Received: 14 June 2024 | Accepted: 14 August 2024

DOI: 10.1002/aaai.12192

HIGHLIGHT

Towards smooth mobile robot deployments in dynamic human environments

Christoforos Mavrogiannis

Department of Robotics, University of Michigan, Ann Arbor, Michigan, USA

Correspondence
Christoforos Mavrogiannis, Department of Robotics, University of Michigan, Ann Arbor, MI, USA.
Email: cmavro@umich.edu

Abstract

Recently, there has been great interest in deploying autonomous mobile robots in airports, malls, and hospitals to complete a range of tasks such as delivery, cleaning, and patrolling. The rich context of these environments gives rise to highly unstructured motion that is challenging for robots to anticipate and adapt to. This results in uncomfortable and unsafe human–robot encounters, poor robot performance, and even catastrophic failures that hinder robot acceptance. Such observations have motivated my work on social robot navigation, the problem of enabling robots to navigate in human environments while accounting for human safety and comfort. In this article, I highlight prior work on expanding the classical autonomy stack with mathematical models and algorithms designed to contribute towards smoother mobile robot deployments in complex environments.

INTRODUCTION

While the conventional robotics paradigm involves robots locked in cages or spaces heavily engineered just for them, in recent years, we have seen many robot deployments in unstructured and dynamic environments. These include autonomous cars on the road, delivery robots in campuses and sidewalks, patrolling robots in malls, guide robots in airports, home robots, and robots in healthcare. While these deployments have underscored the promise of robotics for improving productivity and handling tedious and laborious tasks in real-world settings, they have also uncovered many technical limitations exemplified by failures resulting in discomfort, accessibility and safety hazards, and public dissatisfaction.

The technical challenges encountered during robot deployments have motivated research that goes beyond the classical autonomy stack, and expands towards understanding more deeply the deployment domains, the users, and their dynamics (Mavrogiannis, Baldini et al. 2023). This is part of a broader community effort within the field of computational human–robot interaction (Thomaz, Hoffman, and Cakmak 2016), which addresses research questions related to affording robots with the ability of accounting for social considerations when operating close to humans. Such questions often require a deeper understanding of human–human interactions and the development of mathematical models describing them.

In this paper, which is part of the AAAI New Faculty Highlights Program, I discuss my prior work on social robot navigation, highlighting how insights from human factors led to the development of mathematical models, algorithms, and systems enabling smoother human–robot interactions. I start by discussing the coupling of human and robot motion in densely crowded domains, and describe a mathematical representation that formalizes it. Then, I emphasize the value of understanding the robot's operation domain, which may empower even simple models to enable robust performance. Acknowledging that even the most robust algorithms will experience failures, I discuss practical ways in which a robot may seek human help from bystanders to return to operation. Finally, I highlight the sensitivity of users to robot motion, and discuss how motion could be used to signal important internal robot states and capabilities, ensuring informed users' expectations. I close with open questions and direction for future development.

FORMALIZING THE HUMAN–ROBOT MOTION COUPLING

Social order in pedestrian navigation relies on a form of cooperation: pedestrians assume a shared responsibility in resolving conflicts and ensuring safety and comfort of each other as they navigate crowded spaces. Wolfinger (1995) described this protocol as a "pedestrian bargain," a negotiation that takes place under two social rules: *pedestrians must navigate competently and they must trust that copresent others do the same*. This negotiation takes place through a tight connection between perception and action (Warren 2006) at the core of which is a subtle information exchange leveraging multiple modalities such as body posture, eye gaze, and gestures.

How do we describe this process to a robot? Prior work has leveraged the extraction of domain-specific crowd motion patterns (Ziebart et al. 2009), the goal-directedness of human behavior (Trautman et al. 2015), and a balance between continuous and discrete decision-making rules that humans employ when navigating in crowded spaces (Kretzschmar et al. 2016). Building on this prior work, we are inspired by psychology studies highlighting the tendency of humans to attribute goals to observed actions of others (Csibra and Gergely 2007). In crowd navigation, the notion of a goal is not well defined. A pedestrian's exact destination is not very relevant to a nearby pedestrian in the moment of handling a navigational conflict leading to a possible collision. Our insight is that *the intent over a direction of passing* is much more relevant, and being able to read it quickly may reduce human cognitive load and improve human performance in collaborative tasks (Carton, Olszowy, and Wollherr 2016). In crowded environments, the robot's directionality of passing is coupled in space and time with the respective passing directionality of other pedestrians. We formalize this coupling using tools from topology.

A formalism of passing. Humans intuitively understand the process of *passing* another pedestrian as a mechanism for resolving a potential conflict. Typically, this amounts to making a decision on whether to move to the left, right, before, or after another pedestrian. Such decisions have topological signatures that can be identified using topological invariants (Berger 2001). Based on this idea, we described a formalism of *passing* between two agents as a pairwise winding number (Mavrogiannis, Balasubramanian et al. 2023), an invariant that corresponds to the signed number of times that a pair of pedestrians encircle each other. The absolute value of this number represents the passing *progress* whereas its sign represents the passing *side*. Based on this representation, we designed a cost function that when minimized results in actions expediting pairwise passing maneuvers. This insight has been the foundation for the development of a family of reactive crowd navigation controllers that monitor and expedite the process of passing other agents in crowded scenes (Mavrogiannis et al. 2022; Mavrogiannis, Balasubramanian et al. 2023) (see Figure 1) and driving environments (Roh et al. 2020) (see Figure 2).

Topological braids. The idea of *passing* can be generalized to environments with multiple agents. Each passing leaves a signature that can be symbolically represented as a

FIGURE 1 Honda's experimental ballbot (Honda 2019) navigates next to three users in our lab experiment (Mavrogiannis, Balasubramanian et al. 2023).

