

Fluent Coordination in Proximate Human Robot Teaming

Sachiko Matsumoto and Laurel D. Riek
University of California, San Diego

Abstract—Fluent coordination is important in order for teams to work well together. In proximate teaming scenarios, fluent teams tend to perform more successfully. Recent work suggests robots can support fluency in human-robot teams a number of ways, including using nonverbal cues and anticipating human intention. However, this area of research is still in its early stages. We identify some of the key challenges in this research space, specifically individual variations during teaming, knowledge and task transfer, co-training prior to task execution, and long-term interactions. We then discuss possible paths forward, including leveraging human adaptability, to promote more fluent teaming.

I. INTRODUCTION

As robots venture into more advanced tasks in human-centered environments, it becomes increasingly imperative for them to proximately team with people to complete tasks. Hoffman [19] defines fluency as “a well-synchronized meshing of [agents’] actions,” characterized by agents who coordinate with each others’ actions with “precise and efficient” timing. For example, a fluent interaction might involve two people coordinating to move boxes or hang a banner (c.f. Fig. 1).

Entin and Serfaty [13] showed that well-coordinated teamwork is linked to significantly improved performance in human-human teams. Additionally, fluent, well-coordinated teaming enables human-robot teams to complete tasks more successfully, and may make people more willing to work with robots [15, 23, 25]. Thus, in order for human-robot teams to be effective, it is crucial that robot teammates behave in a way that preserves and promotes fluency.

Many of the environments in which robots could be especially useful are dynamic and unstructured. For instance, robots could be extremely helpful in healthcare, both in hospitals and home settings [43]. However, for robots to enter these environments, and for them to interact with people for long periods of time, their control systems and teaming strategies must be robust. Different people will do the same task differently, and even an individual will often change how they perform a task over time. Thus, when teaming with people, robots must be able to adapt to these variations and support fluent teaming regardless of people’s individual characteristics.

In spite of the importance of fluency in teaming, research on this topic is still in its early stages within the proximate human-robot teaming (HRT) community. In fact, there are currently no widely agreed upon standard measures for fluency [19]. Although some measures are commonly used in the teaming literature, these were only systematically evaluated this year [19]. Additionally, while such measures work well for certain teaming tasks (e.g., shared workspace tasks), they are not all



Fig. 1. Examples of robots engaging in proximate HRT from our prior work (clockwise from top left): collaboratively hanging a banner [51], engaging in rehabilitation exercises [52], implicitly learning from a human [18], and stacking boxes [32]. These platforms act in shared spaces with people, and need to be able to team fluently with their human partners for successful interactions.

as well-suited for others, such as co-manipulation or tasks with more than two agents. Standard fluency measures that can be used across a variety of tasks will aid the HRT community in addressing challenges in teaming.

In this paper, we present some of the current major challenges to achieving robust, fluent HRT. Specifically, we explore human variance in teaming scenarios, knowledge and task transfer, co-training prior to execution, and long-term interactions. We discuss the factors that make these aspects of teaming difficult, and also consider some of the work that investigates solutions to these problems. We then suggest next steps for addressing some of these challenges, such as capitalizing on people’s ability to adapt to robotic partners and expanding current methods to account for unstructured behavior.

II. CHALLENGES IN FLUENT TEAMING

A. Human Variance

Fluent coordination in proximate HRT faces a prominent challenge: the large amount of variance within the human population, which makes developing a general model for human behavior impossible. This is an especially relevant issue in fluent coordination, as the robot’s movements must be closely coordinated with the human’s. Thus, in the absence

of a comprehensive prior, robots must be able to quickly and reliably learn a viable model of an individual partner's behavior. In particular, robust strategies to detect and respond to human intention, role changes, and levels of expertise would enhance robots' abilities to understand and coordinate with humans. However, enabling robots to adapt to these variations is challenging, as it requires detailed sensing and accurate interpretation of people's behaviors, which are open problems in robotics.

Human activity recognition could address one facet of fluent coordination. A wide range of algorithms and sensors are used to classify people's activities [9]. Such classifications can be particularly useful for robots that collaborate with people on tasks that have a clear structure, such as on manufacturing assembly lines or cooking using a recipe. For instance, if a robot detects what step a person on an assembly line is currently performing, it could fetch the component that will subsequently be needed in advance. However, this is still an active area of research, and the classification of fine-grained motions in particular can be difficult [30].

