

Human-Aware Robot Navigation: A Survey

Thibault Kruse^{a,*}, Amit Kumar Pandey^b, Rachid Alami^b, Alexandra Kirsch^c

^a *Human-Centered Artificial Intelligence, Technische Universität München
Boltzmannstr. 3, 85748 Garching, Germany*

^b *LAAS CNRS*

7 Avenue du Colonel Roche, 31077 Toulouse Cedex 4, France

^c *Eberhard Karls Universität Tübingen, Department of Computer Science
Sand 14, 72076 Tübingen, Germany*

Abstract

Navigation is a basic skill for autonomous robots. In the last years human-robot interaction has become an important research field that spans all of the robot capabilities including perception, reasoning, learning, manipulation and navigation. For navigation, the presence of humans requires novel approaches that take into account the constraints of human comfort as well as social rules. Besides these constraints, putting robots among humans opens new interaction possibilities for robots, also for navigation tasks, such as robot guides. This paper provides a survey of existing approaches to human-aware navigation and offers a general classification scheme for the presented methods.

Keywords: autonomous robot, navigation, human-aware, human-centered environment, survey

1. Introduction

Robot technology is moving more and more from factories to everyday applications in households, offices and public places. Almost all such robots have to master the task of navigating through their environment. Human-aware navigation is the intersection between research on human robot interaction (HRI) and robot motion planning. Robot motion planning is an established engineering domain [52, 85]. HRI is a younger sub-domain of robotics [85], [18, 22, 29, 40]. As Feil-Seifer [18] puts it: “Human-robot interaction (HRI) is the interdisciplinary study of interaction dynamics between humans and robots.” There are many possible types of interaction between robots and humans, which depends largely on the robot design and capabilities. For our survey we consider those interaction dynamics that occur as a result of navigation.

While the term “HRI” may mainly evoke images of robots talking to humans or handing over items, it also has implications to robot navigation. By definition any robot has a body. Navigation means for the robot to change the position of its body, to move to a different place. This is often called “base motion”, as opposed to “arm motion”. In navigation, interaction with humans happens either as a purpose of the base motion, such as guiding humans, or as a side effect, such as when a delivery robot gets in the way of a human.

This paper uses a very broad concept of interactions describing any change in a human or a robot as a consequence of the presence of the other. So a human being distracted by a present robot is already an interaction. In this survey the distinction between explicit and implicit interaction is made: Explicit interaction happens when an action is made for the purpose of achieving a reaction by the partner; implicit interaction happens when actions are modified due to the (potential) presence of a partner, but no reaction is intended.

*Corresponding author

Email addresses: kruset@in.tum.de (Thibault Kruse), akpandey@laas.fr (Amit Kumar Pandey), rachid.alami@laas.fr (Rachid Alami), alexandra.kirsch@uni-tuebingen.de (Alexandra Kirsch)

Even though navigation is a basic robot skill, the robotics and HRI communities have not yet produced a holistic approach to human-aware navigation. However, pioneering works on several individual aspects of human-aware navigation exist. This paper introduces human-aware navigation as a research topic and provides an overview of the existing research, including not only navigation techniques, but also evaluation methods.

1.1. Background and Scope

Autonomous robots have been deployed in human environments, an overview of such robots was given in 2005 by Jensen [40]. Among the presented systems, there are several museum tour-guide robots, the museum setting being particularly suitable for deployment of robot prototypes. Here is a list of publications describing prototype robots that operated in public spaces:

- Rhino, by W. Burgard, A. B. Cremers, D. Fox, S. Thrun et al. [9]
- Robox, by K. O. Arras, R. Siegwart, et al. [4, 74, 86]
- Minerva, by S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. B. Cremers, D. Fox, et al. [96]
- Rackham, by R. Alami, R. Chatila, et al. [15]
- Mobot, by I. Nourbakhsh, C. Kunz, and T. Willeke [64]
- Cice, by A. Chella et al. [55]

All those robots had navigation modules that prevented collisions with humans or other obstacles using robust navigation strategies such as a minimal required obstacle distance or stop-and-wait behavior in situations of conflict were used to guarantee safety. A similar early work worth mentioning is by Prassler et al. [77] on an autonomous wheelchair navigating through a trainstation crowd. Thanks to experiences with such robots navigating robustly in human centered environments, researchers across the world have identified new challenges towards developing more human-friendly robots.

The research listed in this paper identifies situations where the acceptance of the robot can be improved.