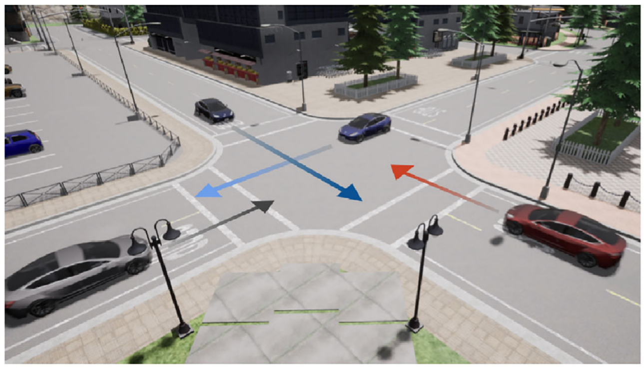

FIGURE 2 Navigating an uncontrolled intersection by reasoning about topological interactions (Roh et al. 2020).

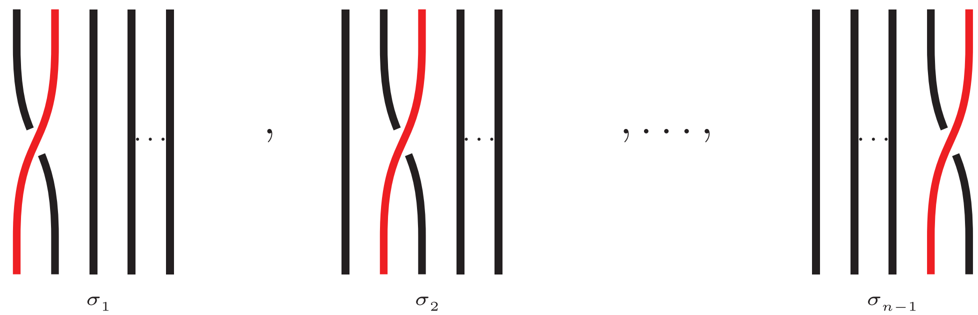

FIGURE 3 Presentation of the braid group, B_n. The group can be generated by the $n-1$ elements shown above, called generators (and their inverses), using an operation called composition.

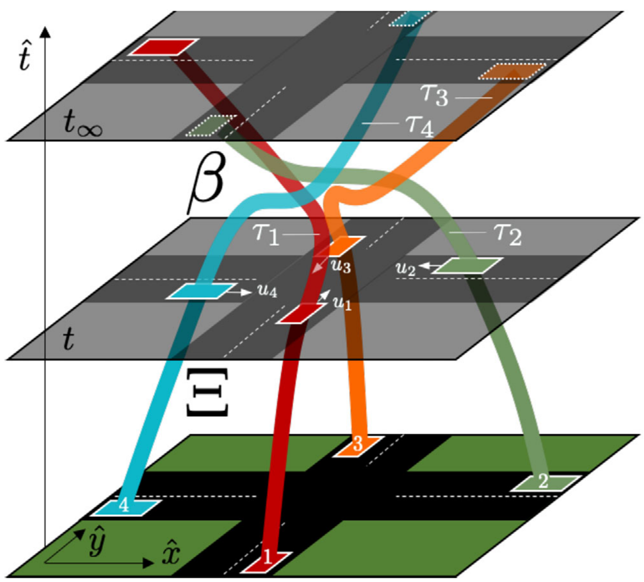

FIGURE 4 Crossing an uncontrolled intersection by reasoning about the topology of the spatiotemporal interaction among agents (Mavrogiannis, DeCastro, and Srinivasa 2023). At time t, given state history Ξ, the ego agent (red), following path τ_1, predicts the topology β of the unfolding multiagent interaction.

topological braid (Birman 1975). Geometrically, a braid is a collection of strings (called strands) anchored between two planes (Thiffeault 2010). The set of all braids on n strands, along with a composition operation forms the braid group on n-strands. The group can be generated via composition of the generators shown in Figure 3. We have used braids as primitives describing the coupling between the motion of the robot and other agents over space and time (see Figure 4). In that sense, a braid is a representative of a strategy that a set of agents engage to pass each other on their way to their destinations (Mavrogiannis and Knepper 2019). By reasoning about the probability over a set of relevant braids, a robot may make navigation decisions without requiring detailed predictions about the motion of other agents. We have explored this idea in the context of crowd navigation (Mavrogiannis and Knepper 2021) and autonomous driving (Mavrogiannis, DeCastro, and Srinivasa 2023), demonstrating practical performance across a series of challenging scenes.

Grouping. A practical way that humans employ to navigate in densely crowded scenes (e.g., the train station of Shinjuku or the crossing of Shibuya in Tokyo) is to detect groups of others and follow them, rather than making individual predictions about each of them. Gestalt theories from psychology (Koffka 1935) suggest that organisms tend to perceive and process *formations of entities*, rather than individual components. Inspired by this idea, we developed a mechanism that first clusters similarly navigating humans into groups (Wang and Steinfeld 2020) and then outputs a prediction of the social space that each group would occupy in the near future. We demonstrated that based on such predictions, a model predictive controller could smoothly navigate through a crowd while avoiding separating groups of conavigating pedestrians (Wang, Mavrogiannis, and Steinfeld 2022).