In addition to detection of human activity, many teaming scenarios will require robots to sense and interpret subtle, visual communication cues. This is a challenging area, particularly given that there is not always a simple mapping between nonverbal signals and meaning, and it can be difficult to even distinguish when a gesture begins and ends [27]. However, if robots can glean some meaning from these cues, they will likely be able to predict human intentions more accurately and subsequently generate plans that better support fluent teaming. For instance, Duarte et al. [12] showed that people can predict where someone will place an object with over 80% accuracy based only on saccadic eye movement and head rotation. Huang and Mutlu [22] used gaze cues to predict human intent and demonstrated that a robot anticipating these cues collaborated more efficiently with people than a reactive robot did. Additionally, several studies indicate that robots can also display nonverbal cues to influence human behavior and make interaction more fluent [44, 11, 6, 8, 35, 12, 36, 1].

Another important area challenged by human variation is understanding agents' roles throughout an interaction. There are several types of roles that can arise during an interaction, including supervisor, worker, and peer [46]. Much of the work on fluent teaming to date either does not explicitly consider roles or assumes static roles, with Peternel et al. [40] and Rezvani et al. [41] being notable exceptions. However, in real-world scenarios, people's roles often shift over the course of an interaction. For instance, when carrying a table, the person facing forward may start off in a leader role, but if that person ends up facing backwards after turning a corner, they might switch to a follower role. Furthermore, in group tasks, individuals may join and leave the group at random intervals, and at such points people's roles will necessarily change [48].

Consequently, robots will need to be able to shift their own role in response to new situations. Peternel et al. [40] explored this by developing an algorithm for a robot to actively change roles by having the robot sense when a person became fatigued

and adjust its role to reduce the force output required from the human. Rezvani et al. [41] also proposed autonomous driving situations in which a robot's role might change and investigated how the robot can best initiate this shift with a person. However, in both of these studies, the role changes occurred in very specific, well-defined cases, which do not reflect the varied scenarios robots will encounter in the real world. Therefore, more work needs to be done to ensure that robots can appropriately initiate and respond to role changes in dynamic environments.

Robots must also be able to recognize a teammate's level of comfort or expertise with a task in order to assume an appropriate role. As they integrate more into human-centered environments, robots will encounter people with varying amounts of knowledge about a given task, which will affect their ability to perform the task. For instance, in a handover task, novice users were less likely to give physical feedback than experienced users [33]. This sometimes caused the robot to keep hold of the object after the person expected a release, leading to failed handovers. Therefore, robots must be able to distinguish between differing levels of expertise and adjust their behavior accordingly. This is important both when performing the same task with different people and when performing different tasks with the same person. If their partner is a novice, the robot may have to provide more assistance or guidance, potentially teaching the person about the task. On the other hand, if the person is very experienced, the robot should leverage their knowledge by asking the person for assistance or allowing them to lead. By taking their partner's abilities into account, robots will be able to achieve greater fluency during the interaction.

B. Knowledge and Task Transfer

As robots enter human-centered environments, they will need to work with multiple people and perform a variety of tasks. In order to perform reliably, they will need to be able to generalize information from teaming experiences to better understand how to navigate new tasks and new partners.

Even for distinct tasks, there are likely some underlying characteristics or patterns that can be used to transfer knowledge from one task to another. For instance, Iqbal and Riek [23, 24] demonstrated that the same group synchrony metric can be used across a variety of scenarios, including playing a cooperative game, marching, and dancing. Additionally, Thomaz and Chao [49] suggested that while turn-taking behaviors are context-dependent, there are still some generic components. They proposed a framework constructed from a domain-specific finite state machine with a generic Markov decision process that models turn-taking. Recent studies also showed that robot grasps can generalize to similar objects [2], and teaching techniques can generalize across topics [4]. Furthermore, Gutierrez et al. [16] investigated how robots can efficiently perform new tasks by learning modifications to similar tasks they can already achieve.

Similarly, robots can leverage characteristics of humans that are largely constant over the population to more quickly adapt

to new partners. Past research shows that certain movement dynamics are somewhat invariant across different people (e.g., the minimum jerk model [21], the 2/3 power law [31], etc.) and that people will often act similarly given the same affordances in their environment [26]. Robots can capitalize on these tendencies to support accurately predicting human behavior more quickly, which will promote fluent teaming with users they have not encountered before.

