












Using Human-Inspired Signals to Disambiguate Navigational Intentions

Justin Hart^(✉), Reuth Mirsky, Xuesu Xiao, Stone Tejada,
Bonny Mahajan, Jamin Goo, Kathryn Baldauf, Sydney Owen,
and Peter Stone

Department of Computer Science, The University of Texas at Austin,
2317 Speedway, Stop D9500, Austin, TX 78712, USA
{hart, reuth, xiao, pstone}@cs.utexas.edu,
{stonetejada, bm.bonny, jgoo, kathrynbaldauf, seowen}@utexas.edu

Abstract. People are proficient at communicating their intentions in order to avoid conflicts when navigating in narrow, crowded environments. Mobile robots, on the other hand, often lack both the ability to interpret human intentions and the ability to clearly communicate their own intentions to people sharing their space. This work addresses the second of these points, leveraging insights about how people implicitly communicate with each other through gaze to enable mobile robots to more clearly signal their navigational intention. We present a human study measuring the importance of gaze in coordinating people’s navigation. This study is followed by the development of a virtual agent head which is added to a mobile robot platform. Comparing the performance of a robot with a virtual agent head against one with an LED turn signal demonstrates its ability to impact people’s navigational choices, and that people more easily interpret the gaze cue than the LED turn signal.

Keywords: Social navigation · Gaze · Human-robot interaction

1 Introduction

When robots and humans navigate in a shared space, conflicts may arise when they choose conflicting trajectories. People are able to resolve these conflicts between each other by communicating, often passively through non-verbal cues such as gaze. When this communication breaks down, the parties involved may do a “Hallway Dance,”¹ wherein they walk into each other—sometimes several times while trying to deconflict each other’s paths—rather than gracefully passing each other. This occurrence, however, is rare and socially-awkward.

Robot motion planners generally generate trajectories which can be difficult for people to interpret, and which communicate little about the robot’s internal state passively [3]. This behavior can lead to situations similar to the hallway

¹ <https://www.urbandictionary.com/define.php?term=Hallway%20dance>.

dance, wherein people and robots clog the traffic arteries in confined spaces such as hallways or even in crowded, but open spaces such as atria. The Building-Wide Intelligence project [10] at UT Austin intends to create an ever-present fleet of general-purpose mobile service robots. With multiple robots continually navigating our Computer Science Department, we have often witnessed these robots interfering with people passing them in shared spaces. The most common issue occurs when a human and a robot should simply pass each other in a hallway, but instead stop in front of each other, thus inconveniencing the human and possibly causing the robot to choose a different path.

Previous work [6] on the Building-Wide Intelligence project sought to prevent these conflicts by incorporating LED turn signals onto the robot. It was found that the turn signals are not easily interpreted by study participants. In response, the work introduced the concept of a “passive demonstration.” A passive demonstration is a training episode wherein the robot demonstrates the use of the turn signal in front of the user without explicitly telling the user that they are being instructed. In the previous study, the robot makes a turn using the turn signal within the field of view of the user. Thus, the user is taught how the signal works. However, limitations of this technique include that it requires (1) that the robot recognize when it is first interacting with a new user, allowing it to perform the demonstration, and (2) that an opportunity arises to perform such a demonstration before the signal must be used in practice.

This work designs and tests a more naturalistic signaling mechanism, in the form of gaze. We hypothesize that this gaze signal does not require a demonstration in order to be understood. Signaling mechanisms mimicking human communicative cues such as gaze may be far more easily understood by untrained users. When walking, a person will look in the direction that they intend to walk simply to assure that the path is safe and free of obstacles. Doing so enables others to observe their gaze, implicitly communicating the walker’s intention. Observers can interpret the trajectory that the person performing the gaze is likely to follow, and coordinate their behavior.

This paper presents two studies exploring gaze as a cue to express navigational intentions. The first study is a human field experiment exploring the use of gaze when navigating a shared space. The second is a human-robot study contrasting a robot using an LED turn signal against a gaze cue rendered on a virtual agent head. These studies support the hypotheses that gaze is an important social cue used to coordinate human behavior when navigating shared spaces and that the interpretation of gaze is more clear to human observers than the LED signal when used to express the navigational intention of a mobile robot.