Figure 1: Distribution of human-aware navigation publications over years. While the list may not be exhaustive or objective, the numbers still indicate the topic only got strong attention after 2000 and interest is still increasing.

We have conducted literature search on the topic of human-aware navigation. The search was limited to English language literature. We used online search engines and manual search of robotics conferences and journals of the last five years, recursively considering referenced work.

Figure 1 shows the result of the literature search, the citations are a subset of the bibliography of this paper, considering only those papers which explicitly deal with human-aware robot navigation. We cannot




Figure 2: Example scenario to explain challenges of human-aware navigation. The task of robot R is to guide person A to person B without causing discomfort to any human present. Dashed area shows example of robot laser range finder coverage. Challenges include passing through the corridor in formation with person A , while avoiding incoming person E , and approaching person B without disturbing person D unnecessarily.

claim our search to be exhaustive, but believe the offset of publications around 2001 and the raising trend are still viable conclusions from our selection. The growing number of publications on the topic indicates the growing interest. In the selection of publications we excluded the works on related topics such as perception of humans, tracking humans, tele-operated robots, legged robots, robot arm trajectory planning and affective signals of robots.

Using the selected papers as reference we identified smaller areas of research within the larger topic of human-aware navigation. In particular many papers focus on one of the aspects. We classified the papers into the categories comfort, naturalness and sociability (see Figure 3). Also it is useful to understand where a research result would fit into a general navigation framework, so we also partitioned the papers by what functionality of a general navigation framework they suggest to modify (as shown in Figure 7). Those two dimensions of challenge focus and navigation functionality serve as a structure for this document.

1.2. Example

To illustrate the special challenges of human-aware navigation, here is a simple example scenario. Figure 2 shows the scenario of a party with a robot butler. In an apartment, several humans are present and socially active. A new guest A has just arrived and the robot R has to lead her to the host B . The robot remembers the most recently observed position of B . It has range-limited sensors to cover an area in front of itself as depicted with a dashed line. Any state-of-the-art navigation framework can move the robot from its current position to a position close to the position of B . However, there are many aspects the robot must consider to avoid any discomfort, harm, and surprise to the persons present in the scene.

For example, the robot must move at a speed that is neither too slow nor too fast for person A to follow. On the way, it might encounter a situation, where person E moves towards C , thus crossing the way of the robot and person A . The robot must ensure that the person can still follow it. When moving through the corridor, the robot should move on one side (which side depends on the country and culture) and be especially aware that other people can come hurrying through the corridor out of the room, not seeing the robot and the guest. When moving around standing persons, the robot must keep an appropriate distance, far enough from the persons not to annoy or frighten them, but close enough for not taking unnecessary detours, which would be awkward for person A . While approaching to B , the robot should make sure that it is visible to B .

When other humans move, such as E moving towards C , the robot can still perceive their current positions. But for a smooth navigation behavior, it is necessary to predict the movement of E . While it is

necessary to react to the movement of E , the robot should not overreact. Just because E is crossing the way of R and A , it might not be necessary to adapt the path to B . Most of the time, it suffices to slow down slightly so that E can pass [47].

These are just example requirements drawn from the situation. More requirements of a similar nature exist, they all center around adaption to human properties, expectations and needs.

1.3. Organization of the Article

The rest of the article is organized as follows: Section 2 describes the challenges of moving among humans in more detail than the example in Figure 2, looking at the specific types of challenges comfort, natural movement and social aspects. Section 3 identifies the key functionalities of a navigation framework and categorizes novel technological approaches according to how they fit into the overall organization of robot navigation. A specific challenge for all HRI research is the evaluation of the methods. In multi-agent navigation specifically situations vary strongly and the psychological nature of the problem make it difficult to apply traditional robotics or psychology methods of evaluation, with additional effort to setup a robotic platform and sensor coverage for a test area for navigation. Evaluation methods used in the literature are therefore presented separately in Section 4. Section 5 concludes with a summary and outlook on open research questions.

2. Challenges of Human-aware Navigation

Research on human-aware robot navigation follows different goal: Most of the papers we collected attempted to minimize annoyance and stress, thus making interaction more comfortable for the human. Others strive to make robots behave more natural within their abilities or make robots behave according to cultural norms. All three goals have in common that they attempt to improve robot acceptance, but the methods vary. These terms like “comfort” and “natural” are used loosely in literature, so in order to classify the papers, we use the following definitions.