ROBUST MOBILITY DOES NOT ALWAYS REQUIRE COMPLEX MODELS

A common recent paradigm in social robot navigation involves the use of deep neural architectures as tools for high-fidelity human motion prediction. While these architectures exhibit impressive performance on offline benchmarks (Rudenko et al. 2020), their poor generalization to novel situations, and their limited interpretability call for alternative approaches when considering safety-critical applications like deployments close to humans. Motivated by such observations and by similar findings from the motion-tracking community (Schöller et al. 2020), we have been exploring the promises and limitations of simpler, domain-specific models for navigation in crowded spaces. Coupled with the appropriate representations of the coupling between human and robot motion (Mavrogiannis, Balasubramanian et al. 2023), our insight is that these models can provide a foundation of performance that can

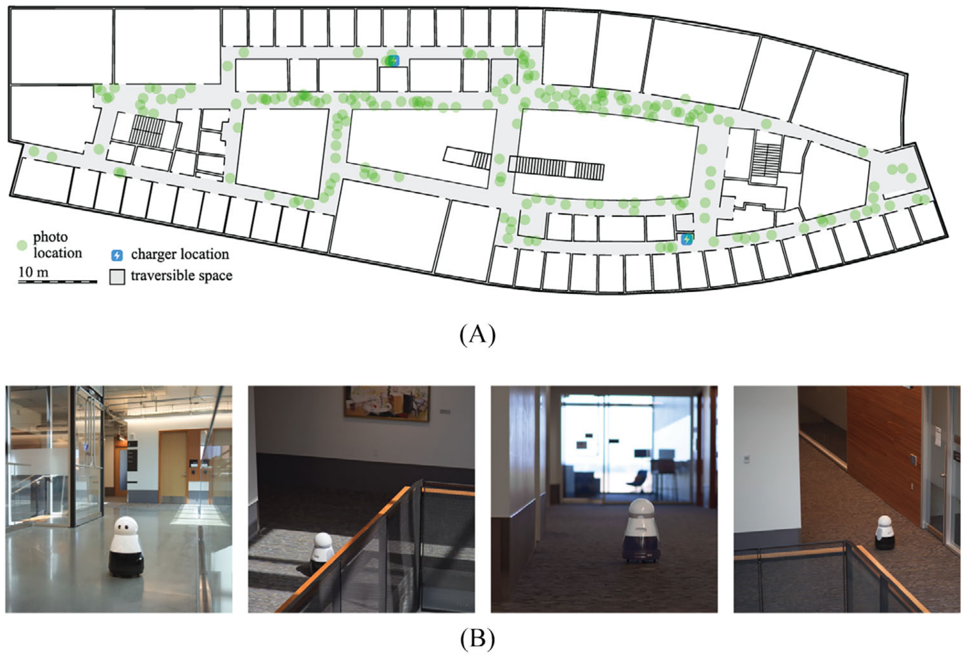

FIGURE 5 Our Kuri robot wandered a 28,000 ft^2 floor for 4 days with minimal human help (Nanavati et al. 2022). (a) Kuri's coverage of the building floor. (b) Photos of Kuri as it wandered in the environment.

be further finetuned online through repeated interactions with users and the environment.

Social robot navigation with constant velocity prediction

Our passing-aware controller (Mavrogiannis, Balasubramanian et al. 2023) integrates multi-agent trajectory prediction to produce adaptive navigation performance in crowded environments. While developing it, we experimented with a series of different models by deploying them on our self-balancing robot platform and subjecting the robot to a diverse range of crowd conditions in the lab (see Figure 1). We found that constant velocity (CV) prediction (Poddar, Mavrogiannis, and Srinivasa 2023) performs comparably to a recent state-of-the-art motion prediction baseline (S-GAN) (Gupta et al. 2018) across a range of crowd behaviors including aggressive or inattentive agents. While our framework could be further expanded through online model improvements and adaptation to different environments and users, its core navigation capabilities could enable a robust performance threshold during the initial stages of deployment.

Localization-free field deployment

One of the challenges preventing the smooth prolonged deployment of mobile robots in complex spaces is the need for consistently accurate robot localization. Despite the important advances in simultaneous localization and mapping (SLAM) (Cadena et al. 2016) over the past few decades, robots deployed in indoor environments are still prone to delocalization, which may require impractical workspace engineering and extensive human interventions to address.

While localization remains an important skill for any robot, there is a long-history of highly effective *localization-free* systems (Brooks 1986; Kinzer 2009; Bennett 2021). By *wandering* in space, such systems are capable of completing a wide range of tasks, especially coverage-based, such as vacuum cleaning (Bennett 2021) or patrolling. Inspired by their effectiveness, we developed a wandering system, which we deployed on a Kuri robot in our academic building (Nanavati et al. 2022). Via Lidar and bump-sensor readings, our system generates a local costmap representing proximity to obstacles. It fixes the direction of lowest cost and passes it for execution on a local controller. If the robot gets stuck due to some obstruction, it updates its costmap to trigger the selection of a new direction. If the robot remains stuck for a prolonged period, it initiates a recovery procedure involving rotation in place and backing off. These behaviors allowed the robot to recover from typical failure modes such as getting stuck on furniture or trapped with a tread off of a cliff. Despite its simplicity, this system enabled the robot to navigate the massive hallways of our academic building (area 28,000 ft^2) for 4 days (see Figure 5).

While alternative tasks involving point-to-point navigation would require a localization system in place, our system can empower robots with even relatively weak compute and sensing features to perform practical coverage tasks. It could also serve as a navigation mechanism to support data collection for building and updating environment maps or for refining the robot's localization system. Finally, it could serve as a backup navigation system in cases of failure of the main localization module.

BYSTANDERS CAN ENABLE SCALABLE ROBOT RECOVERY

When autonomy inevitably fails, human help can be crucial for robot recovery (Weiss et al. 2010; Chi et al. 2020). Typically, researchers and engineers are responsible for ensuring continued robot operation during studies and field deployments. However, this paradigm may not be scalable: while robots can deliver value on many important applications, autonomy can be expected to be brittle and prone to frequent failures. While some of the failures require high expertise and close attention, many of the common failures could be addressed with simple and quick actions (e.g., responding to a robot question, shaking the robot to get it unstuck from a motion planning local minimum, pushing the robot to a new location, moving the robot to its charger). Our insight is that for such types of failures, bystanders could enable scalable robot recovery. This insight was discussed in earlier work (Weiss et al. 2010; Rosenthal, Biswas, and Veloso 2010; Rosenthal, Veloso, and Dey 2011; Thomason et al. 2019) but also motivated by our experience deploying Kuri (Nanavati et al. 2022) in our academic building for 4 days, as part of a user study.

