Experimental Insights from Developing Mobile Robots for Long-term Indoor Deployments*

Christoforos Mavrogiannis[1,2] and Siddhartha Srinivasa[2]
[1]Department of Robotics, University of Michigan
[2]Paul G. Allen School of Computer Science & Engineering, University of Washington

ABSTRACT

The ability to continually learn is crucial for robots designed to work with and around people. But as robots execute their tasks while improving through interactions with users, they will likely exhibit a wide range of behavior in response to novel situations. The way robots handle such situations has significant implications: robot behaviors may endanger or discomfort humans, robot task performance may drop, and robots may even experience catastrophic failures. Besides these immediate implications, such situations have longer-term impact on users whose trust to and expectations from robots are reshaped, resulting in overcautious or unnatural future human reactions to future robot behaviors and eventually hindering robot acceptance. In this paper, we discuss insights and lessons learned from the process of developing systems to support smooth, long-term deployments of mobile robots in real-world environments as they continually improve from interactions with users.

1 INTRODUCTION

The long-term deployment of robots in human environments gives rise to repeated and prolonged interactions with users that influence and shape robot acceptance. There is a long discussion on understanding the mechanisms underlying technology acceptance [2, 8, 36]. One of the most widely referenced models of acceptance is by Davis [8] who highlighted that the dominant predictors for the acceptance of any technology are its *perceived usefulness* and its *perceived ease of use*. From a robotics point of view, Beer et al. [2] identified attributes such as robot function, social capability, and appearance as crucial for the acceptance of a robot system. For robots entering the workplace alongside human co-workers, Wise [36] described the path towards robot acceptance as a series of stages, starting from fear, and continuing with apprehension, curiosity, tolerance, and finally satisfaction.

Motivated by these observations, we discuss models and algorithms that may keep robots towards the path to satisfaction and widespread acceptance. Specifically, we look at elements of autonomous mobility, mechanisms for failure recovery, and models for user perceptions.

Specifically for mobile robots, the path towards satisfaction requires a series of considerations that will ensure high functionality that is perceived and appreciated by the user without extensive effort from the user's part. These considerations include but are not limited to human safety and comfort, the robustness of the autonomy, and the implications of the autonomously generated behavior for users and bystanders.

*This paper presents a summary of work in collaboration with P. Alves-Oliveira, Wil Thomason, Ross Knepper, Krishna Balasubramanian, Sriyash Poddar, Anush Gandra, Siddhartha Srinivasa, Amal Nanavati, Kevin Weatherwax, Leila Takayama, Maya Cakmak, Nick Walker, and Lee Taber.

Figure 1: Honda's experimental ballbot [11] navigates next to three users in our lab experiment [20].

2 ROBUST MOBILITY DOES NOT ALWAYS REQUIRE COMPLEX MODELS

Many important real-world robotics applications involve the deployment of mobile robots in dynamic and unstructured human environments like homes, warehouses and hospitals. One influential system paradigm involves robots adapting their navigation strategies online through lifelong interactions with their environment [24, 32]. However, to keep users safe and comfortable as robots learn, they need a series of core capabilities, including mechanisms for localization, human motion prediction, collision avoidance, and human comfort among others. There is a long history of work across these domains, with recent advances involving the use of complex, large-scale, deep learning architectures demonstrating great performance on familiar instances [29] but also poor generalization and limited interpretability on others. In contrast, our experimental insight is that simple, domain-driven models can provide a foundation of performance that can be further finetuned online through repeated interactions with users and the environment.

2.1 Crowd Navigation with Constant Velocity Prediction

Social order in pedestrian navigation is often the result of cooperation: humans tend to share responsibility for collision avoidance and follow contracts for repair when things go wrong [37]. Roboticists have embraced this viewpoint, developing algorithms that support cooperative decision making in densely crowded human environments [16, 19, 33, 38].

In our recent work [20], we described a formalism of *passing* between two agents as a pairwise winding number. The absolute

(a)

(b)

Figure 2: Our Kuri robot wandered a 28,000ft^2 floor for 4 days with minimal human help [23]. (a) Kuri's coverage of the building floor. (b) Photos of Kuri as it wandered in the environment.

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. We integrated this function within a model predictive controller (MPC) which additionally incorporates considerations of agents' personal space and robot's task efficiency. Using a simple constant-velocity (CV) model as human motion prediction, we deployed this MPC on a self-balancing robot and extensively tested it in the lab under challenging crowd conditions (see Fig. 1). Our experiments demonstrated that our approach significantly outperformed a recent deep reinforcement-learning based baseline in terms of *Safety* and *Efficiency*. Further, we found that CV prediction [26] performs comparably to a recent state-of-the-art motion prediction baseline (S-GAN) [10] across a range of crowd behaviors including aggressive or inattentive agents.

