Interactive Control and Decision-Making for Multi-Robots Systems

Doctoral Consortium

Yiwei Lyu
Carnegie Mellon University
Pittsburgh, United States
yiweilyu@andrew.cmu.edu

ABSTRACT

In order to achieve the ultimate goal of harmonious human-robot co-existence, the key is to build autonomous robots that can safely interact with humans for collaboration and coordination, as well as demonstrate reliable behavior that is acceptable to humans. These two requirements slightly differ from each other, with the former addressing the safety and functionality of robots as task performers, and the latter emphasizing the social compliance of robots as entities in society. In this abstract, I will outline my efforts towards enhancing the safety and reliability of interactive robot autonomy from three progressively advancing perspectives, 1) self-level autonomy, aiming to develop reactive behavior that ensures safety for individual robots when encountering non-cooperative agents, 2) peer-level autonomy, emphasizing the establishment of a safe interaction mechanism within an diverse and unconnected multi-robot system, and 3) human-involved autonomy, highlighting the consideration of human factors in the decision-making process for the design of multi-robot systems.

KEYWORDS

Multi-Robot Systems; Safe Control; Autonomous Driving

ACM Reference Format:
Yiwei Lyu. 2024. Interactive Control and Decision-Making for Multi-Robots Systems: Doctoral Consortium. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

To promote a harmonious coexistence between humans and robots, prioritizing safety is of utmost importance. Various techniques and tools [7, 10] have been explored to enhance robots' ability to avoid potential collisions, including Velocity Obstacles [4, 5], Reachability Analysis [3, 8, 13], and Voronoi Cell Partitioning [27]. These approaches offer formal safety guarantees and ensure collision avoidance. However, current research often assumes perfect information about agent positions and motions for pre-computed safety guarantees or relies on overly conservative safety behaviors to account for uncertainty due to imperfect robot information, which may lead to unintended consequences, such as an autonomous vehicle decelerating and yielding when it is still a considerable distance away from another vehicle merging onto the ramp [6, 20, 24, 26]. Therefore, **(Q1) how to empower robots to demonstrate adaptive safe behavior under uncertainty without being overly conservative** is an important research question to answer in self-level autonomy, which mainly focuses on the safety and functionality of one single robot under our control when interacting with the environment.

While ensuring safety is already a complex task for single robots, it becomes even more intricate when managing a team of robots [11, 12], especially in scenarios without direct communication channels within the multi-robot system. For example, when mobile robots of different model years and manufacturers work in a shared warehouse, ensuring safe interactions becomes challenging. This necessitates the development of interaction mechanisms that don't rely on hardware communication modules and can adapt to diverse capabilities among robots. Control Barrier Function (CBF)-based methods [1, 2] have been investigated for application in heterogeneous robot teams [24]. However, existing research primarily addresses diverse actuation constraints within robot teams, rather than delving into the integration of various social preferences into decentralized controller design to express intended social implications. In my research agenda, peer-level autonomy serves as a bridge between self-level autonomy and full human-involved autonomy. Its primary goal is to enable robots to exhibit basic socially-aware behaviors when interacting with their peers safely. **(Q2) How to achieve control decentralization and handle the heterogeneity of robot characteristics without compromising the formally provable safety guarantees** is the central challenge.

As we progress towards human-involved autonomy, safety considerations should not be limited to physical state configurations but extended to the psychological aspects of human perception and trust [9, 21, 22, 25]. It is crucial to ensure that interactions with robots, such as self-driving cars, inspire a sense of safety and confidence. Therefore, it is equally essential to explore innovative approaches that empower robots to align their actions with human expectations through socially compliant behavior, ultimately enhancing their acceptability among humans. In this stage, the critical challenge lies in **(Q3) how to enable robots to think and act acceptably to humans by introducing human factors into the controller design of multi-robot systems and exploring novel mathematical formulations to characterize terms like social norms and responsibilities**. This includes demonstrating consequence-aware behavior and role reasoning in interactions.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

2 PREVIOUS WORK

Study 1: Enriched Behavior Characterization in Safety-critical Control. Robots exhibiting overly conservative behavior and limited behavior characterization for safe controllers are two major challenges in building safety assurance for self-level autonomy, which can render the control problem infeasible and cause collision even with a possible solution. In my work [15], I developed a novel bi-level Control Barrier Function-based safe controller specifically designed for self-driving cars. With the integrated chance constraints, this approach allows autonomous vehicles to exhibit behaviors that deviate minimally from the nominal task-related controller while considering various levels of uncertainties from the interaction with other non-cooperative vehicles. A tighter probabilistic safety guarantee is obtained compared to existing approaches that prevents robots from being unnecessarily conservative. More importantly, the co-optimization between *behavioral conservativeness* and *feasibility* in my work provides a new way to characterize different degrees of safe behavior from aggressive to conservative agents, allowing for explicit specification of a robot's reaction speed when approaching the boundary of the minimum safety distance. This finding allows for a generalized vehicle behavior design when interacting with drivers who have diverse driving styles. The approach also extends to high-relative degree systems, considering realistic vehicle dynamics [23], greatly enhancing its general applicability to complicated robot systems.

