

5

Spatial Interaction

What is covered in this chapter:

- The importance of the spatial placement of agents in social interaction;
- Basic understanding about human proxemics (i.e., how people manage space around them in a social context);
- How a robot manages the space around it, including interactions such as approaching, initiating interaction, maintaining distance, and navigating around people;
- How the properties of spatial interaction can be used as cues for robots.

In 2012, *Exertion Games Labs* released a drone exercise companion called Joggobot (see Figure 5.1). Runners who feel like they can use a little extra motivation or companionship during their run but don't have a personal trainer or a friend to join them can now have a drone accompany them during their exercise laps. One of the critical features of Joggobot is its placement in space during the run: right in front of the runner, like a carrot tempting a running horse. This position wasn't chosen on a whim. The developers studied where the drone should ideally be in relation to the runner (i.e., above, following, leading, on the side) and how much of a distance it should keep in order to maximize motivation (Graether and Mueller, 2012). They found that having the drone flying behind the jogger made people feel like they were being chased, which decreased their enjoyment in exercising. Users much preferred to take on the chasing role themselves.

This example shows that the placement of the robot with respect to the user is an important aspect of human–robot interaction (HRI). When only taking the need for collision avoidance into account while deciding on the optimal location or path of a robot, one might inadvertently create robot behavior that is considered uncomfortable, rude, or inappropriate. When a Roomba vacuum cleaner treats people as “obstacles” and keeps bumping into them as it tries to avoid them, it can comically seem to be “humping” their feet. Thus, when planning a robot's placement in space, it is important to take into account people's

Figure 5.1 The Joggobot Drone (2012). (Source: Photo provided by Eberhard Gräther and Florian “Floyd” Mueller)



preferences and the social norms that exist regarding such placement in relation to others.

5.1 Use of space in human interaction

When space is available, individuals are strongly expected to adhere to social distance norms. Most people feel that it is inappropriate for a stranger to sit down right next to them in an otherwise empty bus. However, when taking the bus during rush hour, we are forced to step into others' personal space, and it is acceptable to sit or stand close to people. Even though it is not considered impolite to stand next to someone on a busy commute, people still often feel uncomfortable, avoiding eye contact and quickly repositioning themselves at a greater distance when more space becomes available.

5.1.1 Proxemics

Cultural anthropologists coined the term *proxemics* to describe how people take up space in relation to others and how spatial positioning influences attitudes, behaviors, and interpersonal interaction.

Hall et al. (1968) describe four distance zones in their original work: intimate distance, personal distance, social distance, and public distance (see Figure 5.3). When the available space is (relatively) unlimited, these distances indicate the psychological closeness between people (see Figure 5.2).

As the name suggests, intimate distance is reserved for close personal relationships or the sharing of private information. Intimate distance ranges roughly from a few centimeters to about half a meter, depending on one's age and culture. Together with personal distance (which ranges from about half a meter to 1.2 meters), these zones make up the personal space of a person: the amount of space that people generally consider theirs to take up. Under normal circumstances, only friends, relatives, and partners are expected to come this close. For less personal relationships, such as acquaintances or colleagues, one is expected to maintain social distance, which ranges between 1.2 and about 4 meters between persons. Finally, public distance starts at around 4 meters, which is the distance people are expected to keep between them in relatively impersonal settings, such as public speaking at a conference.

Hall considered people's use of space as an often-overlooked dimension of cultural experience and noted that people from different cultures have varying personal proxemic preferences and expectations. For example, in "high-contact cultures" such as those of South America, people will frequently enter each other's personal space and touch, whereas in "low-contact cultures," such as the United States, touching a stranger may be construed as assault. Hall wittily observes that

Figure 5.2
Commuters during rush hour on the Tokyo underground having their personal space violated. We often deal with this by avoiding the gaze of others.