Conversely, robots must also respect people's differences and adapt their policies accordingly. There is no "one size fits all" model for people, and individual differences play a key role in how we understand teaming [17, 18, 44]. General trends and transferring learned knowledge are a good prior when a robot first interacts with someone, but it needs to then adapt to the idiosyncrasies specific to that person. For instance, even if a robot uses common behaviors associated with given affordances as a starting point for interaction, the affordances of an environment vary for different people (e.g., a large step affords stepping for an adult but not a small child) [26].

Robots must also adapt as people change over time. For example, if people have a neurodegenerative condition, such as dementia, their ability to perform certain tasks may change across a range of time scales, throughout the day or across many years [43, 34]. These changes may occur at different rates that are specific to each person. Thus, robots must be able to personalize their behavior to people as they interact with them [52].

C. Co-Training Prior to Task Execution

Many current teaming algorithms require the human and robot to practice together for several iterations before the human-robot team can perform effectively enough for real-world deployment [37, 20]. There are many contexts in which extensive repetition is necessary for safe or effective performance. For instance, in manufacturing settings, where high precision may be necessary when working on a delicate product, or in hospitals, where mistakes can be life-threatening, it is appropriate that the human and robot train together extensively to ensure safe and reliable outcomes. Multiple training iterations enable both the human and robot to more fully understand each other, and ultimately allow the human-robot team to optimize their performance.

However, such training is neither appropriate nor feasible in all settings. For example, in many environments a robot will not have a consistent team, but rather be expected to spontaneously work with new teams or individuals. In these cases, having an extended training period would disrupt the interaction and potentially make the robot a burden to work with rather than an aid. Furthermore, it is impossible to predict all scenarios a human-robot team will encounter in unstructured environments, making the ability to learn on-the-fly a necessity.

For the robot to act appropriately despite a short or non-existent training period means that it must be able to either quickly learn about or have generalizable knowledge of human preferences and tendencies and have a prior for the task. Such

skills will enable robots to promote fluent coordination in teams it has not worked with before. As discussed above, being able to transfer knowledge across tasks and users will likely be helpful in reducing training times.

Koppula et al. [29] approached this problem by training a policy offline before interacting with a person. However, their model did not adapt to the person while interacting with them, though they note that it could be changed to learn online.

Limiting training time is not only detrimental to the robot's efficacy, but the person's as well. When a person trains with a robot, they learn about the robot's limitations and develop a better mental model of the robot. Without this prior experience, people may have unrealistic expectations when they begin working with the robot [45, 39, 7]. Therefore, the robot must be transparent and clearly convey its capabilities to the person in order for them to perform successfully.

D. Long-Term Interaction

Another difficult challenge HRT researchers face is maintaining reliable teaming over long periods of time [42, 10, 28, 14]. For instance, as robots integrate into a wider variety of environments, they will be expected to run for long periods without major faults and maintain long-term interactions with a variety of people. They will also likely encounter situations during teaming that they are unable to cope with on their own [47]. For example, a human may expect a robot to retrieve an item that is out of the robot's reach. In this case, the robot will need to identify that the task is outside of its limitations and react appropriately to mitigate the disruption to the task. This is an inherently challenging problem.

However, if robots have robust control strategies for teaming and are able to sense humans, learn from them, and transfer knowledge across tasks, long-term interaction becomes significantly easier. With these skills, robots could alert people to situations where failure is likely. This would result in fewer disruptions during teaming and avoids situations the robot cannot recover from. By transferring what they have learned across tasks, robots also would be able to effectively collaborate with humans on a wider variety of tasks, a necessary feature for long-term deployment in dynamic environments.

Additionally, if the robot encountered a task it was uncertain about, it could rely on its human partner's knowledge to help it complete the task. Cakmak et al. [5] demonstrated that actively learning robots (i.e., those that ask questions) learn more quickly than passively learning robots. Furthermore, people preferred to interact with robots that actively participated in learning. Thus, by actively curating knowledge, robots could become both better at their task and more engaging partners.

III. RECOMMENDED PATHS FORWARD

A. Bi-Directional Adaptation

The challenges presented above are by no means straightforward to solve. However, one approach to address some of the issues is to leverage human adaptability. Humans are extraordinarily adept at adapting to a wide variety of situations, including new partners or new tasks in teaming scenarios.

