This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grants 433183605 and 501890093 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] Ebel, H.: Distributed Control and Organization of Communicating Mobile Robots: Design, Simulation, and Experimentation. Dissertation, Schriften aus dem Institut für Technische und Numerische Mechanik der Universität Stuttgart, Bd. 69. Düren: Shaker Verlag, 2021. DOI: 10.18419/opus-13315

[2] Jameson, A.; Kreindler, E.: Inverse Problem of Linear Optimal Control. SIAM Journal on Control, Vol. 11, No. 1, pp. 1–19, 1973.

[3] Westermann, K.; Lin, J.F.; Kulìc, D.: Inverse Optimal Control with Time-Varying Objectives: Application to Human Jumping Movement Analysis. Scientific Reports, Vol. 10, No. 1, p. 11174, 2020.

[4] Papamakarios, G.; Nalisnick, E.; Jimenez Rezende, D.; Mohamed, S.; Lakshminarayanan, B.: Normalizing Flows for Probabilistic Modeling and Inference. Journal of Machine Learning Research, Vol. 22, pp. 2617–2680, 2021.