While our 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.

2.2 Localization-free Field Deployment

One of the challenges preventing the smooth prolonged deployment of mobile robots is the need for accurate localization. Despite the important advances in simultaneous localization and mapping (SLAM) [6] 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 [3, 5, 13]. By *wandering* in space, such systems are capable of completing a wide range of tasks, especially coverage-based, such as vacuum cleaning [3] or patrolling. Inspired by their effectiveness, we developed a wandering system which we deployed on a Kuri robot [23]. 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 until it gets stuck due to some obstruction at which point it updates its costmap to trigger the selection of a new direction. If the robot gets stuck for a prolonged period, it initiates a recovery procedure involving rotation in place and backing off. These behaviors were parametrized based on our environment and 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 four days while requiring minimal human help.

While alternative tasks involving point-to-point navigation would require a localization system in place, our system can empower robots with 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.

3 BYSTANDERS CAN ENABLE SCALABLE ROBOT RECOVERY

When autonomy inevitably fails, human help can be crucial for robot recovery [7, 35]. Typically, researchers and engineers are responsible for ensuring continued robot operation during studies and field deployments. However, this paradigm may not always be a 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 [27, 28, 31, 35] but also motivated by our experience deploying Kuri [23] in our academic building for four days, as part of a user study (see Sec. 2.2). In the following paragraphs we elaborate on some of the nuances of leveraging bystander help.

3.1 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 28,000ft^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) [23]. 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 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 Sec. 2.2– enabled the robot to achieve a substantial coverage of the floor (see Fig. 2a). Overall, the researchers did not spend more than 30' helping the robot over the four full days (32 hours) of the study, illustrating the practicality of enabling continued robot operation through periodical, non-technical human help.

3.2 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 [22]. Specifically, we considered a delivery robot that 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 the 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 Fig. 3). 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.

4 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 [30]. Robots driven by purely functional objectives may complete their tasks but in 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 [9], the robot's inability to complete its task [17], or aspects of style and attitude [15]. However, as the robot learns 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.

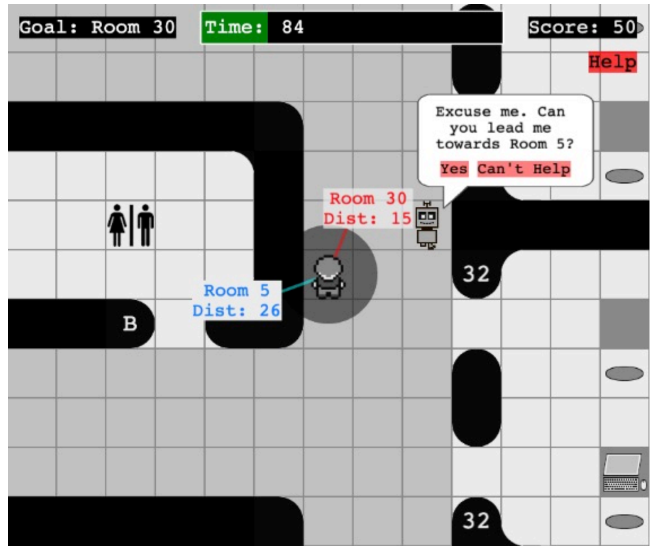

Figure 3: A user performs tasks in our virtual office environment while a robot periodically asks them for help. In this environment, we deployed and tested our system for planning effective human help requests [22].

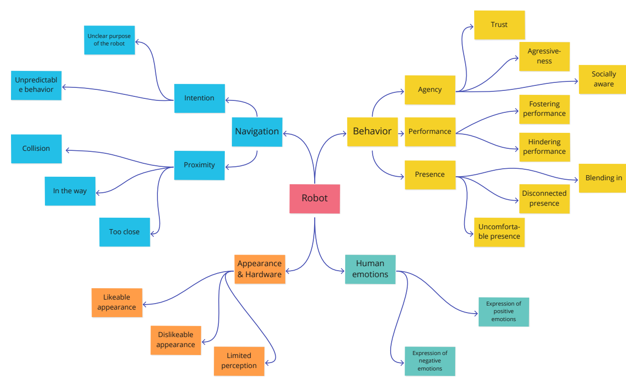

Figure 4: Coding scheme used to analyze users' impressions during robot navigation experiments in the lab [19].

In the following paragraphs, we discuss how different navigation algorithms elicit different user attributions, and how we can extract context-specific attribution models to automatically synthesize robot behaviors that elicit desired user impressions.