Study 2: Socially-aware Coordination for Heterogeneous Robot Teams. Peer-level autonomy, such as smart transportation platooning, presents challenges in real-time coordination among robots with diverse capabilities and limited communication. To address the issues of compatibility and computational efficiency, a novel decentralized control scheme is crucial, allowing heterogeneous robots to make socially-aware decisions individually without sacrificing collective task performance. In my research [14, 18], I propose a unified framework for enabling collaboration among heterogeneous robots with various dynamics and characteristics, i.e., willingness to cooperate. The framework [17] relies on implicit coordination, utilizing a novel metric based on Social Value Orientation to quantify relative differences between individuals. This metric is used to compose a decentralized control policy characterizing the admissible control space based on each robot's relative capabilities and characteristics. Importantly, this framework allows each robot to make decisions solely based on local information, eliminating the need to predict other robots' movements. By employing this framework, unconnected heterogeneous robots can effectively coordinate tasks in a socially-aware manner in runtime, while still achieving formally provable safety guarantees. Peer-level autonomy, operating in this phase, acknowledges the differences among heterogeneous individuals and accommodates them within the decentralized controller design. It plays a pivotal role in connecting self-level autonomy with human-involved autonomy, establishing the foundation for developing multi-robot systems capable of making responsible decisions in the future.

Study 3: Responsible Decision-making in Human-robot Interaction. Robots face significant challenges in achieving harmonious human-robot interaction due to fundamental differences from humans in thinking and acting. First, robots need to develop the ability to mimic humans' impact-aware decision-making process. This entails considering the potential consequences and impact of their actions on humans, accounting for factors such as safety, ethics, and social norms. Second, robots must adapt to unpredictable human behavior by dynamically adjusting their responses. To overcome these challenges, my work in [19] proposed the first responsibility-oriented safety-critical controller, that enables robots to make responsible decisions by explicitly reasoning over the propagation of behavior among multiple entities and its impact on the rest of the group. With the integration of a learned human behavior model [16] and risk measurement based on Conditional Value at Risk, it allows robots to dynamically evaluate the inter-agent influence and determine the appropriate portion of responsibility to contribute towards the collective goal without excessive risk or dominance, leading to improved task performance with formal safety guarantees.

3 FUTURE WORK

Study 4: Accountability Reasoning in Interaction. In a fully human-involved autonomous scenario, robots must not only determine their roles and responsibilities during interactions but also consider the level of accountability they should bear if an unintended event, such as a collision, occurs. This retrospective perspective is essential because while previous research focuses on achieving "correct-by-design" robot behavior, accountability reasoning delves into identifying where things went wrong and how to improve when undesirable outcomes happen. Similar to responsibility reasoning, accountability analysis should be based on the joint dynamics and their resulting consequences. However, the key difference lies in the origins of these concepts: responsibility is grounded in robots' capabilities to achieve objectives within a collective goal, while accountability is rooted in counterfactual reasoning, assessing how outcomes might differ if a robot did not take specific actions. Formulating the problem of tracing accountability in a mathematically quantifiable manner offers a valuable tool for future policy or legal studies concerning incidents involving human-robot interaction, such as self-driving cars sharing roads with humans.

Study 5: Explainability of Social Implication in Multi-HRI. In the pursuit of enabling multi-robot systems to seamlessly coexist with humans, another goal is to improve the explainability of the social implications of robots' behavior, ensuring their alignment with prevailing societal expectations. One significant challenge is to establish a cohesive perspective for defining social norms or common human expectations regarding how robots should behave across a range of Human-Robot Interaction (HRI) applications. One of my ongoing effort involves harnessing the concept of accumulated risk introduced in Study 3 to formulate a foundational guideline for socially compliant behavior, drawing inspiration from Isaac Asimov's "Three Laws of Robotics." This unified framework does not necessitate task-specific specifications and empowers us to enhance the clarity of robots' decision-making processes, allowing them to elucidate why they opt for a particular course of action over alternative choices.

REFERENCES

[1] Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. 2019. Control barrier functions: Theory and

applications. In *18th European Control Conference (ECC)*.

[2] Aaron D Ames, Jessy W Grizzle, and Paulo Tabuada. 2014. Control Barrier Function based Quadratic Programs with Application to Adaptive Cruise Control. In *53rd IEEE Conference on Decision and Control*.

[3] Andrea Bajcsy, Somil Bansal, Eli Bronstein, Varun Tolani, and Claire J Tomlin. 2019. An efficient reachability-based framework for provably safe autonomous navigation in unknown environments. In *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 1758–1765.

[4] Daman Bareiss and Jur Van Den Berg. 2015. Generalized reciprocal collision avoidance. *The International Journal of Robotics Research* 34, 12 (2015), 1501–1514.

[5] Andrew Best, Sahil Narang, and Dinesh Manocha. 2016. Real-time reciprocal collision avoidance with elliptical agents. In *IEEE International Conference on Robotics and Automation (ICRA)*. 298–305.