Extended operation via human help

Our study with Kuri involved the robot wandering the 2nd floor of the Gates Center at the University of Washington (approximately 28,000 ft^2), taking pictures of its surroundings and asking in real-time users for feedback through a chatbot application deployed in a department-wide digital workspace (Slack) (Nanavati et al. 2022). To ensure good coverage of the possible artistic themes present in the environments, the robot needed to keep running throughout the workday. Thus, we expanded the chatbot to message the research team when it would get physically stuck or when it would be low on battery, and added a simple diagnostic tool streaming Kuri's front camera feed. In case, something was wrong, a researcher would attend to the matter and put the robot back to operation. These simple monitoring tools—in conjunction with the localization-free system discussed in the previous section—enabled the robot to achieve a substantial coverage of the floor (see Figure 5a). Overall, the researchers did not spend more than 30' helping the robot over the 4 full days (32 h) of the study, illustrating the practicality of enabling continued robot operation through periodical, nontechnical human help. Broadening this idea, we envision that with the right human–robot communication interface (e.g., dialog, display, motion) robots could solicit and leverage human help from bystanders when necessary to keep delivering a productive performance for prolonged periods.

Effectively soliciting bystander help

Soliciting help from a bystander is different than soliciting help from a researcher dedicated to a study or engineering goal: a bystander typically has no clear incentive to help the robot. Thus, it is important for the robot to reason about bystanders' internal states and context. For instance, a robot that asks for help too often or at the wrong times might end up annoying users and quickly stop getting help from them. Considering an office environment setting, we developed a system that plans effective help requests based on past interactions with users (Nanavati et al. 2021). We instantiated this setting in a virtual world where a delivery robot is tasked with visiting offices to deliver mail while a human worker performs computer-repair tasks in the same space. We modeled the robot's task as a Bayes-Adaptive Markov Decision Process (BAMDP) where the robot's goal, expressed in its reward function, is to maximize the number of offices it visits while minimizing the number of human help requests it makes. The transition function returns the probability of the user helping given contextual factors (i.e., the human's—assumed observed—busyness and the frequency of past help requests) and individual factors (i.e., the user's latent helpfulness, estimated from past interactions with them). The model was estimated using Generalized Linear Mixed Models (GLMM) regression from a dataset collected in a virtual office environment, created using the Phaser3 framework (see Figure 6). Through an evaluation user study, we found that our system, integrating both individual and contextual factors significantly outperformed baseline systems (using help models using either only contextual or only individual factors) in terms of accrued rewards, while managing to generate more effective help requests.

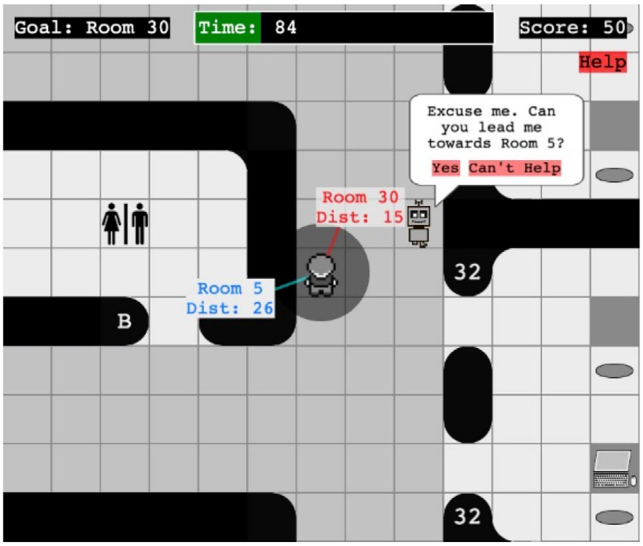

FIGURE 6 A user performs tasks in our virtual office environment while a robot periodically asks them for help. In this environment, we deployed and tested our system for planning effective human help requests (Nanavati et al. 2021).

ROBOT MOTION CAN ADJUST USER EXPECTATIONS

While robots are increasingly entering homes, airports, and streets, users and bystanders often have limited mental models about how robots make decisions. Naturally, users tend to make attributions—often anthropomorphic—mapping robot behavior to possible robot capabilities, intentions, or internal states (Sung et al. 2007). Robots driven by purely functional objectives may complete their tasks but while doing so, they may produce behaviors that confuse users or mislead them about the robot's capabilities and incentives. Prior work has shown that integrating models of human inference into motion planning may enable an observer to guess the robot's goal (Dragan and Srinivasa 2014), the robot's inability to complete its task (Kwon, Huang, and Dragan 2018), or aspects of style and attitude (Knight and Simmons 2016). However, as the robot operates in the presence of humans, its behaviors will also communicate global, long-term behavioral attributes about its decision making mechanism, incentives, and internal states. By managing the types of attributions that a robot elicits from observers as it completes a task, it may manage users' expectations, and shape their impressions as desired.