2 Related Work

The study of humans and robots navigating in a shared space [2, 6, 19, 20] has been of recent interest to the robotics community. Many works focus on how can a robot recognize human signals and react to them [11]. This work focuses on the robot communicating its intentions so nearby pedestrians can change their

course. Baraka and Veloso [2] used an LED configuration on their CoBot to indicate a number of robot states—including turning—focusing on the design of LED animations to address legibility. Their study shows that the use of these signals increases participants’ willingness to aid the robot. Szafr et al. [20] equipped quad-rotor drones with LEDs mounted in a ring at the base, providing four different signal designs along this strip. They found that their LEDs improve participants’ ability to quickly infer the intended motion of the drone. Shrestha et al. [19] performed a study similar to ours, in which a robot crosses a human’s path, indicating its intended path with an arrow projected onto the floor.

In this work, we contrast a gaze cue made by a virtual agent head on our mobile robot with an LED turn signal. Whereas a person who is attending to an object or location may do so without the intention of communication; to an onlooker, communication nonetheless takes place. This communication is said to be “implicit,” as the gaze is performed for perception. This role of gaze as an important communicative cue has been studied heavily in HRI [1], and it is a common hypothesis that gaze following is “hard-wired” into the brain [4].

Following this line of thought, Khambhaita et al. [9] propose a motion planner which coordinates head motion to the path a robot will take 4s in the future. In a video survey in which their robot approaches a T-intersection in a hallway, they found that study participants are significantly more able to determine the intended path of the robot in terms of the left or right branch of the intersection when the robot uses the gaze cue as opposed to when it does not. Using different gaze cues, several works [13,14] performed studies in a virtual environment in which virtual agents used gaze signaling when crossing a study participant’s path in a virtual corridor. Our work leverages a similar cue, but differs from theirs in that we use a physical robot in a human-robot study.

Recent work has studied the use of gaze as a cue for copilot systems in cars [7,8], with the aim of inferring the driver’s intended trajectory. Gaze is also often fixated on objects being manipulated, which can be leveraged to improve learning from human demonstrations [18]. Though the use of instrumentation such as head-mounted gaze trackers or static gaze tracking cameras is limiting for mobile robots, the development of gaze trackers which do not require instrumenting the subject [17] may soon allow us to perform the inverse of the robot experiments presented here, with the robot reacting to human gaze.

3 Human Field Study

In a human ecological field study, we investigate the effect of violating expected human gaze patterns while navigating a shared space. In a 2×3 study, research confederates look in a direction that is congruent with the direction that they intend to walk; opposite to the direction in which they intend to walk; or look down at their cell phone providing no gaze cue. The opposite gaze condition violates the established expectation in which head pose (often a proxy for gaze) is predictive of trajectory [15,21]. The no gaze condition simply eliminates the gaze stimulus. Along the other axis of this study, we vary whether the interaction happens while the hallway is relatively crowded or relatively uncrowded.

We hypothesize that violating expected gaze cues, or simply not providing them, causes problems in interpreting the navigational intent of the confederate, and can lead to confusion or near-collisions. Specifically

Hypothesis 1. *Pedestrian gaze behavior that violates the expectation that it will be congruent with their trajectory leads to navigational conflicts.*

Hypothesis 2. *The number of navigational conflicts will increase when gaze behavior is absent.*

3.1 Experimental Setup

This study was performed in a hallway at UT Austin which becomes crowded during class changes (see Fig. 1). Two of the authors trained each other to proficiently look counter to the direction in which they walk. They acted as confederates who interacted with study participants. Both of the confederates who participated in this study are female. A third author acted as a passive observer (and recorder) of interactions between these confederates and other pedestrians.



Fig. 1. The hallway in which the human field study took place.

In a 2×3 study, we control whether the interaction occurs in a “crowded” or “uncrowded” hallway, and whether the confederate looks in the direction in which they intend to go (their gaze is “congruent”), opposite to this direction (their gaze is “incongruent”), or focuses unto a mobile phone (“no gaze” is available). Here, “crowded” is defined as a state in which it is difficult for two people to pass each other in the hallway without coming within 2 m of each other. It can be observed that the busiest walkways at times form “lanes” in which pedestrians walk directly in lines when traversing these spaces. This study

was not performed under these conditions, as walking directly toward another pedestrian would require additionally breaking these lanes.

The observer annotated all interactions in which the confederate and a pedestrian walked directly toward each other. If the confederate and the pedestrian encountered problems walking around each other or nearly collided, the interaction is annotated as a “conflict.” Conflicts are further divided into “full,” in which the two parties (gently) bumped into each other; “partial,” in which the confederate and pedestrian brushed against each other; or “shift,” in which the two parties shifted to the left or right to pass after coming into conflict.