4.1 User Impressions of Different Crowd Navigation Strategies

To study human impressions of different robot navigation strategies, we developed a fictional factory setting mockup in the lab, where three users navigated between a set of machines to perform maintenance tasks while one robot was moving around inspecting their work [19]. 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

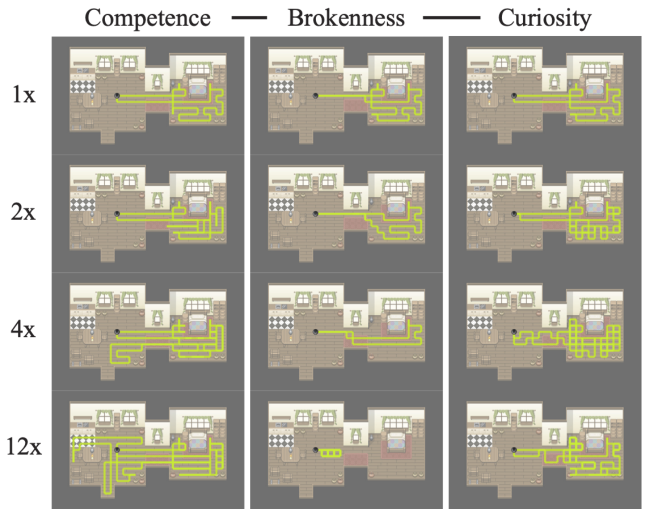

Figure 5: Robot trajectories that communicate desired attributions generated by our trajectory optimization framework [34]. Columns indicate the type of attribution and rows indicate the intensity of the attribution.

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 [19], an algorithm designed to generate legible motion in multiagent domains) enabled users to navigate with lower accelerations next to our robot. Interestingly, 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 Fig. 4) is indicative of the range of human impressions when interacting closely with mobile robots.

4.2 Shaping Users' Impressions of Navigating Robots

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 [34].

Considering a coverage navigation task (vacuum-cleaning robot) in a virtual home environment, we asked users to rate a wide range of robot behaviors [1]. 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 evaluation user study, we demonstrated that our framework was able to autonomously generate robot motion eliciting desired attributions of desired intensity from users (see Fig. 5).

5 DISCUSSION

We discussed experimental insights extracted from our experience developing models and algorithms meant to support long-term deployments and lifelong learning [32] of interactive mobile robots in indoor environments. These insights may be useful for guiding users through a path towards satisfaction [36] and acceptance [2, 8] as robots continue to improve through interactions with their environment. 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 as the robot learns. There is an extensive body of work on approaches that directly address aspects of safety in human-robot interaction [18] and relevant work in safe reinforcement learning [4] 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 [12].

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 [14], 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 [25], additional research is required to develop realistic simulators, evaluation criteria, and benchmark experiments design of realistic simulators but also the definition of benchmark experiments [21].

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

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.

REFERENCES

[1] Christoph Bartneck, Dana Kulić, Elizabeth A. Croft, and Susana Zoghbi. 2009. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* 1 (2009), 71–81.

[2] Jenay M Beer, Akanksha Prakash, Tracy L Mitzner, and Wendy A Rogers. 2011. *Understanding robot acceptance.* Technical Report. Georgia Institute of Technology.

[3] Brian Bennett. 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/

[4] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. 2017. Safe Model-based Reinforcement Learning with Stability Guarantees. In *Advances in Neural Information Processing Systems*, Vol. 30.

[5] Rodney Brooks. 1986. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation* 2, 1 (1986), 14–23.

[6] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J. Leonard. 2016. Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age. *IEEE Transactions on Robotics* 32, 6 (2016), 1309–1332.

[7] Ta-Chung Chi, Mihail Eric, Seokhwan Kim, Minmin Shen, and Dilek Hakkani-Tür. 2020. Just Ask: An Interactive Learning Framework for Vision and Language Navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence.*

[8] Fred D. Davis. 1989. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly* 13, 3 (1989), 319–340. http://www.jstor.org/stable/249008

[9] Anca Dragan and Siddhartha Srinivasa. 2014. Integrating Human Observer Inferences into Robot Motion Planning. *Autonomous Robots: Special Issue on RSS '13* 37, 4 (2014), 351 – 368.

[10] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre 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).* 2255–2264.

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

[12] Taylor A. Kessler Faulkner, Elaine Schaertl Short, and Andrea L. Thomaz. 2020. Interactive Reinforcement Learning with Inaccurate Feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA).* 7498–7504.

[13] Kacie Kinzer. 2009. Tweenbots. http://www.tweenbots.com/

[14] Rachel Kirby. 2010. *Social Robot Navigation.* Ph. D. Dissertation. Carnegie Mellon University.