Understanding users' impressions of robot motion

To study human impressions of different robot navigation strategies, we developed a fictional factory setting

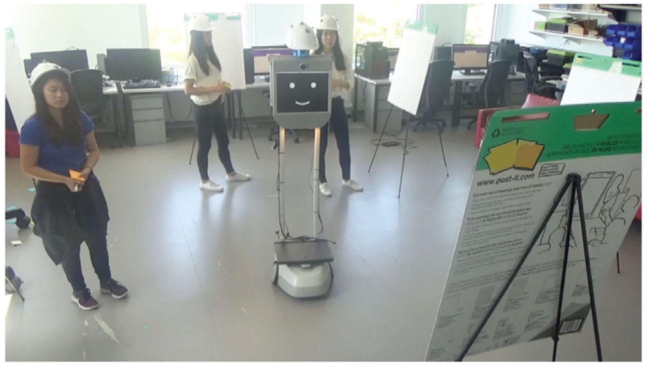

FIGURE 7 Benchmarking social robot navigation in the lab (Mavrogiannis et al. 2022). The setup comprises a telepresence robot and a set of six easels representing machines in a fictional factory workspace. Three participants, wearing tracking helmets navigate between stations to perform fictional maintenance tasks on the machines.

mockup in the lab, (Figure 7) where three users navigated between a set of machines to perform maintenance tasks while one robot was moving around inspecting their work (Mavrogiannis et al. 2022). This setting allowed us to motivate complex navigation encounters between the users and the robot while ensuring natural human walking. Considering a within-subjects design where conditions represented navigation strategies, we compared users' performance and correlated it with their self-reported impressions. A highlight of our findings was that our algorithm (Social Momentum Mavrogiannis et al. 2022, an algorithm designed to generate legible motion in multiagent domains) enabled users to navigate with lower accelerations next to our robot. This was reflected in their open-form responses in which they often noted that our robot was *not noticeable*, whereas baselines elicited responses referring to violations of personal space or unpredictability of robot motion. The coding scheme used to analyze users' open-form responses (see Figure 8) is indicative of the range of human impressions when interacting closely with mobile robots.

Shaping users' impressions via robot motion

In the previous work, user impressions were a byproduct of robot navigation strategies but not explicitly accounted for. To enable robots to control for the types and intensities of attributions they broadcast to human observers, we developed a data-driven methodology for mapping robot motion to attributions (Walker et al. 2021).

Considering a coverage navigation task (e.g., vacuum-cleaning robot) in a virtual home environment, we asked

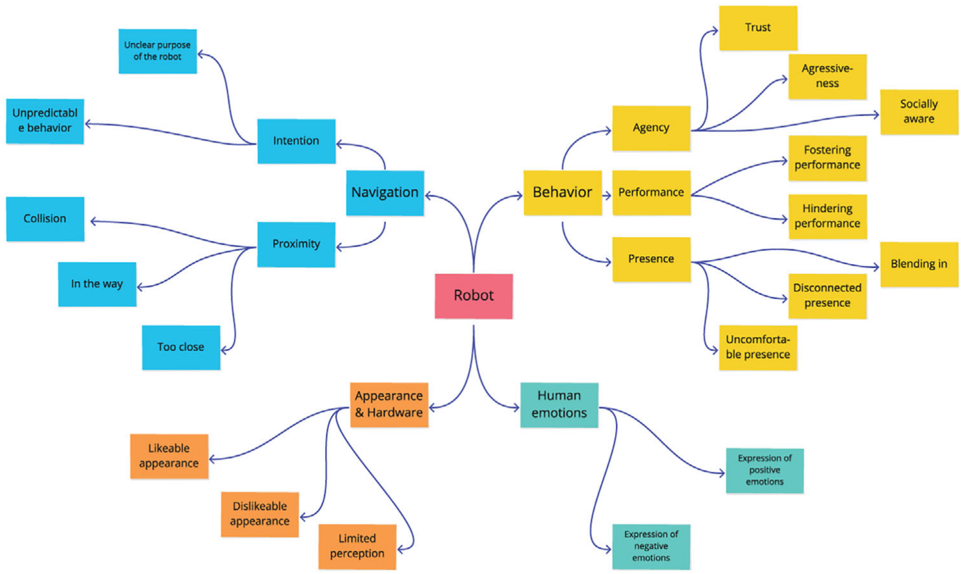

FIGURE 8 Coding scheme used to analyze users' impressions during robot navigation experiments in the lab (Mavrogiannis et al. 2022).

users to rate a wide range of robot behaviors (Bartneck et al. 2009). Through a factor analysis, we extracted a space of attributions that users typically made; these were related to *Competence*, *Curiosity*, and *Brokenness*. Through an active-learning methodology, we guided additional data collection steps that enabled us to train probabilistic models (Mixture Density Networks) mapping robot motion to attributions that an observer would make to describe it. Using these models, we developed a trajectory optimization framework that balanced between the task-related objective of coverage and the communicative objective of eliciting a desired attribution from the user. Through an online evaluation user study instantiated in Amazon Mechanical Turk, we demonstrated that our framework was able to autonomously generate robot motion eliciting desired attributions of desired intensity from users (see Figure 9). Our complete framework is visualized in Figure 10.

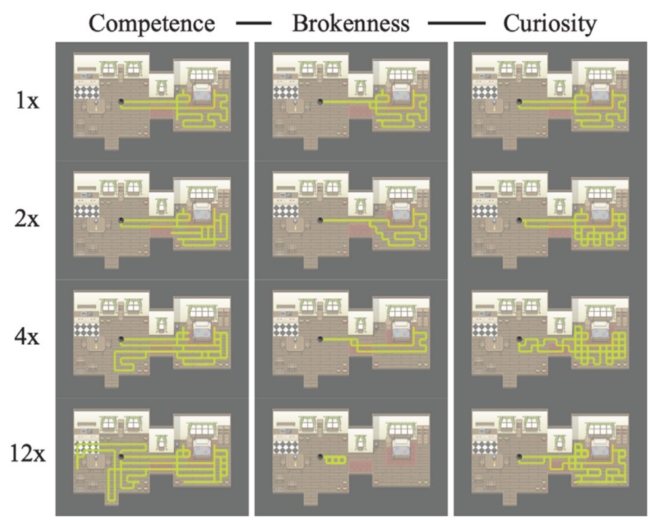

FIGURE 9 Robot trajectories that communicate desired attributions generated by our trajectory optimization framework (Walker et al. 2021). Columns indicate the type of attribution and rows indicate the intensity of the attribution.