If robots can explicitly model and take advantage of the way humans adjust to them, they may, in turn, be able to coordinate their adaptations with the person to promote fluent, coordinated teaming. For instance, if a robot and human are carrying a table together and the robot detects that the person has adapted to the speed of its trajectory, the robot might keep its speed the same for the rest of the interaction so as not to confuse the person. However, it could still adapt the shape of its trajectory to better mesh with the person.

There has been some work to date looking into human adaptability during HRT. For instance, Amirshirzad et al. [3] investigated the effects of robotic behavior on human behavior in a ball-balancing task. They compared people's performance when using a purely teleoperated robot to that when collaborating with a shared-control robot that attempted to infer its partner's intention throughout execution. Although the teleoperated robot employed a simpler controller (e.g., it precisely followed the person's input), people naturally adapted to the shared-control robot, and quickly outperformed the teleoperation group. While the robot's control strategy was static in this experiment, these results suggest that human adaptability can be leveraged to improve team performance.

Thomaz et al. [50] also found that people adapt to robots over the course of an interaction. In a reinforcement learning scenario, in which people assigned rewards to the robot to help it learn, people adjusted the number of rewards they gave throughout the task. They also changed the type of reward they gave to better accommodate the robot's learning. These results again indicate that humans will adapt to robots to produce more effective teaming outcomes.

This issue was also partly addressed by Nikolaidis and Shah [37] with a cross-training algorithm, in which humans and robots repeatedly switched roles to learn how to better complete a task. This enabled both the human and the robot to adapt to their partner throughout training, and the authors showed that the "mental models" of both agents converged during training. However, this process requires a significant amount of training, and the robot did not explicitly consider how the human changed throughout the process. Nor did the robot subsequently adapt after the training stage if the human changed their policy during execution.

In contrast, Nikolaidis et al. [38] proposed the Bounded-Memory Adaptation Model (BAM), which explicitly models human adaptability and adjusts robot behavior based on that adaptability throughout the interaction. The model includes a latent variable that models the person's level of adaptability. If the person is not very adaptable, they are more likely to continue with their policy, regardless of the robot's behavior, so the robot must switch to the human's policy or "mode" to make progress on the task. On the other hand, if the human is more adaptable, they are likely to switch to the robot's mode. However, as the authors note, this model assumes a fully observable world-state, an assumption that often does not hold in real-world scenarios.

The above studies represent a good starting point for investigating how robots can take advantage of human adaptability.

However, HRI researchers are only beginning to understand the ways in which people adapt to robots. Additionally, the current models that incorporate human adaptability are somewhat limited in their applicability due to either their training requirements or world-state assumptions. Thus, more work needs to be done to investigate the dynamics of bi-directional adaptation, as well as how robots can best use people's adaptability to support fluent teaming.

B. Unstructured Behavior

While we try to make robots able to generalize knowledge to new tasks, we should also consider how we can generalize to tasks with even less structure. Many current techniques focus on periodic [24] or structured tasks, like constructing an object [20]. However, tasks that people regularly perform often do not have such a well-defined structure. As robots enter human-centered environments, they will be expected to collaborate on such tasks with people.

If robots can interpret human motion and transfer knowledge across tasks, it is possible that the switch from structured to unstructured tasks will not be such a difficult one. Most tasks will likely have some underlying structure, even if it is not as obvious as the structure of the tasks currently focused on. Therefore, it is possible that robots could still use some of the same information and techniques from structured tasks to promote fluent teaming in unstructured tasks. For instance, if a robot is collaboratively moving furniture with people from a moving truck to a house, it may not have complete information about the task. It might not know what piece of furniture they will pick up next or where they want to move it. However, it could still realize that there is a pattern to the movement, namely that the person repeatedly goes from the truck to the house and back, and it could use this information to coordinate with the person, perhaps building on the work of Iqbal and Riek [24] to better synchronize with the people involved. Furthermore, even if there is little structure in the task, there may still be parallels in the human dynamics. Thus, if the robot could robustly sense the person's motion, it may be able to predict approximately what the person will do next.

IV. DISCUSSION

In this paper, we discussed several of the major challenges in HRT, namely human variance in teaming, knowledge and task transfer, co-training prior to task execution, and long-term interactions. While these are all difficult challenges, work to date provides a good foundation to address them. Current methods could also potentially be adapted to support fluent teaming in unstructured environments. Furthermore, by utilizing human skills, such as adaptability, robots will be able to more readily engage in fluent coordination with human partners. With the realization of better models to support fluent coordination in teaming, robots will become more effective teammates for people, enabling them to augment people's capabilities in a variety of situations.

