ALYSSA PIERSON

apierson@mit.edu alyssapierson.com 970.846.8787

Research Statement

Robotics will transform our everyday lives, from home service and personal mobility, to large-scale warehouse management and agriculture monitoring. Across these applications, robots need to interact with humans and other robots in complex, dynamic environments. My research studies how to integrate these underlying cooperation and interaction models into the design of the robots. Understanding how robots interact with other robots and humans allows us to design safer and more robust systems. For example, an autonomous vehicle must understand the intent of human drivers for safe driving. By estimating the personality of the human drivers, the autonomous vehicle can better predict their intentions, resulting in lower errors and safer driving. These interactions also apply to robot teams. Consider a team of drones tasked with taking pictures and monitoring an environment. If we assume all drones are equal and cooperative, performance deficits of a single robot affect the group's outcome. Instead, if we model how those drones interact, the team can compensate for under-performing individuals, improving the group's performance.

Creating a team of capable, collaborative robots requires insight into several challenges in decision-making. The correct response of the robot is intrinsically linked to its task, surroundings, and cooperation of other robots. For applications with mixed human-robot teams, the algorithms must be socially-compliant by design. Robots also need to understand how to quantify cooperation dynamics, and adapt to evolving relationships. As these robots work in complex and dynamic environments, they may encounter unknown robots and obstacles, yet still need navigate safely. My current research focuses on three key directions to address these challenges: (i) how to define cooperation, (ii) how to adapt to evolving cooperation, and (iii) how to cooperate with unknown humans and robots.

Current Research

My research designs new capabilities for collaborative robotic teams by integrating trust, cooperation, and competition models into their design. Exploring the nuances in interaction enables richer control policies with proven performance guarantees. I use tools from social psychology and behavioral decision theory to design the interaction models. Leveraging geometric tools, like Voronoi tessellations, for environment partitioning allows for fast and decentralized decision-making. I draw upon my expertise in game theory, nonlinear control, and Lyapunov stability theory to develop distributed control strategies with provable properties. My research spans mathematical theory, algorithmic design, simulations, and hardware implementations on ground and aerial robots.

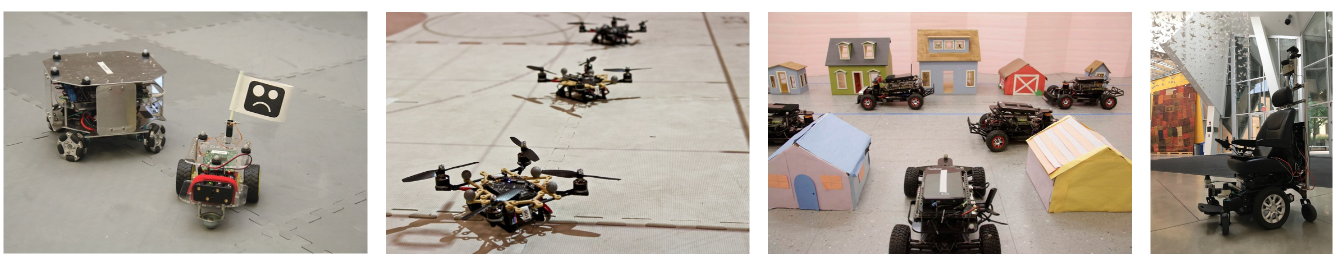

Figure: Experimental validation is crucial for understanding the performance of multi-robot algorithms. My research has been implemented on a variety of ground and aerial platforms that allow us to study cooperation in robot teams.