DISCUSSION

Prepared in the context of the AAAI New Faculty Highlights Program, this article is a summary of past work conducted by myself and collaborators in developing theory, models, and systems targeting the problem of robot navigation in crowded environments. The topics discussed build a bridge between the classical robot autonomy stack with social sciences and human factors in an effort to lead the community towards the development of robots with a greater human awareness. The insights extracted from our studies are meant to support the deployment of interactive mobile robots in indoor environments by guiding users through a path towards satisfaction (Wise 2018) and acceptance (Davis 1989; Beer et al. 2011) as robots continue to improve through lifelong interactions with their environment (Thrun 1994). However, many additional considerations must be made to ensure safe, smooth, and effective mobile robot deployments involving close-interaction settings:

Safety. Crucially, when interacting with users, it is important to develop safety assurances for the user, especially as the robot learns. There is an extensive body of work on approaches that directly address aspects of safety in human–robot interaction (Lasota, Fong, and Shah 2017)

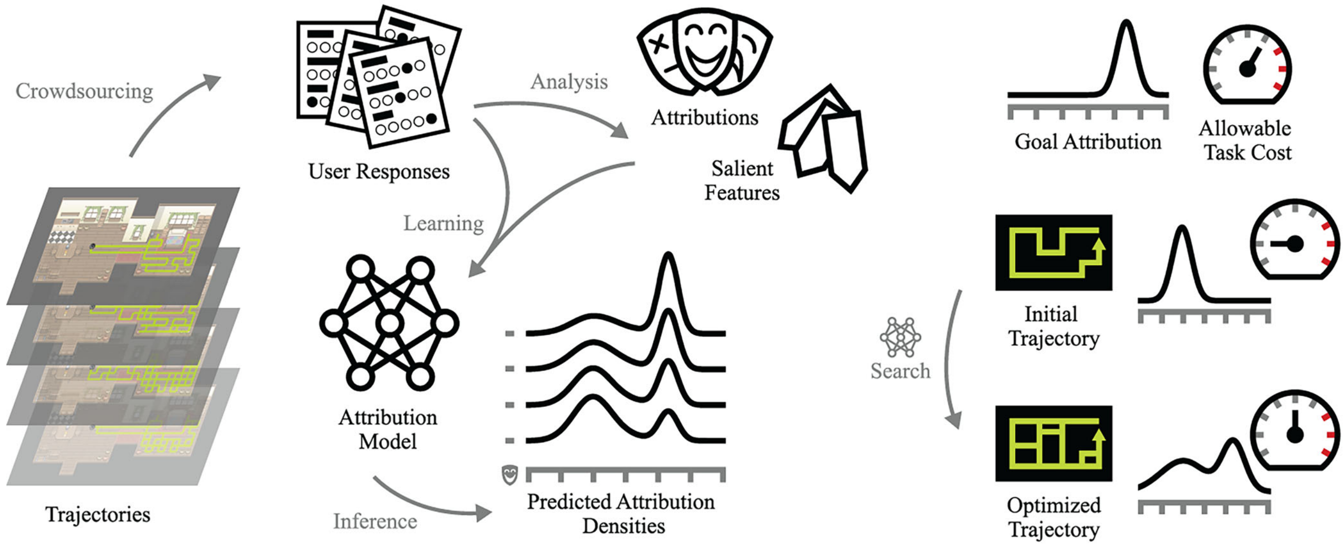

FIGURE 10 Generating robot motion that communicates desired behavioral attributes (Walker et al. 2021). From left to right: user responses to robot trajectories are first analyzed to extract salient features and attributions; then they are used to train a model that probabilistically maps robot trajectories to human attributions; the acquired model is integrated into an optimizer to generate robot trajectories that elicit a desired attribution.

and safe reinforcement learning (Berkenkamp et al. 2017) that is relevant for real-world deployments of continually learning robots.

Imperfect user feedback. While bystanders can be an effective source of feedback for lifelong learning robots, it is important to account for the fact that their feedback will often be imperfect and even inaccurate. Recent work on the development of interactive reinforcement learning could be applicable to enable robots to reason about the quality of human feedback (Kessler Faulkner, Schaertl Short, and Thomaz 2020).

Social awareness. Understanding and reacting to the dynamic social context of a complex environment like a pedestrian domain, a warehouse, or a hospital remains an open challenge. While aspects like proxemics have been increasingly integrated in the design of navigation algorithms (Kirby 2010), additional considerations must be made including cultural and individual adaptation, and accommodation of the requirements of the deployment domain.

Benchmarking. It is important that baseline policies deployed in critical real-world domains are already sufficiently advanced before interacting with real users. Doing so requires mature validation methodologies that capture critical aspects of real-world interaction. While there have been efforts towards formalizing protocols for the validation of social navigation policies (Stratton, Hauser, and Mavrogiannis 2024), additional research is required to develop realistic simulators, evaluation criteria, and benchmark experiments design of realistic simulators but also the definition of benchmark experiments (Mavrogiannis, Baldini et al. 2023).

Technological challenges. Many technological limitations get in the way of smooth robot deployments. For instance, despite the maturity of perception approaches for localization and people tracking, robots frequently get delocalized and errors in human pose estimates may give rise to unsafe maneuvers that are challenging to handle. Finally, there are several robot design challenges to be addressed, including decisions on robot kinematics/dynamics, degrees of freedom, and even anthropomorphism.

ACKNOWLEDGMENTS

I would like to thank all my co-authors and mentors without whom this research would not have been possible. This research has been supported by grants from: the Honda Research Institute USA; the National Science Foundation NRI (#2132848), CHS (#2007011), IIS-1526035, and IIS-1563705; DARPA RACER (#HR0011-21-C-0171); the Office of Naval Research (#N00014-17-1-2617-P00004 and #2022-016-01 UW); Amazon.