[6] Federico Celi, Li Wang, Lucia Pallottino, and Magnus Egerstedt. 2019. Deconfliction of motion paths with traffic inspired rules. *IEEE Robotics and Automation Letters* 4, 2 (2019), 2227–2234.

[7] Jianyu Chen, Wei Zhan, and Masayoshi Tomizuka. 2019. Autonomous driving motion planning with constrained iterative LQR. *IEEE Transactions on Intelligent Vehicles* 4, 2 (2019), 244–254.

[8] Mo Chen, Jennifer C Shih, and Claire J Tomlin. 2016. Multi-vehicle collision avoidance via hamilton-jacobi reachability and mixed integer programming. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 1695–1700.

[9] Jong Kyu Choi and Yong Gu Ji. 2015. Investigating the importance of trust on adopting an autonomous vehicle. *International Journal of Human-Computer Interaction* 31, 10 (2015), 692–702.

[10] Laurene Claussmann, Marc Revilloud, Dominique Gruyer, and Sébastien Glaser. 2019. A review of motion planning for highway autonomous driving. *IEEE Transactions on Intelligent Transportation Systems* 21, 5 (2019), 1826–1848.

[11] Raffaello D'Andrea. 2012. Guest editorial: A revolution in the warehouse: A retrospective on kiva systems and the grand challenges ahead. *IEEE Transactions on Automation Science and Engineering* 9, 4 (2012), 638–639.

[12] Valerio Digani, Lorenzo Sabattini, Cristian Secchi, and Cesare Fantuzzi. 2013. Towards decentralized coordination of multi robot systems in industrial environments: A hierarchical traffic control strategy. In *2013 IEEE 9th International Conference on Intelligent Computer Communication and Processing (ICCP)*. IEEE, 209–215.

[13] Anjian Li, Liting Sun, Wei Zhan, Masayoshi Tomizuka, and Mo Chen. 2021. Prediction-based reachability for collision avoidance in autonomous driving. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 7908–7914.

[14] Yiwei Lyu, John M Dolan, and Wenhao Luo. 2023. Decentralized Safe Navigation for Multi-agent Systems via Risk-aware Weighted Buffered Voronoi Cells. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*.

[15] Yiwei Lyu, Wenhao Luo, and John M Dolan. 2021. Probabilistic safety-assured adaptive merging control for autonomous vehicles. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10764–10770.

[16] Yiwei Lyu, Wenhao Luo, and John M Dolan. 2022. Adaptive safe merging control for heterogeneous autonomous vehicles using parametric control barrier functions. In *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 542–547.

[17] Yiwei Lyu, Wenhao Luo, and John M Dolan. 2022. Responsibility-associated Multi-agent Collision Avoidance with Social Preferences. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 3645–3651.

[18] Yiwei Lyu, Wenhao Luo, and John M Dolan. 2023. CBF-inspired Weighted Buffered Voronoi Cells for Distributed Multi-agent Collision Avoidance. In *2023 American Control Conference (ACC)*. IEEE.

[19] Yiwei Lyu, Wenhao Luo, and John M Dolan. 2023. Risk-aware Safe Control for Decentralized Multi-agent Systems via Dynamic Responsibility Allocation. In *2023 IEEE International Conference on Intelligent Robots and Systems (IROS)*. IEEE.

[20] Quan Nguyen and Koushil Sreenath. 2016. Exponential control barrier functions for enforcing high relative-degree safety-critical constraints. In *2016 American Control Conference (ACC)*. IEEE, 322–328.

[21] Luis Oliveira, Karl Proctor, Christopher G Burns, and Stewart Birrell. 2019. Driving style: How should an automated vehicle behave? *Information* 10, 6 (2019), 219.

[22] Faisal Riaz, Sohail Jabbar, Muhammad Sajid, Mudassar Ahmad, Kashif Naseer, and Nouman Ali. 2018. A collision avoidance scheme for autonomous vehicles inspired by human social norms. *Computers & Electrical Engineering* 69 (2018), 690–704.

[23] Spencer Van Koevering, Yiwei Lyu, Wenhao Luo, and John Dolan. 2022. Provable probabilistic safety and feasibility-assured control for autonomous vehicles using exponential control barrier functions. In *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 952–957.

[24] L. Wang, A. D. Ames, and M. Egerstedt. 2017. Safety Barrier Certificates for Collisions-Free Multirobot Systems. *IEEE Transactions on Robotics* 33, 3 (2017), 661–674.

[25] Adam Waytz, Joy Heafner, and Nicholas Epley. 2014. The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of experimental social psychology* 52 (2014), 113–117.

[26] Jun Zeng, Bike Zhang, and Koushil Sreenath. 2021. Safety-critical model predictive control with discrete-time control barrier function. In *2021 American Control Conference (ACC)*. IEEE, 3882–3889.

[27] Dingjiang Zhou, Zijian Wang, Saptarshi Bandyopadhyay, and Mac Schwager. 2017. Fast, on-line collision avoidance for dynamic vehicles using buffered voronoi cells. *IEEE Robotics and Automation Letters* 2, 2 (2017).