5.1 Use of space in human interaction

71

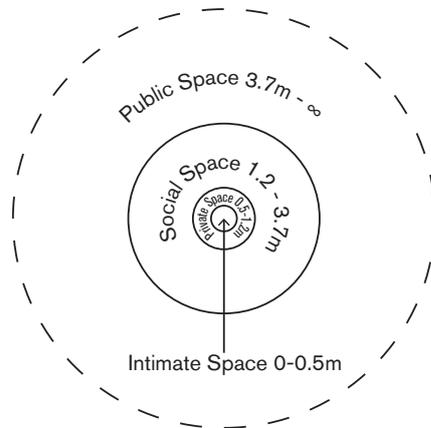


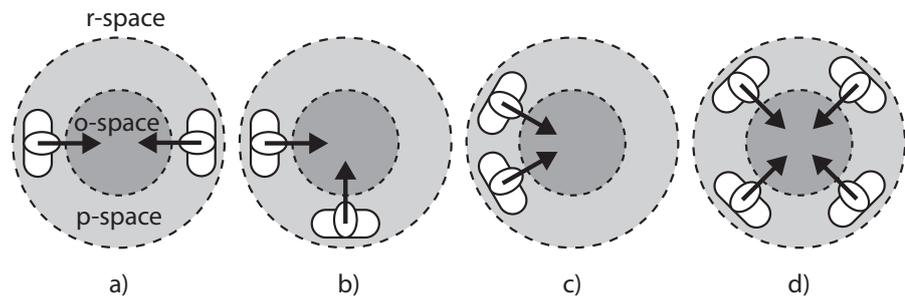
Figure 5.3
Intimate, personal, social, and public distance according to Hall et al. (1968).

North Americans visiting South America will find themselves “barri-
caded behind desks, using chairs and typewriters to keep the Latino at
what is to us a comfortable distance.”

Sometimes slight breaches of proxemic norms are made on purpose
by individuals, for instance, to create more psychological closeness or
perhaps to intimidate. For example, a man who nonchalantly places
his arm first on the backrest of the sofa where his date is sitting and
then cautiously inches closer and closer is making a transition from
personal distance to the intimate zone. The friend who touches your
arm when you are telling a personal story does the same, although with
a different underlying motive. However, these moves have to be made
very cautiously and under continuous assessment and reassessment of
the reaction of the other person. Few people would be charmed if the
hopeful suitor had abruptly placed himself right on his date’s lap. Like-
wise, when we attempt to comfort a colleague by giving a hug at the
wrong moment, the interaction can turn awkward rather quickly. This
is because the meaning of spatial interaction cues is highly contextual.
Unlike the friendly moves just mentioned, an investigator questioning
a suspect may “get in the suspect’s face” by moving as close to him or
her as possible to seem more threatening.

Not only the distance at which we interact with each other but also
our placement in relation to interaction partners are bound by social
norms. For example, researchers found that people who sat next to
each other were more cooperative, whereas people sitting opposite each
other behaved more competitively. During conversations, people usually
position themselves at an angle to each other (Cook, 1970). The way in
which people place themselves with respect to each other is therefore an
important aspect of the dynamics of interaction (Williams and Bargh,
2008).

Figure 5.4
Kendon's F-formations come in several variants, all of which include the components of o-, p-, and r-space, namely (a) the face-to-face, (b) the L, (c) the side-by-side, and (d) the circular formation.



5.1.2 Group spatial interaction dynamics

The importance of spatial dynamics goes beyond one-on-one interaction and is also salient in group interaction scenarios. The spatial orientation of people in a group in relation to others can make the group seem as if it were inviting more members or seeking to keep others out. For example, at a cocktail party, when people stand in a tight-knit circle, it can seem difficult to join in the conversation. However, if the group notices people wanting to join and opens up the circle so that there is space for new members to fill, it can be construed as an invitation to participate. This type of information can be useful for robots to gauge which groups of people they can approach in public spaces like museums or malls or if they want to affect the interaction dynamics of human groups.

Group spatial dynamics such as these were described by Adam Kendon as the “facing formation,” or “F-formation” (Kendon, 1990) (see Figure 5.4). These formations are created through the positioning of two or more people in relation to each other, such that the areas of space that they are facing and on which they focus their attention are overlapping. The space between these people, which is “one to which they have equal, direct, and exclusive access,” is termed the o-space. The group participants themselves are said to occupy the p-space, and they are surrounded by r-space. People can modify their positions to maintain this space or to include other participants in the group conversation, as in the previous example. Different configurations of the F-formation are possible, based on people’s orientation to each other, and are termed the face-to-face, L-shape, and side-by-side formation for two people and the circular formation and other shapes for larger groups.

These group formations have been used to understand people’s interactions with technology (Marshall et al., 2011) in general and with robots more specifically (e.g., Hüttenrauch et al., 2006; Yamaoka et al., 2010). In navigation around people, Pérez-Hurtado et al. (2016) found that a robot needs to be aware of people movements and cognizant of people engaged in conversation and not walk between them even if there is enough space.