Towards Socially-Compliant Robots

Integrating robots into our everyday life requires these robots to interact with humans. We can improve the capabilities of these robots by having the robots assess the personality of humans. To better model these interactions, we embed tools from social psychology into the robot's decision making. In particular, we use Social Value Orientation (SVO), a metric that quantifies human social preferences and their corresponding levels of cooperation in social dilemmas. Social dilemmas are conflicts between agents where an individual's short-term self-interest is at odds with the group's longer-term collective interest. An agent's SVO preference informs how they act in social dilemmas by measuring how an individual weights their reward against the reward of others. This translates to social preferences like altruistic, prosocial, egoistic, and competitive. By integrating SVO into our autonomous decision-making, we enable the robots to understand nuances in human decisions, and generate more socially-compliant behavior.

For autonomous vehicles operating among human drivers, it is imperative the vehicles understand human intentions and respond in a predictable manner. Current state-of-the-art vehicles struggle with social dilemmas in traffic, like unprotected left turns or highway merging. An altruistic driver may yield, while an egoistic driver only cares about their own interests and does not yield. In [1], we integrate SVO preferences into a game-theoretic decision-making model. Our model observes human drivers, estimates their SVO preference, predicts their behavior, and responds in a predictable manner. Incorporating social preferences enables the autonomous vehicle to identify when an altruistic driver yields, knowing when to take the merge. We can apply these concepts to improving intersection management [2], as well as prioritized collision avoidance [3], where SVO is a proxy for agent selfishness and priority in the system.

Adapting to Changes in Cooperation

Individual variations in performance and cooperation have a significant impact on the success of the team. If we can encode how cooperation evolves for individual robots, we can derive new policies that improve the team's performance. My work explores how variations in the cooperation dynamics between robots changes the capabilities of the team. One common problem in multi-robot systems is coverage control, which examines how to deploy a group of robots across an environment for tasks such as sensing, resource foraging, or surveillance. For example, a group of drones may be tasked with monitoring crops in a field. One solution assigns robots to cover their Voronoi cell. While this provides all robots with equitable partitions, it assumes all agents are cooperative with equal performance. Any individual errors, such as a drone with a faulty camera, compromise the group's performance. In [4, 5, 6], we incorporate sensing and actuation variations into the robots' definition of the Voronoi partition. The robots have no prior knowledge of their relative performance, but compare information with their neighbors and calculate a trust weighting online. These trust weightings quantify the relative performance and modify the Voronoi partition, such that lower-performing robots receive a smaller region of the environment. By accounting for individual performances, we improve the group's performance in their coverage task.

Another application of multi-robot teams is robotic herding and wildlife management. Currently, large-scale cattle mustering requires expert helicopter pilots to move herds with low-flying and dangerous maneuvers. Robotic wildlife management has the promise to reduce human accidents, while simultaneously providing more information about the animals. Here, our multi-robot team comprises two types of agents: the robotic herders under our control, and the "sheep" agents, which are not under our control. The challenge then becomes designing control policies for our robotic herders. While we don't control the herd, we do know they react to our herders by running away. Our key design insight is to enforce geometric relationships that allow for the combined system dynamics of the herders and herd to reduce to a single vehicle [7, 8]. Under this simplified system, finding the control policies of the herders is a straightforward mapping of the proxy vehicle controls, allowing us to drive the herd to a desired goal or follow paths.

Multi-robot teams can also be used in tracking and surveillance. This may include: finding a lost herd member, studying wildlife, or tracking a malicious person in a city. In an environment with no obstacles, we can use Voronoi tessellations to guarantee that one or more pursuers will capture any number of evaders in finite time [9]. Here, the Voronoi cell of the evader defines its safe-reachable area. The pursuers choose a strategy that minimizes the area of the evader, guaranteeing the evader's capture. We can also modify the Voronoi definition to account for static obstacles [10], such that the pursuers can track the target evader while avoiding collisions with obstacles. In some environments, the pursuers may not be able to follow the target to all areas of the environment. For instance, pursuer drones may not be able to enter buildings in cities or thick forests, but the target can freely move throughout the environment. When the target enters a no-fly zone, the pursuers use strategies from Voronoi-based coverage control to distribute themselves about the no-fly zone boundary, minimizing the expected time to capture once the evader re-emerges [11]. For these pursuer-evader games, we not only provide theoretical guarantees on the robot's performance, but also implement the algorithms on ground robots and quadrotors to demonstrate their real-time efficiency.