3.2 Results

A total of 220 interactions were observed (130 female/90 male), with 112 in crowded conditions and 108 in uncrowded conditions. In the crowded condition, the confederate looked in the congruent direction in 31 interactions, in the incongruent direction in 41 interactions, and with no gaze in 40. In the uncrowded condition, the confederate looked in the congruent direction in 29 interactions, in the incongruent direction in 44 interactions, and with no gaze in 35. A one-way ANOVA found no significant main effect between the crowded and uncrowded conditions ($F_{1,218} = 1.49, p = 0.22$), or based on gaze direction assuming the confederate looked either left or right ($F_{1,143} = 1.28, p = 0.26$). The “no gaze” condition was excluded from this analysis, in order to isolate the effect of the direction of gaze.

Whether the confederate goes to the pedestrian’s right or left during the interaction has a significant main effect ($F_{1,218} = 9.44, p = 0.002$), demonstrating a significant bias for walking to the right-hand side. Whether gaze is congruent, incongruent, or non-existent ($F_{2,217} = 5.02, p = 0.007$) also has a significant main effect. Post-hoc tests of between-groups differences using the Bonferroni criteria show significant mean differences between the congruent group and the other groups (congruent versus incongruent: $md = 0.221, p = 0.017$, congruent versus no gaze: $md = 0.191, p = 0.033$), but no significant mean difference between the incongruent and no gaze conditions ($md = -0.030, p = 1.00$). Perhaps the reason that there is no significant difference between the incongruent and no gaze conditions is that people use more caution when passing the no gaze pedestrian. Giving a the faulty signal of incongruent gaze, however, increases the number of conflicts. Our results support Hypothesis 1, that pedestrian gaze behavior violating the expectation of congruent gaze leads to conflicts; but not Hypothesis 2, that the absence of gaze will also lead to conflicts. A full breakdown of conflicts based on the congruent condition versus the incongruent gaze condition and the no-gaze condition can be found in Table 1.

Table 1. Human field study results

Conflict type	Congruent	Incongruent	No gaze
Partial	9 (15%)	11 (13%)	7 (9%)
Quick shift	6 (10%)	27 (32%)	13 (17%)
Full	0 (0%)	3 (4%)	1 (1%)
Any conflict	15 (25%)	41 (48%)	21 (28%)
No conflict	45 (75%)	44 (52%)	54 (72%)
Total	60 (100%)	85 (100%)	75 (100%)

4 Gaze as a Navigational CUE for HRI

We engineered a system in which a robot uses a virtual agent head to gaze in the direction that the robot will navigate toward. This study tests whether study participants understand this cue more readily than an LED turn signal. A study participant starts at one end of a hallway and the robot starts at the other end. The participant is instructed to traverse the hallway to the other end. The robot also autonomously traverses it. As a proxy for measuring understanding of the cue, the number of times that the human and robot come into conflict with each other is measured for two conditions: one in which the robot uses a turn signal to indicate the side of the hallway that it intends to pass the person on, and one in which it uses a gaze cue to make this indication.

This study tests the following hypothesis:

Hypothesis 3. *The gaze signal results in fewer conflicts than the LED signal.*

The robot’s navigation system models the problem of traversing the hallway as one in which the hallway is divided into three lanes, similar to traffic lanes on a roadway, as illustrated in Fig. 3. If the human and the robot come within 1 m of each other as they cross each other’s path, they are considered to be in conflict with each other. This distance is based on the 1 m safety buffer engineered into the robot’s navigational software, which also causes the robot to stop.

GAZE CUE. To display gaze cues on the BWIBot we developed a 3D-rendered version of the Maki 3D-printable robot head.² The decision to use this head is motivated by the ability to both render it as a virtual agent and, in future work, to 3D print and assemble a head that can be contrasted against the virtual agent. The virtual version of the head was developed by converting the 3D-printable STL files into Wavefront .obj files and importing them into the Unity game engine.³ To control the head and its gestures, custom software was developed using ROSBridgeLib.⁴ The head is displayed on a 21.5 inch monitor mounted

² <https://www.hello-robo.com/maki>.

³ <https://unity.com/>.

⁴ <https://github.com/MathiasCiarlo/ROSBridgeLib>.