[15] Heather Knight and Reid 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).* 2881–2888.

[16] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. 2016. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research* 35, 11 (2016), 1289–1307.

[17] Minae Kwon, Sandy H. Huang, and Anca D. Dragan. 2018. Expressing Robot Incapability. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction.* 87–95.

[18] Przemyslaw A. Lasota, Terrence Fong, and Julie A. Shah. 2017. A Survey of Methods for Safe Human-Robot Interaction. *Foundations and Trends® in Robotics* 5, 4 (2017), 261–349.

[19] Christoforos Mavrogiannis, Patrícia Alves-Oliveira, Wil Thomason, and Ross A. Knepper. 2022. Social Momentum: Design and Evaluation of a Framework for Socially Competent Robot Navigation. *ACM Transactions on Human-Robot Interaction* 11, 2 (2022), 37 pages.

[20] Christoforos Mavrogiannis, Krishna Balasubramanian, Sriyash Poddar, Anush Gandra, and Siddhartha S. Srinivasa. 2023. Winding Through: Crowd Navigation via Topological Invariance. *IEEE Robotics and Automation Letters (RA-L)* 8, 1 (2023), 121–128.

[21] Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. 2023. Core Challenges of Social Robot Navigation: A Survey. *ACM Transactions on Human-Robot Interaction* (2023).

[22] Amal Nanavati, Christoforos Mavrogiannis, Kevin Weatherwax, Leila Takayama, Maya Cakmak, and Siddhartha S Srinivasa. 2021. Modeling Human Helpfulness with Individual and Contextual Factors for Robot Planning. In *Proceedings of Robotics: Science and Systems (R:SS).*

[23] Amal Nanavati, Nick Walker, Lee Taber, Christoforos Mavrogiannis, Leila Takayama, Maya Cakmak, and Siddhartha 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).* 422–431.

[24] Rohan Paul and Paul Newman. 2013. Self-help: Seeking out perplexing images for ever improving topological mapping. *The International Journal of Robotics Research* 32, 14 (2013), 1742–1766.

[25] Sören Pirk, Edward Lee, Xuesu Xiao, Leila Takayama, Anthony Francis, and Alexander Toshev. 2022. A Protocol for Validating Social Navigation Policies. *arXiv e-prints* (2022). arXiv:2204.05443 [cs.RO]

[26] Sriyash Poddar, Christoforos Mavrogiannis, and Siddhartha S. Srinivasa. 2023. From Crowd Motion Prediction to Robot Navigation in Crowds. *arXiv e-prints*, Article arXiv:2303.01424 (March 2023), arXiv:2303.01424 pages. https://doi.org/10.48550/arXiv.2303.01424 arXiv:2303.01424 [cs.RO]

[27] Stephanie Rosenthal, Joydeep Biswas, and Manuela 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).* 915–922.

[28] Stephanie Rosenthal, Manuela Veloso, and Anind K. Dey. 2011. Learning Accuracy and Availability of Humans Who Help Mobile Robots. In *Proceedings of the AAAI Conference on Artificial Intelligence.* 1501–1506.

[29] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. 2020. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research* 39, 8 (2020), 895–935.

[30] Ja-Young Sung, Lan Guo, Rebecca E. Grinter, and Henrik I. Christensen. 2007. "My Roomba Is Rambo": Intimate Home Appliances. In *UbiComp 2007: Ubiquitous Computing*, John Krumm, Gregory D. Abowd, Aruna Seneviratne, and Thomas Strang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 145–162.

[31] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. 2019. Vision-and-Dialog Navigation. In *Proceedings of the Conference on Robot Learning (CoRL).*

[32] S. Thrun. 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 vol.1.

[33] Peter Trautman, Jeremy Ma, Richard M. Murray, and Andreas Krause. 2015. Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation. *International Journal of Robotics Research* 34, 3 (2015), 335–356.

[34] Nick Walker, Christoforos Mavrogiannis, Siddhartha S. Srinivasa, and Maya Cakmak. 2021. Influencing Behavioral Attributions to Robot Motion During Task Execution. In *Proceedings of the Conference on Robot Learning (CoRL).*

[35] Astrid Weiss, Judith Igelsböck, Manfred Tscheligi, Andrea Bauer, Kolja Kühnlenz, Dirk Wollherr, and Martin 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.

[36] Melonee Wise. 2018. *The 5 stages of acceptance as robots enter the workforce.*

[37] Nicholas H. Wolfinger. 1995. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography* 24, 3 (1995), 323–340.

[38] Brian D. Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J. Andrew Bagnell, Martial Hebert, Anind K. Dey, and Siddhartha Srinivasa. 2009. Planning-based Prediction for Pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).* 3931–3936.