Comfort is the absence of annoyance and stress for humans in interaction with robots.

Naturalness is the similarity between robots and humans in low-level behavior patterns.

Sociability is the adherence to explicit high-level cultural conventions.

Comfort is different from safety: even when a robot is moving safely, an observer may still feel that the motion is not safe, lacking trust in the technology. So research on comfort attempts to not only make the robot move safely, but additionally in a way that also feels safe to humans. The opposite is possible as well of course, robot motion that feels perfectly safe to an observer, can still end up in a collision, if safety constraints are violated.

The most common approach for research in comfort is to identify robot behavior aspects that are clearly uncomfortable behavior, and to modify the behavior. This is similar to the approach taken for safety, where risks are identified and mitigated. This approach is very successful where a behavior aspect is clearly impacting acceptance in a negative way. However this approach struggles a lot with improving the robot behavior in a positive way, where the status quo is not clearly discomforting, but an improvement might be achievable still.

The other two categories thus focus on modifying the robot behavior to satisfy a given ideal of behavior. The category of “naturalness” strives to imitate nature, such as human motion, as the target behavior to recreate with a robot. This works well with several lower-level parameters of behavior, like dynamics, shapes and velocities, where a continuous distance-measure can be applied between behaviors. However due to the different abilities of robots compared to humans, not all aspects are suitable to be transferred from humans to robots. For high-level decisions, the problem is usually not how similar the robot behaves, but whether a robot makes the same decision as a human or not. So the category “sociability” instead uses known culture-dependent norms for civilized human behavior, which is appropriate for high-level decisions, where

individual action sequences or distinct decision alternatives can be formulated. The protocols we consider for sociability are constraints imposed by society, sometimes as solutions for conflict situations, sometimes to support civil order in society. One example is the rule to walk on the right-hand side in corridors.

Research on both “naturalness” and “sociability” thus is also usually expected to contribute to comfort, but we classify according to the authors’ primary intent, which can either be to measure and reduce known discomfort, or to modify behavior to either increase similarity to nature or obey defined social conventions. The well-known phenomenon of the “uncanny valley” shows that greater similarity does not generally increase acceptance of robots, thus the difference of focus is meaningful.

To clarify the distinction between the categories, here are examples for the challenge of the scenario in section 1.2 illustrated in Figure 2. The task of the robot is to guide person *A* to person *B*. For comfort, the robot should avoid to annoy any person involved. To do so, it should keep a proper distance to all persons, and avoid moving too fast or too slow. In order for the guiding task to be intuitive, the robot should maintain a natural formation with Person A, and move in smooth and goal directed behavior. A jittery path and unnecessary breaks would make the motion in formation intuitive. Finally the robot should respect social norms, such as moving on the right in the corridor, or giving priority to humans. Adhering to such cultural norms makes the robot behave in a civilized way, thus being more sociable.

The distribution of papers to topic is shown in Figure 3, with comfortable and natural motion being the dominant research topics. We show a fourth category “Other” there for papers dealing with diverse other aspects such as providing new interaction capabilities for robot in navigation.

Comfort	[1, 10, 13, 17, 19, 21, 27, 30, 34, 39, 42, 43, 44, 48, 49, 51, 53, 54, 56, 79, 60, 63, 65, 66, 67, 68, 69, 70, 71, 72, 81, 82, 83, 84, 87, 88, 89, 90, 91, 93, 97, 100, 105, 106]	44
Naturalness	[1, 6, 7, 13, 27, 28, 35, 37, 38, 43, 45, 48, 49, 50, 53, 54, 79, 62, 63, 66, 67, 68, 69, 70, 71, 72, 81, 82, 83, 84, 87, 88, 89, 90, 91, 93, 98, 99]	38
Sociability	[1, 13, 34, 35, 43, 44, 46, 56, 79, 63, 69, 70, 72, 81, 82, 93]	16
Other	[8, 12, 13, 30, 33, 37, 41, 42, 56, 57, 59, 60, 62, 69, 70, 71, 92, 104, 26]	19

Figure 3: The categories that the selected papers focused on.

Note that outside human-robot interaction, time-efficiency or energy-efficiency of robot actions are common optimization goals in robotics. With human-robot interaction, it may happen that social aspects take precedence over efficiency, or that efficiency is defined using social factors rather than just time or energy.