CONFLICT OF INTEREST STATEMENT

The author declares that there is no conflict.

ORCID

Christoforos Mavrogiannis https://orcid.org/0000-0003-4476-1920

REFERENCES

Bartneck, C., D. Kulić, E. A. Croft, and S. Zoghbi. 2009. "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots." *International Journal of Social Robotics* 1: 71–81.

Beer, J. M., A. Prakash, T. L. Mitzner, and W. A. Rogers. 2011. "Understanding Robot Acceptance." Technical report, Georgia Institute of Technology.

Bennett, B. 2021. "Why Your Roomba Takes a Weird Path to Keep Your Floors Clean." https://www.cnet.com/home/kitchen-and-household/this-is-why-your-roombas-random-patterns-actually-make-perfect-sense/.

Berger, M. A. 2001. "Topological Invariants in Braid Theory." *Letters in Mathematical Physics* 55(3): 181–92.

Berkenkamp, F., M. Turchetta, A. Schoellig, and A. Krause. 2017. "Safe Model-Based Reinforcement Learning With Stability Guarantees." In *Advances in Neural Information Processing Systems*, vol. 30.

Birman, J. S. 1975. *Braids Links And Mapping Class Groups*. Princeton, NJ: Princeton University Press.

Brooks, R. 1986. "A Robust Layered Control System for a Mobile Robot." *IEEE Journal on Robotics and Automation* 2(1): 14–23.

Cadena, C., L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard. 2016. "Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age." *IEEE Transactions on Robotics* 32(6): 1309–32.

Carton, D., W. Olszowy, and D. Wollherr. 2016. "Measuring the Effectiveness of Readability for Mobile Robot Locomotion." *International Journal of Social Robotics* 8(5): 721–41. ISSN 1875-4805.

Chi, T.-C., M. Eric, S. Kim, M. Shen, and D. Hakkani-Tür. 2020. "Just Ask: An Interactive Learning Framework for Vision and Language Navigation." In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Csibra, G., and G. Gergely. 2007. "'Obsessed With Goals': Functions and Mechanisms of Teleological Interpretation of Actions in Humans." *Acta Psychologica* 124(1): 60–78.

Davis, F. D. 1989. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology." *MIS Quarterly* 13(3): 319–40.

Dragan, A., and S. Srinivasa. 2014. "Integrating Human Observer Inferences Into Robot Motion Planning." *Autonomous Robots: Special Issue on RSS '13* 37(4): 351–68.

Gupta, A., J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. 2018. "Social GAN: Socially Acceptable Trajectories With Generative Adversarial Networks." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2255–64.

Honda. 2019. "Honda P.A.T.H. Bot." https://global.honda/innovation/CES/2019/path_bot.html.

Kessler Faulkner, T. A., E. Schaertl Short, and A. L. Thomaz. 2020. "Interactive Reinforcement Learning With Inaccurate Feedback." In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 7498–504.

Kinzer, K. 2009. "Tweenbots." http://www.tweenbots.com/.

Kirby, R. 2010. *Social Robot Navigation*. PhD thesis, Carnegie Mellon University.

Knight, H., and R. Simmons. 2016. "Laban Head-Motions Convey Robot State: A Call for Robot Body Language." In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2881–88.

Koffka, K. 1935. *Principles of Gestalt Psychology*. Harcourt, Brace: New York.

Kretzschmar, H., M. Spies, C. Sprunk, and W. Burgard. 2016. "Socially Compliant Mobile Robot Navigation via Inverse Reinforcement Learning." *The International Journal of Robotics Research* 35(11): 1289–307.

Kwon, M., S. H. Huang, and A. D. Dragan. 2018. "Expressing Robot Incapability." In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 87–95.

Lasota, P. A., T. Fong, and J. A. Shah. 2017. "A Survey of Methods for Safe Human-Robot Interaction." *Foundations and Trends® in Robotics* 5(4): 261–349.

Mavrogiannis, C., and R. A. Knepper. 2021. "Hamiltonian Coordination Primitives for Decentralized Multiagent Navigation." *The International Journal of Robotics Research* 40(10-11): 1234–54.

Mavrogiannis, C., P. Alves-Oliveira, W. Thomason, and R. A. Knepper. 2022. "Social Momentum: Design and Evaluation of a Framework for Socially Competent Robot Navigation." *ACM Transactions on Human-Robot Interaction* 11(2): 1–37.

Mavrogiannis, C., K. Balasubramanian, S. Poddar, A. Gandra, and S. S. Srinivasa. 2023. "Winding Through: Crowd Navigation via Topological Invariance." *IEEE Robotics and Automation Letters (RA-L)* 8(1): 121–28.

Mavrogiannis, C., F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh. 2023. "Core Challenges of Social Robot Navigation: A Survey." *ACM Transactions on Human-Robot Interaction* 12: 1–39.

Mavrogiannis, C., J. A. DeCastro, and S. S. Srinivasa. 2023. "Abstracting Road Traffic via Topological Braids: Applications to Traffic Flow Analysis and Distributed Control." *The International Journal of Robotics Research*. https://doi.org/10.1177/02783649231188740.

Mavrogiannis, C. I., and R. A. Knepper. 2019. "Multi-Agent Path Topology in Support of Socially Competent Navigation Planning." *The International Journal of Robotics Research* 38(2-3): 338–56.

Nanavati, A., C. Mavrogiannis, K. Weatherwax, L. Takayama, M. Cakmak, and S. S. Srinivasa. 2021. "Modeling Human Helpfulness With Individual and Contextual Factors for Robot Planning." In *Proceedings of Robotics: Science and Systems (R:SS)*.