REFERENCES

- [1] H. Admoni and B. Scassellati. Social eye gaze in human-robot interaction: a review. *Journal of Human-Robot Interaction*, 6(1):25–63, 2017.
- [2] S. Alexandrova, Z. Tatlock, and M. Cakmak. Roboflow: A flow-based visual programming language for mobile manipulation tasks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5537–5544. IEEE, 2015.
- [3] N. Amirshirzad, A. Kumru, and E. Oztop. Human adaptation to human–robot shared control. *IEEE Transactions on Human-Machine Systems*, 49(2):126–136, April 2019. ISSN 2168-2291. doi: 10.1109/THMS.2018.2884719.
- [4] M. Cakmak and A. L. Thomaz. Eliciting good teaching from humans for machine learners. *Artificial Intelligence*, 217:198–215, 2014.
- [5] M. Cakmak, C. Chao, and A. L. Thomaz. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118, June 2010. ISSN 1943-0604. doi: 10.1109/TAMD.2010.2051030.
- [6] M. Cakmak, S. S. Srinivasa, M. K. Lee, S. Kiesler, and J. Forlizzi. Using spatial and temporal contrast for fluent robot-human hand-overs. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 489–496, March 2011. doi: 10.1145/1957656.1957823.
- [7] E. Cha, A. D. Dragan, and S. S. Srinivasa. Perceived robot capability. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 541–548. IEEE, 2015. URL https://ri.cmu.edu/pub/_files/2015/8/roman2015_{_}137.pdf.
- [8] C. Chao and A. Thomaz. Timed petri nets for fluent turn-taking over multimodal interaction resources in human-robot collaboration. *The International Journal of Robotics Research*, 35(11):1330–1353, 2016. doi: 10.1177/0278364915627291. URL <https://doi.org/10.1177/0278364915627291>.
- [9] M. Cornacchia, K. Ozcan, Y. Zheng, and S. Velipasalar. A survey on activity detection and classification using wearable sensors. *IEEE Sensors Journal*, 17(2):386–403, Jan 2017. doi: 10.1109/JSEN.2016.2628346.
- [10] K. Dautenhahn. Methodology & themes of human-robot interaction: A growing research field. *International Journal of Advanced Robotic Systems*, 4(1):15, 2007. doi: 10.5772/5702. URL <https://doi.org/10.5772/5702>.
- [11] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa. Effects of robot motion on human-robot collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI ’15*, pages 51–58, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696473. URL <http://doi.acm.org/10.1145/2696454.2696473>.
- [12] N. F. Duarte, M. Raković, J. Tasevski, M. I. Coco, A. Billard, and J. Santos-Victor. Action anticipation: Reading the intentions of humans and robots. *IEEE Robotics and Automation Letters*, 3(4):4132–4139, Oct 2018. doi: 10.1109/LRA.2018.2861569.
- [13] E. E. Entin and D. Serfaty. Adaptive team coordination. *Human Factors*, 41(2):312–325, 1999. doi: 10.1518/001872099779591196. URL <https://doi.org/10.1518/001872099779591196>.
- [14] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. C. Schultz, et al. Designing robots for long-term social interaction. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1338–1343. IEEE, 2005.
- [15] M. C. Gombolay, R. A. Gutierrez, S. G. Clarke, G. F. Sturla, and J. A. Shah. Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams. *Autonomous Robots*, 39(3):293–312, Oct 2015. doi: 10.1007/s10514-015-9457-9. URL <https://doi.org/10.1007/s10514-015-9457-9>.
- [16] R. A. Gutierrez, E. S. Short, S. Niekum, and A. L. Thomaz. Learning from corrective demonstrations. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 712–714, March 2019. doi: 10.1109/HRI.2019.8673287.
- [17] C. J. Hayes. *Building adaptable robots: How behavioral metrics can enable robots to implicitly learn from humans*. PhD thesis, 2015. URL <https://search.proquest.com/docview/1764221242?accountid=14524>.
- [18] C. J. Hayes, M. Moosaei, and L. D. Riek. Exploring implicit human responses to robot mistakes in a learning from demonstration task. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 246–252, Aug 2016. doi: 10.1109/ROMAN.2016.7745138.
- [19] G. Hoffman. Evaluating fluency in human–robot collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3):209–218, 2019. doi: 10.1109/THMS.2019.2904558.
- [20] G. Hoffman and C. Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of the ACM/IEEE International Conference on Human-robot Interaction, HRI ’07*, pages 1–8, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228718. URL <http://doi.acm.org/10.1145/1228716.1228718>.
- [21] N. Hogan. An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, 4(11):2745–2754, 1984. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.04-11-02745.1984. URL <http://www.jneurosci.org/content/4/11/2745>.
- [22] C.-M. Huang and B. Mutlu. Anticipatory robot control for efficient human-robot collaboration. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, HRI ’16*, pages 83–90, Piscataway, NJ, USA, 2016. IEEE Press. ISBN 978-1-4673-8370-7. URL <http://dl.acm.org/citation.cfm?id=2906831.2906846>.