2.1. Human Comfort

When a robot moves to a defined goal, it can cause discomfort in human observers. It can do so by moving too close, too fast, too slow, too loud, being in the way, etc. This section presents research on identifying what causes discomfort and changing the robot behavior to reduce discomfort. The larger part of publications approaches this by avoiding negative emotional responses such as fear or anger. Fear is interesting does not only depend on the robot being safe but also to make it appear safe. The literature shows different strategies ranging from maintaining proper distance, proper approaching strategy, control strategies to avoid being noisy, and planning to avoid interference.

In the given literature the comfort requirement is mostly inspected regarding a distance a robot needs to keep from persons. This distance does not merely serve collision avoidance, but prevents the emotional discomfort humans may feel when approached closer than they like. The concept of a virtual personal space around a person that humans mutually respect - called “Proxemics” - was proposed by Edward T. Hall [32] for human-human interaction. He found different social distances humans chose for interaction depending on the relationship and intention as shown in Table 1.

The idea is that for each kind of interaction with other agents, humans will feel annoyed if the other is too close or too far away, greater distance than the expected one maybe indicating a dislike. The original proxemics model describes only one aspect of comfort in only one kind of situation: a distance to chose for explicit interaction when standing. Robot navigation deals with many more situations than static explicit interaction, and has to produce many more aspects of behavior. Other situations are: circumventing a

Designation	Specification	Reserved for ...
Intimate distance	0 - 45cm	Embracing, touching, whispering
Personal distance	45 - 120cm	Friends
Social distance	1.2 - 3.6m	Acquaintances and strangers
Public distance	> 3.6m	Public speaking

Table 1: Proxemic interpersonal distances as found by E. T. Hall [32]. How those categories and the values are to be adopted for the human-robot context is an open research question

human, standing in queue (or in an elevator) without explicit communication, and also when moving in formation. Also the proxemics model does not describe the orientation, pose and gaze agents choose or the trajectories agents choose when moving. But the general idea of proxemics of making the robot choose an appropriate social distance for any explicit or implicit interaction can be applied to other situations as well, and to other parameters than distance.

Generally from the perspective of making a person feel safe, the robot should try to avoid passing very closely and entering into the intimate or personal space, unless explicit interaction is required. This is regardless of objective safety, as the goal of comfort is to improve emotional quality, thus even avoiding fears that are objectively unfounded. Studies by Pacchierotti et al. [67, 68] show that even if there is no physical collision, to make a person feel safer the robot should deviate at a proper signaling distance. See the section 2.1.1 for more details.

Other than by keeping a proper distance as suggested by the Proxemics model, authors have attempted to reduce discomfort by modulating the velocity, approaching for explicit interaction from the front, etc. In the following such alternatives are presented. So research on comfort largely follows the concept of Proxemics, but other aspects of behavior than distance are also researched.

The following subsection describes how user studies attempt to quantify different aspects of behavior, and how robot behavior can be improved.

2.1.1. User studies involving comfort

Dautenhahn et al. [17] and Koay et al. [46] present studies in which the robot approaches sitting test participants several times from different directions to hand over an object. Afterwards the participants rate their discomfort. The studies showed that users in the test situation prefer a robot to approach from the front to hand-over the object, rather than diagonally from the left or right.

Pacchierotti et al. [67, 68] performed studies for a robot passing a human in a hallway with the human walking in the opposite direction. The study suggests the need of prompting signals and a proper signaling distance. In their study they applied a control strategy for robots in corridors, as a reaction to persons coming the other way. The robot deviated to one of three lateral distances, where generally larger distances (0.4m between robot hull and human center) felt better to participants, while on a few occasions, a large lateral distance was judged unnatural.

Butler and Agah [10] present several experiments for a robot moving in a room with a standing person. They tested approaching the person at different velocities, circumventing a person with different trajectories and robot bodies, and the robot exploring the room in different ways. They found that for approaching, a velocity of 1m/s made subjects uncomfortable, while 0.5m/s was acceptable. To some subjects a humanoid robot felt less comfortable than a small robot for the circumventing trials at the same distance of 0.5m, which the authors attribute to the robot taking more space in the room, thus maybe appearing more menacing. For exploring the room, no significant difference was found between a structured and a random approach.

The study of Shi [84], measures how much humans evaded from an approaching robot when they felt scared. The robot would as an example move towards the subject at 4.5m/s and safely stop at a distance of

0.5m, and the subjects felt unsafe. So even if the robot controller prevents all risks of collision, approaches may still scare humans in the environment and thus cause discomfort. From these observations they designed a velocity constraint for moving toward a person at distance $v(r) \leq kr$ for $r < r_{max}$ such that the robot's maximum permissible velocity within some social radius r_{max} decreases by a scalar k as the robot approaches a person. A similar approach was also used by other authors [1].