5.2 Spatial interaction for robots

Robots will often share physical space with humans. Some robots are mobile, moving over the ground or through the air. Some of them have arms and manipulators so that they can interact with objects and users. The placement and movement of such robots with respect to people must be considered when designing human–robot interactions. Robots that do not respect the personal space of the user will evoke negative reactions or even rejection and withdrawal by the user. Robot designers can attempt to increase acceptance of the robot by having it keep an appropriate distance (assuming that they can code the robot in such a way that it knows what the “appropriate distance” is at a given point in time and space) and adjusting its position to create a fitting interaction experience. For example, a security robot might initially keep a polite distance but enter a person’s intimate space at some point in the interaction in an attempt to intimidate the person.

5.2.1 *Localization and navigation*

Before going into HRI, let us briefly explain basic techniques inherited from robotics that are required for a robot in order to engage in spatial interactions with humans. When a robot wants to interact with people, it needs to locate itself in space with regard to the people it aims to interact with. Thus, one of the basic techniques required for mobile robots is localization; a robot needs to know where it is. This is not a trivial problem. A typical robot is equipped with an odometer, a sensor that records the distance traveled by the robot’s wheels. However, as the robot travels, these lose accuracy, and the robot therefore needs to correct the information that the odometry provides about its location. The typical solution to this is to let the robot build a map of its environment and then cross-reference information on its location and orientation from the odometry with information from other sensors, such as a laser range finder or camera, to locate itself on the map. This process is known as simultaneous localization and mapping, or SLAM (Davison et al., 2007; Thrun et al., 2005).

In addition to reporting the robot’s location, localization can help the robot know what type of space it is in (e.g., whether it is in the living room or bathroom). However, it will not reveal anything about the whereabouts of any people in that space.

Identifying the location and orientation of people interacting with the robot thus is the next challenge. For detecting people at a short range, the robot will carry sensors such as two-dimensional (2D) cameras and depth cameras, that enable it to identify nearby people. The software processing the camera images can not only detect and track humans but also can report on the location of body parts such as arms, legs,

and heads. For tracking people at longer distances, there are techniques that use laser range finders (also known as light detection and ranging [LIDAR]). A motion-capturing system is sometimes used. By placing reflective or fiducial markers on people and objects, motion capture can be used to identify and locate the markers (and by extension, the people or objects they were initially attached to). However, these marker-based approaches are difficult to use outside a lab setting. Finally, researchers can also mount sensors, such as cameras, in the environment to track people (Brscić et al., 2013).

Moving the robot through a crowded environment, also known as robot navigation, is a well-studied problem in mobile robotics. To avoid collisions between the robot and objects or people, techniques such as the dynamic window approach (DWA) are often used (Fox et al., 1997). The idea behind this technique is that a system computes its future location based on the current velocity of the robot while at the same time considering whether to keep or alter its velocity within the limitation of its actuation capability—and while calculating a future velocity that does not result in a collision. Over longer time scales, there are techniques based on path planning. In these techniques, if a given goal of a robot is not within immediate view of the robot, a path-planning algorithm computes a set of way-points or paths for the robot that will let it reach its goal. In HRI, most path-planning algorithms that work well for navigating around obstacles will result in socially inappropriate behavior when tried around people.

Localization and navigation can also take various elements of interaction with a user into account. For instance, Spexard et al. (2006) developed a robotic mapping technique that uses input from dialogue with users to learn about new places in an environment. To develop a human-friendly mapping technique, Morales Saiki et al. (2011) had a robot explore the environment while collecting visual landmarks to build a cognitive map from a humanlike perspective; this enabled the robot to generate route instructions that people could easily comprehend. Researchers have also worked toward developing techniques to understand human spatial descriptions, such as route directions. For instance, Kollar et al. (2010) developed a technique to associate a user’s instructions and visual information about the environment to help the robot interpret the location mentioned by a user.

5.2.2 Socially appropriate positioning

Even though there are basic techniques for perception and navigation that allow robots to move around without colliding with obstacles, robots still often lack the capabilities to navigate in a socially appropriate way in the presence of other people. Suppose we want a robot to move through a corridor in an office building. What would happen if it consid-

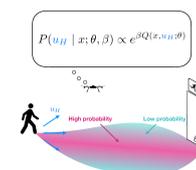
ers people as obstacles? When a person walked toward the robot from the other end of the corridor, the robot would continue to move straight down the corridor until inches before colliding and then move out of the way. Although it would avoid the person eventually, this behavior is very different from what humans would do in a similar situation: we yield to each other well in time, nonverbally showing which side of the corridor we will walk on, and will avoid entering each other's personal space. Thus, a robot waiting until the last moment before moving out of the way may be seen as confrontational or aggressive, even though it still avoids collision with a person.