Fig. 2. The hallway constructed for this experiment (left) and a snapshot of the experiment during execution (right).

to the front of the robot. When signaling, the robot turns its head 16.5° and remains in this pose. The eyes are not animated to move independently of the head. The head turn takes 1.5 s. These timings and angles were hand-tuned and pilot tested on members of the laboratory. The gaze signal can be seen in Fig. 4 (left).

LED CUE. The LED cue is a re-implementation of the LED turn signals from [6]. Strips of LEDs 0.475 m long with 14 LEDs line the 8020 extrusion on the chassis of the front of the BWIBot. They are controlled using an Arduino Uno microcontroller, and blink twice per second with 0.25 s on and 0.25 s off each time they blink. In the condition that the LEDs are used, the monitor is removed from the robot. The LED signals can be seen in Fig. 4 (right).

4.1 Experimental Setup

To test the effectiveness of gaze in coordinating navigation through a shared space, we conducted a human-robot interaction study in a hallway test environment. The environment is built from cubicle furniture and is 17.5 m long by 1.85 m wide (see Fig. 2).

After obtaining informed consent and, optionally, media release, participants are guided to one end of the hallway, where the robot is already set up at the




Fig. 3. The robot’s map of the hallway divided into lanes.




Fig. 4. Gaze and LED Signals. From left to right: Neutral gaze; Gaze signaling to the left lane; Neutral LED; LED blinking to signal a move to the left lane.

opposite end. The participant is instructed to navigate to the opposing end of the hallway. Both the participant and the robot start in the “middle lane,” as per the three traffic lanes in the robot’s navigational model. When the participant starts walking down the hallway, the robot is also started. This study follows an inter-participant design, in which each participant sees exactly one of the two cues—gaze or LED—and in which each participant traverses the hallway with the robot exactly once. The reason for the choice of this design is the strong evidence from [6] that there is a teaching effect as it relates to navigational cues, at least as far as LED turn signals are concerned. Given the presence of such an effect, we could expect all subsequent trials after the first to result in the participant and robot passing each other without conflict (in the LED case).

After completing the task of walking down the hallway, each participant responds to a brief post-interaction survey.⁵ The survey comprises 44 questions, consisting of 8-point Likert and cognitive-differences scales, and one free-response question. Five demographic questions on the survey ask whether people in the country where the participant grew up drive or walk on the left or right-hand side of the road and about their familiarity with robots.

⁵ The questions from the survey are available online. https://drive.google.com/drive/folders/1qVj-gU1aFwY6Eq2a_l9ZdesfmQ-QQOC8?usp=sharing. The first question “Condition” was filled in before participants responded.

To ensure that the study results are reflective of the robot’s motion signaling behavior, rather than the participants’ motion out of the robot’s path, the study is tuned to give the participant enough time to get out of the robot’s way if the decision is based on reacting to its gaze or LED cue, but not enough time if the decision is based on watching the robot’s motion. Three distances are used to tune the robot’s behavior: d_{signal} , $d_{execute}$ and $d_{conflict}$. The distance d_{signal} (4 m) is the distance at which the robot signals its intention to change lanes, which is based on the distance at which the robot can accurately detect a person in the hallway using a leg detector [12] and its on-board LiDAR sensor. The distance at which the robot execute its turn, $d_{execute}$ (2.75 m), is hand-tuned to be at a range at which it is unlikely that the participant would have time to react to the robot’s motion. If the participant has not already started changing lanes by the time the robot begins its turn, it is highly likely that the person and robot will experience a conflict. Thus, this study tests interpretation of the cue, not reaction to the turn. The distance at which the robot determines that its motion is in conflict with that of the study participant is $d_{conflict}$, which is set to 1 m. This design is based on the safety buffer used when the robot is autonomously operating in our building. For this study, the robot is tuned to move at a speed of 0.75(m/s), which is the speed the robot moves when deployed in our building. Average human walking speed is about 1.4(m/s). The range $d_{execute}$ was tuned empirically by the authors by testing on themselves prior to the experiment.

The robot also always moves into the “left” lane. In North America, pedestrians typically walk to the right-hand side of shared spaces. This behavior is demonstrated to be significant in our human field study as well, in Sect. 3.2. Fernandez et al. also conducted a pilot study [6] to test how often humans and the robot come into conflict when the robot goes to the left lane, rather than into the right lane, with no cue, such as an LED or gaze cue. That study showed that the pedestrians and the robot came into conflict 100% of the time under this regime.