Risk and Cooperation with Unknown Agents

Current large-scale multi-robot systems often separate the robots from interacting with humans. If we want to integrate robotic systems into our everyday life, the robots must learn to safely interact with unknown humans, other robots, and environments. An autonomous wheelchair designed for personal mobility needs to adapt to the human's environment, not expect the human to adapt. New advancements in machine learning and semantic reasoning are allowing these robots to operate in more complex environments, however, the robot must also safely maneuver even if it can't identify what it sees. When the robot operates around unknown agents and entities, it is necessary to understand the risk associated with its decisions.

For autonomous vehicles and personal mobility platforms, if we can define an allowable risk tolerance for the vehicles, then we know the vehicle will make safe decisions when the risk is below a threshold. We anchor our risk by reasoning about the probabilities of collision for the vehicle. In [12], we construct a risk model that computes the probability of an incident or collision, based on factors such as other vehicles, attention spans, noisy sensors, and occlusions. Through experiments on one-tenth scale race cars, we evaluate risk at intersections during left turns across traffic, without relying relying on explicit communication with other vehicles. We demonstrate our approach is effective for safe traversal of intersections by autonomous vehicles, and as a shared control system with human drivers.

In some cases, the autonomous vehicle will need to avoid collisions with objects it cannot identify or track. For this, we consider what a "safety net" might look like for the system. We introduced Risk Level Sets [13, 14, 15], which provide a rapid assessment of how risky a location is to the agent, using limited dynamics information. Intuitively, a point is high-risk if it is an obstacle, or if it is in the direct path of a moving obstacle. Risk Level Sets allow the agent to account for the unknown clutter of an environment by approximating a risk of collision. We implemented our algorithm on an autonomous wheelchair navigating crowded hallways, demonstrating that the robot can avoid collisions with pedestrians, tables, strollers, and all the other random clutter of the human environment without the robot needing to explicitly detect, track, or identify these objects.

Future Research Directions

My future research aims to explore how these multi-robot systems can enhance capabilities of humans and enable new exploration frontiers. Building upon my current research, I aim to design socially-aware algorithms, create adaptable and robust teams, and increase the capabilities of multi-agent teams operating in unknown environments.

Socially-Aware Planning

In [1], we began to explore how encoding Social Value Orientation (SVO) preferences enhances the performance of robotic systems. We demonstrated that for autonomous vehicles, SVO preferences create more socially-compliant behavior in mixed human-autonomous systems. For my future work, I plan to continue extending the integration of social psychology and behavioral sciences into the fundamental design of robotic systems. SVO provides one metric of assessing cooperation preferences, but does not provide a full definition of personalities or explain all behavioral decisions. To create socially-aware systems, robots will need to understand these social preferences, how they evolve over time, and how they change depending on the group dynamics. By using tools from psychology and game theory, we can create robots with personalities that have nuanced, clever, and efficient behavior. Furthermore, the decisions of these systems will be intuitive for humans to understand.

Emergent Collaborative Teams

My research explores how cooperation impacts applications such as coverage control [4, 5, 6], herding [7, 8], and tracking [9, 11, 10]. Across these applications, most of the prior work assumes that robot teams are predefined. As future work, I will study how to design control policies that encourage emergent team-forming, instead of being assigned a team. For example, consider a network of wildlife-monitoring drones, tasked to study animal populations and take photos. A local farmer wishes to deploy a robot to survey their fields, and some of the wildlife-monitoring drones overlap with the farmland. These two types of robots have different capabilities and different objectives. With emergent collaboration, robots would be allowed to talk, exchange information, and assist other teams of robots. To enable this technology, robots need to not only understand the cooperativeness of other robots, but also task alignment between the groups. This requires learning policies online of other robots, modeling their decision-making, and modifying individual behavior to maximize the benefit from other robots while minimizing inefficiencies that may arise due to competition or duplication of tasks.