Mumm and Mutlu [63] performed studies where the robot base was not moving, but the robot's head moved to establish or avoid gaze. By modulating the robot greeting of the subjects between friendly and rude, they created an initial likability of the robot. As the next step they measured the self-distancing behavior of participants in different robot gaze conditions, intensive vs. light. The goal was to establish which of four competing distancing models would best fit the human behavior. A very brief summary of these models is:

Reciprocity : Human distance themselves more when the robot gaze is intense

Compensation : Human distance themselves less when the robot gaze is intense

Attraction Mediation : Human distance themselves more when the robot was rude

Attraction-transformation : Human distance themselves more in trials with a rude robot with intense gaze, but will distance themselves less with a friendly robot with intense gaze

The data they collected partly supported both the compensation and the attraction-transformation models. The authors derive three design recommendations: The robot should be designed to make a good first impression, when good rapport is difficult to achieve the interaction should not be too intense, and knowledge about how much an interacting person likes robots should be used to gauge intensity.

Takayama and Pantofaru [93] similarly extract robot design principles for approaching distance and gaze depending on social context. Regarding robot gaze, they recommend averting gaze when approaching a person closer than a threshold to avoid conveying a threat, in particular when approaching women. Else they recommend that the familiarity with robots and attitude towards robots should be taken into account for social distance selection.

Walters et al. [100] present a study showing that over weeks, the preferences of participants for a social distance between them and the robot decreased.

Young et al. [103] compare several relative poses for a robot that is led on a leash by a participant, with the results indicating that the robot moving in front of the person was the most comfortable joint motion approach.

All the user studies above require constrained settings to provide measurable results. The focus is on finding quality criteria of comfortable behavior. Several studies use Wizard-of-Oz experiments and robots with scripted rather than autonomous navigation. Given the results of such studies, the question remains how an autonomous robot can be made to exhibit the same quality of behavior. The next section looks at novel algorithmic solutions to produce such improved behavior.

2.1.2. Techniques to improve comfort

The most common approach to avoid discomfort due to small distance is to define areas around humans as cost functions or potential fields [34, 44, 83, 89, 91]. Using either, a path can be found that balances distance traveled with caused discomfort. Social Cost models are described in more detail in section 3.4.1. Such cost functions are superior to solutions defining forbidden zones around humans [39, 51], because in confined spaces it can be necessary and useful for the robot to move very close to a person. In [89], the human's comfort has been taken into account in a cost grid for static humans. Whereas in [71] proximity guidelines have been used to plan a smooth path around a dynamic human or a group of people moving together. This approach has been compared with purely reactive robot navigation from the aspect of physical comfort [69].

Works such as [71, 83, 89] explicitly try to avoid discomfort and surprise caused by navigating in the region behind the human, the approach is to increase virtual costs in that region to make the robot prefer a greater distance or a different path.

There are special cases where the comfort distance to the same person changes dynamically, such as presented in Zender et al. [105]. They present a robot that increases the distance to the human before leaving a room through a door, such that opening the door does not become clumsy due to the robot proximity. A special form of following is tracking the human from the front, such as investigated by Jung et al., and also by Young et al. [41, 103]. By remaining in front of a walking person, the person does not have to turn the head to survey the robot, which adds to the walking comfort in this particular situation.

To take the aspect of comfort beyond the definition of “a distance to maintain”, Martison [59] takes into account the noise generated by the robot motion itself and presents an approach to generate an acoustic hiding path while moving around a person. A model of human acoustic awareness is presented in [58], to navigate the robot in a way that reduces the discomfort because of noise. This has been further used to move the robot to another location to reduce the ambient noise exposure on the human during the course of human-robot interaction [60].

One way to be comfortable for others is ‘not disturbing’ them if not necessary. This has been addressed in Tipaldi et al. [97] by avoiding the robot to navigate in the areas causing potential interference with others, while performing tasks like cleaning the home. Their approach is to maintain a “spatial affordance map” which contains probabilities for human activities (like having dinner) in the environment (like a home) for time intervals as Poisson processes. This map is populated from observations of persons in their environment and allows for the robot to reason about whether its navigation activity is likely to occur in the same room a human occupies.