Most mapping techniques for robots only provide geometrical maps, where people are considered obstacles. They do not contain information on which direction people are facing, if they are having a conversation or just standing close to each other, or how people are moving. Hence, there are several techniques that allow a robot to acquire a more human-aware representation of its environment.

One of the focuses in investigating proxemics in HRI has been identifying appropriate interaction distances between users and robots. These include questions like the following: How close do people prefer to stand relative to a robot? And how close should a robot approach people before it is considered rude or inappropriate or makes people feel uncomfortable (see Figure 5.5). Walters et al. (2005) measured the distance at which people feel comfortable when they are approached by a robot. They reported that the majority of people prefer a personal or social distance when interacting with a robot, although some people prefer to stand even closer. Hüttenrauch et al. (2006) reported that people preferred that the robot stand at distances derived from human proxemics. Investigating interactions between a robot and a group of people, Kuzuoka et al. (2010) reported that a robot can change the conversational F-formations of the group by changing its body orientation, and they also found that movement of the robot's whole body was more effective than having the robot just move its head.

Relational position is also important when people and robots interact while they are moving. To enhance a robot's social acceptability, techniques have been developed for robot navigation based on human proxemics. For instance, when a robot follows a user from behind, the robot can either follow the same trajectory as the user, or it can move directly to the user's current location, which might be a shorter and faster pathway. Gockley et al. (2007) showed that users perceive the first behavior as more natural. Morales Saiki et al. (2012) developed a technique that allows a robot to navigate side by side with its user, for which they found it important for the robot to anticipate the user's future motion. Furthermore, people's perceived safety does not necessarily correspond to what a robot computes to be safe. For instance, in the corridor passing problem, it was found that a robot needs to

Figure 5.5 The drone calculates a probabilistic model of where the human will go and plans a safe route around (Fisac et al., 2018).



maintain enough distance to avoid entering a person's intimate zone (Pacchierotti et al., 2006). Alternatively, a robot can mimic how people avoid colliding into each other. Luber et al. (2012) and Shiomi et al. (2014), for example, developed a pedestrian model that implemented collision avoidance for dynamic environments. Considerations of comfort and perceived safety can also be integrated into path planning. Sisbot et al. (2007) developed a path planner for a mobile robot that plans how to reach a given goal while avoiding situations that might make people uncomfortable. The planner takes into account aspects such as whether people are sitting or standing and whether the robot might surprise them by suddenly appearing from behind an obstacle. Fisac et al. (2018) used a probabilistic model of a human walking to plan and execute a safe trajectory for an indoor drone (see Figure 5.5).

Planning a motion path that people will perceive as safe and comfortable is also necessary when only a part of the robot enters the user's personal space. For example, when a robot arm is used near a person, such as when a person and an industrial robot collaborate on a shared task, the robot must take the socially appropriate distance into account when computing a path for its end effector (e.g., hand) to reach its given goal (e.g., grasp an object or hand an object to a person) (Kulic and Croft, 2005). This may make the robot's movement inefficient from a purely functional standpoint, but it will lead to a more positive evaluation of the interaction by the user (Cakmak et al., 2011).

5.2.3 Spatial dynamics of initiating HRI

Every social interaction has to be initiated by someone, perhaps by hovering in the vicinity of the person you want to talk to at a cocktail party while orienting your body toward the person, for example, or by approaching a colleague to hand over the annual report. How you approach each other and how the approach is perceived have implications for the ensuing interaction.

Approaching behavior is generally expected to have positive effects on both parties in the interaction. The approacher makes an effort to attract and share attention, which signals interest in the person being approached. At the same time, initiating an interaction triggers neural activity in reward-related brain areas, resulting in positive affect in the initiator (Schilbach et al., 2010). Initiating interaction is, furthermore, a sign of being assertive and having faith in one's capability to conduct a successful social encounter. What may be more surprising is that this runs the other way too. People who approach others are seen by their peers as having more personal control (Kirmeyer and Lin, 1987).

Imagine the moment when a person meets a robot for the first time. Either of them could approach the other to initiate the interaction.

Whereas this can be rather trivial for a person, a robot needs to be carefully designed to appropriately initiate an interaction. Approaching behavior for robots has been studied from early on in the field of HRI. For instance, in a situation where a robot joins a queue, the robot needs to respect the personal space of other people who are also waiting (Nakauchi and Simmons, 2002). When a robot encounters people, it needs to switch its navigation mode from purely functional to considering social distance and spatial configuration (Althaus et al., 2004).