Uncertainty-Aware Cooperation

The current state of the art multi-robot algorithms often assume perfect sensing, localization, and communication abilities of the team. Similarly, applications like warehouse management robots require the environment to be engineered to the robots. I aim to design robots that do not require engineering an environment; instead, the robots should adapt by understanding how risk and uncertainty affect their decision-making. Our prior work used risk to define safety nets for autonomous wheelchairs [14], but it knew the environment map and path it should follow. If the robot doesn't know the environment, we need to give it the tools to explore, reduce its uncertainty about its surroundings, and better navigate. The quality of information available to the robots directly impacts its ability to make decisions. By combining information-seeking behavior with game-theoretic reasoning, we find the agents will take actions that increase their information gain, which improves their ability to make decisions with reduced uncertainty [16]. It is not only important to reason about a robot's own uncertainty, but also the uncertainty of other robots and humans in the environment, and how actions propagate changes in uncertainty. This co-design of control algorithms with complex uncertainty models will enable robots that adapt to the environments around them, creating more agile and robust teams.

Selected References

[1] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, “Social behavior for autonomous vehicles,” *Proceedings of the National Academy of Sciences*, 2019, accepted.

[2] N. Buckman, A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Sharing is caring: Socially-compliant autonomous intersection negotiation,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, November 2019.

[3] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Weighted buffered voronoi cells for distributed semi-cooperative behavior,” 2020, submitted.

[4] A. Pierson and M. Schwager, “Adaptive inter-robot trust for robust multi-robot sensor coverage,” in *Robotics Research: The 16th International Symposium (ISRR)*. Switzerland: Springer International Publishing, Dec 2013, pp. 167–183.

[5] A. Pierson, L. C. Figueiredo, L. C. A. Pimenta, and M. Schwager, “Adapting to performance variations in multi-robot coverage,” in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, May 2015, pp. 415–420.

[6] A. Pierson, L. Figueiredo, L. C. Pimenta, and M. Schwager, “Adapting to sensing and actuation variations in multi-robot coverage,” *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 337–354, 2017.

[7] A. Pierson and M. Schwager, “Bio-inspired non-cooperative multi-robot herding,” in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, May 2015.

[8] A. Pierson and M. Schwager, “Controlling noncooperative herds with robotic herders,” *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 517–525, April 2018.

[9] A. Pierson, Z. Wang, and M. Schwager, “Intercepting rogue robots: An algorithm for capturing multiple evaders with multiple pursuers,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 530–537, April 2017.

[10] A. Pierson and D. Rus, “Distributed target tracking in cluttered environments with guaranteed collision avoidance,” in *2017 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, Dec 2017, pp. 83–89.

[11] A. Pierson, A. Ataei, I. C. Paschalidis, and M. Schwager, “Cooperative multi-quadrotor pursuit of an evader in an environment with no-fly zones,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 320–326.

[12] S. G. McGill, G. Rosman, T. Ort, A. Pierson, I. Gilitschenski, B. Araki, L. Fletcher, S. Karaman, D. Rus, and J. J. Leonard, “Probabilistic risk metrics for navigating occluded intersections,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4322–4329, Oct 2019.

[13] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Navigating congested environments with risk level sets,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 1–8.

[14] A. Pierson, C. Vasile, A. Gandhi, W. Schwarting, S. Karaman, and D. Rus, “Dynamic risk density for autonomous navigation in cluttered environments without object detection,” in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 5807–5814.

[15] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Learning risk level set parameters from data sets for safer driving,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, June 2019, pp. 273–280.

[16] W. Schwarting, A. Pierson, S. Karaman, and D. Rus, “Stochastic dynamic games in belief space,” 2019, arXiv:1909.06963.