Important to note in the interpretation of the results from this experiment is that **we expect the study participant and the robot to come into conflict 100% of the time** unless the robot’s cue—the LED or gaze cue—are correctly interpreted. This is because:

- The robot moving into the left-hand lane is expected to result in conflict 100% of the time.
- The robot’s motion occurs too late for the study participant to base their lane choice on it.

4.2 Results

We recruited 38 participants (26 male/12 female), ranging in age from 18 to 33 years. The data from 11 participants is excluded from our analysis due to software failure or failure in participation of the experimental protocol.

The remaining pool of participants includes 11 participants in the LED condition and 16 in the gaze condition. Table 2 shows the results from the

robot signaling experiment in these two conditions. A pre-test for homogeneity of variances confirms the validity of a one-way ANOVA for analysis of the collected data. A one-way ANOVA shows a significant difference between the two means ($F_{1,26} = 10.185$, $p = 0.004$). None of the post-interaction survey responses revealed significant results.⁶ These results support the hypothesis that the robot’s gaze can be more readily interpreted in order to deconflict its trajectory from that of another pedestrian.

Table 2. Experimental results

	LED	Gaze
No conflict	0 (0%)	8 (50%)
Conflict	11 (100%)	8 (50%)
Total	11	16

5 Discussion and Future Work

The goal of these studies is to evaluate whether a gaze cue outperforms an LED turn signal in coordinating the behavior of people and robots when navigating a shared space. Previous work found that LED turn signals are not readily interpreted by people when interacting with the BWIBot, but that a brief, passive demonstration of the signal is sufficient to disambiguate its meaning [6]. This study investigates whether gaze can be used without such a demonstration.

The human field study presented in Sect. 3 validates the use of gaze as a communicative cue that can impact pedestrians’ choices for coordinating trajectories. Gaze appears to be an even more salient cue than a person’s actual trajectory in this interaction.

In the human-robot study that follows, we compared the performance of an LED turn signal against a gaze cue presented on a custom virtual agent head. In this condition, the robot turns its head and “looks” in the direction of the lane that it intends to take when passing the study participant. Our results demonstrate that the gaze cue significantly outperforms the LED signal in preventing the human and robot from choosing conflicting trajectories. We interpret this result to mean that people naturally understand this cue, transferring their knowledge of interactions with other people onto their interaction with the robot.

Although it significantly outperforms the LED signal, the gaze cue does not perform perfectly in the context of this study. There are several potential contributing factors. The first is that, while the entire head rotates, the eyes do not verge upon any point in front of the robot. Furthermore, interpreting gaze direction on a virtual agent head may be difficult due to the so-called “Mona Lisa

⁶ Example interactions can be seen in the companion video to this paper, posted at <https://youtu.be/rQziUQro9BU>.

Effect” [16].⁷ Thus, the embodiment of the head may play into the performance of displayed cues. Follow-up studies could both tune the behavior of the head, and contrast its performance against a 3D printed version of the same head. The decision to use a virtual agent version of the Maki head is driven by our ability to contrast results in future experiments (upon construction of the hardware) between virtual agents and robotic heads.

Lastly, one might claim that if the LED condition resulted in 100% conflict, then simply changing the semantics of the LED turn signal would improve performance to be 100% “no conflict.” This, however, is not the case. Fernandez [5] compared an LED acting as a “turn signal” against one acting as an “instruction;” indicating the lane that the study participant should go into. Fernandez’s findings indicate no significant difference between these two conditions (“turn signal - conflict” $m = 0.9$, “instruction - conflict” $m = 0.8$, $p > 0.5$).

The overall results of this study are encouraging. Many current-generation service robots avoid what may be perceived by their designers as overly-humanoid features. The findings in this work indicate that human-like facial features and expressions may be more readily interpreted by people interacting with these devices than non-humanlike cues.

Acknowledgments. This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CPS-1739964, IIS-1724157, NRI-1925082), ONR (N00014-18-2243), FLI (RFP2-000), ARO (W911NF-19-2-0333), DARPA, Lockheed Martin, GM, and Bosch. Peter Stone serves as the Executive Director of Sony AI America and receives financial compensation for this work. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research. Studies in this work were approved under University of Texas at Austin IRB study numbers 2015-06-0058 and 2019-03-0139.