Initiating an interaction is also context and task dependent. Satake et al. (2009) show how a robot offering information about the stores in a mall will fail to initiate an interaction if the approach is poorly planned and executed. The planned trajectory needs to be both effective and acceptable to human visitors (Satake et al., 2009; Kato et al., 2015). Whereas approaching from the front was found to be desired when a robot was trying to initiate a conversation, approaching from the front when the robot was delivering an object to a person was less preferred and resulted in more failures (Dautenhahn et al., 2006b; Shi et al., 2013).

Some recent work incorporates machine learning to generate appropriate approaching behaviors that fit with a context. Liu et al. (2016) designed approaching and initiating behavior for a store clerk robot using a fully automated analysis of observed human behavior. The researchers first recorded how people moved and talked in a camera store scenario and then used machine learning to extract typical speech behavior and spatial formations. These behaviors were then transferred to the robot. A user study showed that the learned speech and motion behavior was considered to be socially appropriate by users.

Even in the case where a person approaches a robot, the robot should respond at just the right moment. If it fails to do so, the user could find the interaction unnatural and awkward and might even give up initiating interactions in the future (Kato et al., 2015). Human proxemics studies, particularly observational studies on the interactions of humans with either one another or with robots, can provide more contextually attuned and relevant models. For instance, Michalowski et al. (2006) developed a categorical model of human spatial interaction and engagement with a receptionist robot from observations of people's interactions with the robot. They defined the appropriate timing and types of behavior (e.g., turning toward a person, saying hello) that the robot could perform with people in different spatial zones, in order to both be perceived as more approachable and to successfully initiate an interaction when appropriate.

Social navigation has become particularly relevant in the context of self-driving cars. The story goes that the first self-driving cars at Google drove optimal trajectories following the highway code, but they

frequently startled other road users by driving too close or cutting them off. Only when politeness was explicitly added as an optimization criterion did the cars drive in a way that was socially acceptable.

5.2.4 Informing users of the robot's intent

Robot motion trajectories are often used to convey the intent and goal of the robot. Path-planning algorithms have been developed to explicitly convey information through the robot's trajectory. For instance, by slowly passing a few meters from a visitor, a mobile robot is able to express whether it is available for an interaction (Hayashi et al., 2012). Similarly, trajectories have been used as a means to allow a robot with few options to express itself, such as cleaning robots and drones, to communicate their intent to users (Szafir et al., 2015).

During handover in HRI, that is, when a robot hands an object to its user, users prefer a robot to behave with “legibility”—in a way that allows users to understand its goal and intention (Koay et al., 2007a). Hence, researchers have developed algorithms to control a robot arm to generate legible motions while reaching a given goal. A robot could hand over an object to a person in many different ways, but the most energy-efficient way may be incomprehensible to a person, so it is better to perform a motion that is easier to interpret (Dragan et al., 2013).

When a robot works closely with a person, it needs to have the capability to understand how the person is perceiving the space around him or her. An important related capability is spatial perspective-taking (Trafton et al., 2005). Imagine a situation where two people are working together. One might ask the other to pass an object by saying “give me that object.” The referent of “object” will be obvious if there is only one object available. But what if there are several objects? For people, inferring the intended referent of “object” is often easy. We may use a complex set of cues, including gaze direction, body orientation, the prior context of the interaction, knowledge about the person and his or her preferences, task information, and other cues to disambiguate the request. For a robot, however, this can be rather complicated. Several approaches exist that allow the robot to take the perspective of the user. These often rely on geometric models that keep track of the location of people, robots, and objects and which of these are visible and reachable by whom (Lemaignan et al., 2017; Ros et al., 2010).