- [23] T. Iqbal and L. D. Riek. A method for automatic detection of psychomotor entrainment. *IEEE Transactions on Affective Computing*, 7(1):3–16, Jan 2016. ISSN 1949-3045. doi: 10.1109/TAFFC.2015.2445335.
- [24] T. Iqbal and L. D. Riek. Coordination dynamics in multi-human multirobot teams. *IEEE Robotics and Automation Letters*, 2(3):1712–1717, July 2017. ISSN 2377-3766. doi: 10.1109/LRA.2017.2673864.
- [25] T. Iqbal and L. D. Riek. Human-robot teaming: Approaches from joint action and dynamical systems. *Humanoid Robotics: A Reference*, pages 2293–2312, 2019.
- [26] R. W. Isenhower, M. J. Richardson, C. Carello, R. M. Baron, and K. L. Marsh. Affording cooperation: Embodied constraints, dynamics, and action-scaled invariance in joint lifting. *Psychonomic Bulletin & Review*, 17(3):342–347, 2010. URL <https://link.springer.com/article/10.3758/PBR.17.3.342>.
- [27] K. Kahol, P. Tripathi, S. Panchanathan, and T. Rikakis. Gesture segmentation in complex motion sequences. In *Proceedings 2003 International Conference on Image Processing (Cat. No.03CH37429)*, volume 2, pages II–105, Sep. 2003. doi: 10.1109/ICIP.2003.1246627.
- [28] C. D. Kidd and C. Breazeal. Robots at home: Understanding long-term human-robot interaction. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3230–3235. IEEE, 2008.
- [29] H. S. Koppula, A. Jain, and A. Saxena. *Anticipatory Planning for Human-Robot Teams*, pages 453–470. Springer International Publishing, Cham, 2016. doi: 10.1007/978-3-319-23778-7_30. URL https://doi.org/10.1007/978-3-319-23778-7_30.
- [30] A. Kubota, T. Iqbal, J. A. Shah, and L. D. Riek. Activity recognition in manufacturing: The roles of motion capture and semg+inertial wearables in detecting fine vs. gross motion. In *2018 IEEE International Conference on Robotics and Automation*, 2019.
- [31] F. Lacquaniti, C. Terzuolo, and P. Viviani. The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54(1):115–130, 1983. ISSN 0001-6918. doi: [https://doi.org/10.1016/0001-6918\(83\)90027-6](https://doi.org/10.1016/0001-6918(83)90027-6). URL <http://www.sciencedirect.com/science/article/pii/0001691883900276>.
- [32] S. Matsumoto, A. Washburn, and L. D. Riek. Human-robot co-manipulation demonstration. Technology Showcase for the Contextual Robotics Institute Forum, 2018.
- [33] S. Meyer zu Borgsen, J. Bernotat, and S. Wachsmuth. Hand in hand with robots: Differences between experienced and naive users in human-robot handover scenarios. In A. Kheddar, E. Yoshida, S. S. Ge, K. Suzuki, J.-J. Cabibihan, F. Eyssel, and H. He, editors, *Social Robotics*, pages 587–596, Cham, 2017. Springer International Publishing.
- [34] S. Moharana, A. E. Panduro, H. R. Lee, and L. D. Riek. Robots for joy, robots for sorrow: Community based robot design for dementia caregivers. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 458–467, March 2019. doi: 10.1109/HRI.2019.8673206.
- [35] A. Moon, D. M. Troniak, B. Gleeson, M. K. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft. Meet me where i’m gazing: How shared attention gaze affects human-robot handover timing. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction, HRI ’14*, pages 334–341, New York, NY, USA, 2014. ACM. doi: 10.1145/2559636.2559656. URL <http://doi.acm.org/10.1145/2559636.2559656>.
- [36] B. Mutlu, T. Shiwa, T. Kanda, H. Ishiguro, and N. Hagita. Footing in human-robot conversations: How robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction, HRI ’09*, pages 61–68, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514109. URL <http://doi.acm.org/10.1145/1514095.1514109>.
- [37] S. Nikolaidis and J. A. Shah. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, pages 33–40. IEEE Press, 2013. URL <https://dl.acm.org/citation.cfm?id=2447563>.
- [38] S. Nikolaidis, D. Hsu, and S. Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *International Journal of Robotics Research*, 36(5-7):618–634, 2017. doi: 10.1177/0278364917690593.
- [39] S. Ososky, D. Schuster, E. Phillips, and F. Jentsch. Building appropriate trust in human-robot teams. In *AAAI Spring Symposium: Trust and Autonomous Systems*, pages 60–65, 2013.
- [40] L. Peternel, N. Tsagarakis, D. G. Caldwell, and A. Ajoudani. Robot adaptation to human physical fatigue in human-robot co-manipulation. *Autonomous Robots*, 42:1011–1021, 2018. doi: 10.1007/s10514-017-9678-1. URL <https://doi.org/10.1007/s10514-017-9678-1>.
- [41] T. Rezvani, K. Driggs-Campbell, D. Sadigh, S. S. Sastry, S. A. Seshia, and R. Bajcsy. Towards trustworthy automation: User interfaces that convey internal and external awareness. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 682–688. IEEE, 2016. URL <https://ieeexplore.ieee.org/abstract/document/7795627/>.
- [42] L. D. Riek. The social co-robotics problem space : Six key challenges. In *Robotics: Science and Systems*, 2010.
- [43] L. D. Riek. Healthcare robotics. *Communications of the ACM*, 60(11):68–78, Oct. 2017.
- [44] L. D. Riek, T.-C. Rabinowitch, P. Bremner, A. G. Pipe, M. Fraser, and P. Robinson. Cooperative gestures: Effective signaling for humanoid robots. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction, HRI ’10*, pages 61–68, Piscataway, NJ, USA, 2010. IEEE Press. ISBN 978-1-4244-4893-7. URL <http://dl.acm.org/citation.cfm?id=1734454.1734474>.