Nanavati, A., N. Walker, L. Taber, C. Mavrogiannis, L. Takayama, M. Cakmak, and S. Srinivasa. 2022. "Not All Who Wander are Lost: A Localization-Free System for in-the-Wild Mobile Robot Deployments." In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 422–31.

Poddar, S., C. Mavrogiannis, and S. S. Srinivasa. 2023. "From Crowd Motion Prediction to Robot Navigation in Crowds." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6765–72.

Roh, J., C. Mavrogiannis, R. Madan, D. Fox, and S. S. Srinivasa. 2020. "Multimodal Trajectory Prediction via Topological Invariance for Navigation at Uncontrolled Intersections." In *Proceedings of the Conference on Robot Learning*, 2216–27.

Rosenthal, S., J. Biswas, and M. Veloso. 2010. "An Effective Personal Mobile Robot Agent Through Symbiotic Human-Robot Interaction." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 915–22.

Rosenthal, S., M. Veloso, and A. K. Dey. 2011. "Learning Accuracy and Availability of Humans who Help Mobile robots." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1501–6.

Rudenko, A., L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. 2020. "Human Motion Trajectory Prediction: A Survey." *The International Journal of Robotics Research* 39(8): 895–935.

Schöller, C., V. Aravantinos, F. Lay, and A. Knoll. 2020. "What the Constant Velocity Model Can Teach Us About Pedestrian." *IEEE Robotics and Automation Letters* 5(2): 1696–703.

Stratton, A., K. Hauser, and C. Mavrogiannis. 2024. "Characterizing the Complexity of Social Robot Navigation Scenarios." *IEEE Robotics and Automation Letters*: in press.

Sung, J.-Y., L. Guo, R. E. Grinter, and H. I. Christensen. 2007. ""My Roomba is Rambo": Intimate Home Appliances." In *UbiComp 2007: Ubiquitous Computing*, edited by J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang, 145–62. Berlin, Heidelberg: Springer Berlin Heidelberg.

Thiffeault, J.-L. 2010. "Braids of Entangled Particle Trajectories." *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20(1): 017516.

Thomason, J., M. Murray, M. Cakmak, and L. Zettlemoyer. 2019. "Vision-and-Dialog Navigation." In *Proceedings of the Conference on Robot Learning (CoRL)*.

Thomaz, A., G. Hoffman, and M. Cakmak. 2016. "Computational Human-Robot Interaction." *Foundations and Trends®in Robotics* 4(2-3): 105–223. ISSN 1935-8253.

Thrun, S. 1994. "A Lifelong Learning Perspective for Mobile Robot Control." In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'94)*, vol. 1, 23–30.

Trautman, P., J. Ma, R. M. Murray, and A. Krause. 2015. "Robot Navigation in Dense Human Crowds: Statistical Models and Experimental Studies of Human-Robot Cooperation." *International Journal of Robotics Research* 34(3): 335–56.

Walker, N., C. Mavrogiannis, S. S. Srinivasa, and M. Cakmak. 2021. "Influencing Behavioral Attributions to Robot Motion During Task Execution." In *Proceedings of the Conference on Robot Learning (CoRL)*.

Wang, A., and A. Steinfeld. 2020. "Group Split and Merge Prediction With 3D Convolutional Networks." *IEEE Robotics and Automation Letters* 5(2): 1923–30.

Wang, A., C. Mavrogiannis, and A. Steinfeld. 2022. "Group-Based Motion Prediction for Navigation in Crowded Environments." In *Proceedings of the Conference on Robot Learning (CoRL)*, vol. 164, 871–82.

Warren, W. H. 2006. "The Dynamics of Perception and Action." *Psychological Review* 113(2): 358–89.

Weiss, A., J. Igelsböck, M. Tscheligi, A. Bauer, K. Kühnlenz, D. Wollherr, and M. Buss. 2010. "Robots Asking for Directions – the Willingness of Passers-by to Support Robots." In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 23–30.

Wise, M. 2018. "The 5 Stages of Acceptance as Robots Enter the Workforce." *World Economic Forum*, October 25, 2018. https://www.weforum.org/agenda/2018/10/robots-are-coming-to-your-workplace-here-s-how-to-get-along-with-them/.

Wolfinger, N. H. 1995. "Passing Moments: Some Social Dynamics of Pedestrian Interaction." *Journal of Contemporary Ethnography* 24(3): 323–40.

Ziebart, B. D., N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. Bagnell, M. Hebert, A. K. Dey, and S. Srinivasa. 2009. "Planning-Based Prediction for Pedestrians." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3931–36.

How to cite this article: Mavrogiannis, C. 2024. "Towards smooth mobile robot deployments in dynamic human environments." *AI Magazine* 45: 419–28. https://doi.org/10.1002/aaai.12192

AUTHOR BIOGRAPHY

Christoforos Mavrogiannis is an Assistant Professor of Robotics at the University of Michigan, working at the intersection of human–robot teaming and multi-agent systems. His goal is to enable robots to fluently work with and around people in unstructured, dynamic environments. To this end, he develops systems and algorithms for human–robot coordination and puts them to the test via extensive studies with users. He has been recognized as an outstanding young scientist by the Heidelberg Laureate Forum, a best-paper finalist at the International Conference on Human–Robot Interaction (HRI), and a Pioneer at the HRI and RSS (Robotics: Science and Systems) conferences. He has been a Hackaday Prize finalist and a winner of the Robotdalen International Innovation Award for his open-source initiative OpenBionics, and currently leads MuSHR, an open-source project aspiring to democratize robotics through the development of an affordable, highly functional, small-scale robot racecar. Prior to Michigan, Christoforos was a Postdoctoral Research Associate at the University of Washington. He holds a Ph.D. from Cornell University, and a Diploma in mechanical engineering from the National Technical University of Athens.