5.3 Conclusion

The study of spatial interaction in HRI is often inspired by our understanding of human proxemics, conversational relations, and relational positioning and approach behaviors, although we cannot expect the effects to always be the same. However, norms and understandings that

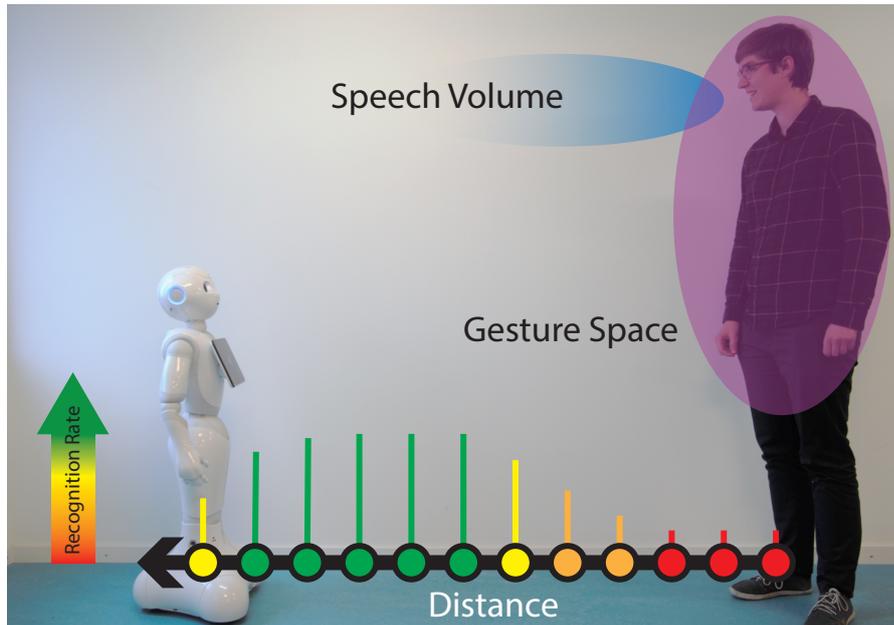


Figure 5.6 A lab setup for proxemic study of HRI.

are common knowledge for people—to the point where they may not even be aware of them anymore—often turn out to be not so trivial to incorporate into robot behavior. For instance, people will unconsciously and effortlessly adjust the distance to their conversation partner to an appropriate amount; however, a robot would need to conduct a careful computation to decide what distance it should keep during an interaction with its human counterpart. Even more difficulties are involved when the interaction is more complex, for example, when a robot has to approach a person, when it has to maintain spatial formation during a conversation, or when it has to navigate together with a person on the move. These considerations are important not only for achieving socially acceptable and comfortable HRI but also for ensuring that people understand the robot’s intentions and can engage with robots safely in their physical space.

Questions for you to think about:

- Let’s role play: To understand how much social information is involved in creating socially appropriate navigation, try to behave like a dumb robot that does not process any social information about space when interacting with a friend (maybe inform your friend beforehand, or “forget” to do so for a more natural response). What happened? How long could you keep this up?
- Think back on a situation when somebody violated your personal space. How did you notice? What was your reaction?

- Imagine you are an engineer building a robot. This robot will come to the market in Japan, Mexico, and the United States. Will the product be the same for every country? Will the robot's spatial navigation behaviors differ? If so, how?
- Think about the use of a robot in various daily situations (e.g., at home, at the office, and on a crowded train). Now, think about how you need to adapt the spatial-navigation behavior of the robot to fit each of these contexts. What would be important factors to consider in these different contexts?

Future reading:

Textbook to learn basic techniques for robot navigation:

- Howie M. Choset, Seth Hutchinson, Kevin M. Lynch, George Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun. *Principles of robot motion: Theory, algorithms, and implementation*. MIT Press, Cambridge, MA, 2005. ISBN 978-026203327. URL <http://www.worldcat.org/oclc/762070740>

More reading about space-related studies in HRI:

- Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013. doi: 10.1016/j.robot.2013.05.007. URL <https://doi.org/10.1016/j.robot.2013.05.007>
- Jonathan Mumm and Bilge Mutlu. Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *Proceedings of the 2011 ACM/IEEE International Conference on Human-Robot Interaction*, pages 331–338. ACM, 2011. ISBN 978-1-4503-0561-7. doi: 10.1145/1957656.1957786. URL <https://dl.acm.org/citation.cfm?doid=1957656.1957786>
- Satoru Satake, Takayuki Kanda, Dylan F. Glas, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. How to approach humans? Strategies for social robots to initiate interaction. In *4th ACM/IEEE International Conference on Human-Robot Interaction*, pages 109–116. IEEE, 2009. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514117. URL <https://doi.org/10.1145/1514095.1514117>
- Michael L. Walters, Kerstin Dautenhahn, René Te Boekhorst, Kheng Lee Koay, Dag Sverre Syrdal, and Chrystopher L. Nehaniv. An empirical framework for human-robot proxemics. *Proceedings of New Frontiers in Human-Robot Interaction*, 2009. URL <http://hdl.handle.net/2299/9670>