- [45] P. Robinette, W. Li, R. Allen, A. M. Howard, and A. R. Wagner. Overtrust of robots in emergency evacuation scenarios. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction - HRI '16*, pages 101–108, Christchurch, New Zealand, 2016. IEEE Press. URL <https://dl.acm.org/citation.cfm?id=2906851>.
- [46] J. Scholtz. Theory and evaluation of human robot interactions. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*, Jan 2003. doi: 10.1109/HICSS.2003.1174284.
- [47] V. Srinivasan and L. Takayama. Help Me Please: Robot Politeness Strategies for Soliciting Help From Humans. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4945–4955, New York, New York, USA, 2016. ACM Press. doi: 10.1145/2858036.2858217. URL <http://dl.acm.org/citation.cfm?doid=2858036.2858217>.
- [48] A. Taylor, H. R. Lee, A. Kubota, and L. D. Riek. Coordinating clinical teams: Using robots to empower nurses to stop the line. In *In review*, 2019.
- [49] A. L. Thomaz and C. Chao. Turn-taking based on information flow for fluent human-robot interaction. *AI Magazine*, 32(4), 2011.
- [50] A. L. Thomaz, G. Hoffman, and C. Breazeal. Reinforcement learning with human teachers: Understanding how people want to teach robots. In *ROMAN 2006 - The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pages 352–357, Sep. 2006. doi: 10.1109/ROMAN.2006.314459.
- [51] A. Washburn, A. Adeleye, T. An, and L. D. Riek. Robot errors in proximal hri: Expectations and experience. under review.
- [52] B. Woodworth, F. Ferrari, T. E. Zosa, and L. D. Riek. Preference learning in assistive robotics: Observational repeated inverse reinforcement learning. In F. Doshi-Velez, J. Fackler, K. Jung, D. Kale, R. Ranganath, B. Wallace, and J. Wiens, editors, *Proceedings of the 3rd Machine Learning for Healthcare Conference*, volume 85 of *Proceedings of Machine Learning Research*, pages 420–439, Palo Alto, California, 17–18 Aug 2018. PMLR. URL <http://proceedings.mlr.press/v85/woodworth18a.html>.