References

1. Admoni, H., Scassellati, B.: Social eye gaze in human-robot interaction: a review. *J. Hum.-Robot Interact.* **6**(1), 25–63 (2017)
2. Baraka, K., Veloso, M.M.: Mobile service robot state revealing through expressive lights: formalism, design, and evaluation. *Int. J. Soc. Robot.* **10**(1), 65–92 (2018)
3. Dragan, A., Lee, K., Srinivasa, S.: Legibility and predictability of robot motion. In: *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, March 2013
4. Emery, N.J.: The eyes have it: the neuroethology, function and evolution of social gaze. *Neurosci. Biobehav. Rev.* **24**(6), 581–604 (2000)
5. Fernandez, R.: Light-based nonverbal signaling with passive demonstrations for mobile service robots. Master’s thesis, The University of Texas at Austin, Austin, Texas, USA (2018)

⁷ Difficulties have been observed in the interpretation of the gaze direction of virtual agents. This effect is so named because it looks as though the Mona Lisa painting by Leonardo da Vinci is always looking at the observer.

6. Fernandez, R., John, N., Kirmani, S., Hart, J., Sinapov, J., Stone, P.: Passive demonstrations of light-based robot signals for improved human interpretability. In: Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 234–239. IEEE, August 2018
7. Jiang, Y.S., Warnell, G., Munera, E., Stone, P.: A study of human-robot copilot systems for en-route destination changing. In: Proceedings of the 27th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), Nanjing, China, August 2018
8. Jiang, Y.S., Warnell, G., Stone, P.: Inferring user intention using gaze in vehicles. In: Proceedings of the 20th ACM International Conference on Multimodal Interaction (ICMI), Boulder, Colorado, October 2018
9. Khambhaita, H., Rios-Martinez, J., Alami, R.: Head-body motion coordination for human aware robot navigation. In: Proceedings of the 9th International Workshop on Human-Friendly Robotics (HFR 2016), p. 8p (2016)
10. Khandelwal, P., et al.: BWIBots: a platform for bridging the gap between AI and human-robot interaction research. *Int. J. Robot. Res.* **36**(5–7), 635–659 (2017)
11. Kruse, T., Pandey, A.K., Alami, R., Kirsch, A.: Human-aware robot navigation: a survey. *Robot. Auton. Syst.* **61**(12), 1726–1743 (2013)
12. Leigh, A., Pineau, J., Olmedo, N., Zhang, H.: Person tracking and following with 2D laser scanners. In: Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA), pp. 726–733. IEEE (2015)
13. Lynch, S.D., Pettré, J., Bruneau, J., Kulpa, R., Crétual, A., Olivier, A.H.: Effect of virtual human gaze behaviour during an orthogonal collision avoidance walking task. In: Proceedings of the 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR), pp. 136–142. IEEE, March 2018
14. Nummenmaa, L., Hyönä, J., Hietanen, J.K.: I’ll walk this way: eyes reveal the direction of locomotion and make passersby look and go the other way. *Psychol. Sci.* **20**(12), 1454–1458 (2009)
15. Patla, A.E., Adkin, A., Ballard, T.: Online steering: coordination and control of body center of mass, head and body reorientation. *Exp. Brain Res.* **129**(4), 629–634 (1999)
16. Ruhland, K., et al.: A review of eye gaze in virtual agents, social robotics and HCI: behaviour generation, user interaction and perception. *Comput. Graph. Forum* **34**(6), 299–326 (2015)
17. Saran, A., Majumdar, S., Short, E.S., Thomaz, A., Niekum, S.: Human gaze following for human-robot interaction. In: Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 8615–8621. IEEE, November 2018
18. Saran, A., Short, E.S., Thomaz, A., Niekum, S.: Understanding teacher gaze patterns for robot learning. arXiv preprint [arXiv:1907.07202](https://arxiv.org/abs/1907.07202), July 2019
19. Shrestha, M.C., Onishi, T., Kobayashi, A., Kamezaki, M., Sugano, S.: Communicating directional intent in robot navigation using projection indicators. In: Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 746–751, August 2018
20. Szafir, D., Mutlu, B., Fong, T.: Communicating directionality in flying robots. In: Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI), HRI 2015, pp. 19–26. ACM, New York (2015)
21. Unhelkar, V.V., Pérez-D’Arpino, C., Stirling, L., Shah, J.A.: Human-robot co-navigation using anticipatory indicators of human walking motion. In: Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA), pp. 6183–6190. IEEE (2015)