The comfort aspect of navigation deals with eliminating obvious causes of discomfort. Beyond that, the interaction between robots and humans can still be improved. The next section shows how interaction quality can be improved by making motions more natural.

2.2. Natural motion

Several publications attempt to make robots navigate more acceptably near humans by making robots move more similar to how humans move. This is often called “natural” behavior. The assumption is that if a robot behaves like a human (or like an animal), the interaction between humans and the robot becomes easier and more intuitive for the humans. Related adjectives are “predictable”, “understandable”, “readable” or “legible”. All those concepts rely on the fact that robot motion is interpreted by human observers, who attempt to judge the motive and the future behavior of the robot. As an example this is necessary as a human might need to proactively interrupt the robot if the safety of the human, the robot or the environment are at risk. The likelihood that a human observer interprets the motion of the robot correctly is assumed to be higher if the robot moves naturally.

While increasing similarity to humans in general increases acceptability, there are two known risks. The “uncanny valley” [61] describes the phenomenon that robots can cause discomfort by being too human-like in some aspect. The other risk is to induce false assumptions in observers about the robot capabilities. However, apparently in navigation both these risks have not had an impact so far.

One aspect of natural motions is smoothness. This refers to both the geometry of the path taken as well as the velocity profile. Human motion is influenced by a principle of energy optimization, among other things, so unless some other factor takes precedence, it is possible to empirically observe human motions with near-minimum jerk [3]. Another aspect of naturalness is the motion relative to other agents. Proxemics deals with distances. But there is more to motion than distance. As an example two persons walking together walk in formation, such as side by side or one following the other. Gockley et al. [27] looked at a similar example of a robot following a person, and found that it was preferable to make the robot share the direction, but not necessarily the path of the person followed.

Althaus et al. [1] present a behavior design for approaching a group of people and maintaining formation. They suggest to maintain a certain distance to the closest person, and regulate the robot speed as a function of the distance. Also when standing in formation with more than one person, they suggest that the robot body faces the middle of the group. For a benign demonstration with subjects, those control functions achieved behavior that appeared natural to the subjects.

To start an explicit interaction Saulnier et al. [82] investigated how different kinds of nonverbal cues in the robot body language can catch the attention of a person in a natural way. In their study a robot

performed several behaviors while subjects were interviewing another person. As conditions, the robot just passed by, approached with quick and erratic motions, or approached slowly. The reactions of the subjects confirm that navigation behavior conveys messages and can be used for non-verbal communication, and must be considered to avoid misunderstandings.

Similar ideas were tested for robots to be deployed in shopping malls, by Hayashi [35] and Satake [81], investigating how the trajectory of the robot can adapt to that of walking humans, to catch their attention without disturbing them. Both found they could improve the acceptability by making the robot approach from the front.

Garcia et al. [57] investigate the situation of having several robots guiding several persons, using a Social Force Model (SFM) as a means to guide a group of humans in a natural way. This model is also used by Tamura et al [94] changing the model to allow for different forces when an opportunity to follow another human exists. Social force models represent moving agents like robots or humans as masses under virtual gravitational forces. Thus an agent can be drawn towards a goal or a leading person while at the same time being repelled by obstacles and other agents. This can be used as a control scheme for a robot where the virtual forces are added up and used as input to the robot motion control.

A particular problem for naturalness is the motion within crowds. The motion that evolves in human crowds is an emergent system arising from the individual decisions of all humans in the crowd, yet the decisions taken may also depend on the overall behavior, density and shape of the crowd. Research work on robots in crowds is usually done using simulation. Examples are the work of Berg et al. [99] investigating reciprocal avoidance as a working social norm, and of Henry et al. [37] investigating global path planning using only local information about the crowd. In environments with restricted space potential conflicts may also occur frequently. Robots and humans may be in the way of each other, and a mutual planning effort is required to allow robots and humans to go where they want to go. To reduce the involved effort, it is important for the robot to avoid creating ambiguous signals about its intention. In Kruse et al. [47] the crossing situation was analyzed and the legibility of the robot motion improved by making the robot adopt a trajectory that is similar to that of humans in such situations. The key is to adapt the velocity of the robot rather than the path, allowing the path to indicate the intention of the robot. A generalized investigation of close human-human encounters is done by Luber et al. [54] based on overhead video data. In simulation they show how assumed fixed social cost models around humans generate less human-like behavior than cost models that vary over time depending on the distance between robot and human.

Another special challenge for robots is to move in densely crowded areas. Müller et al. [62] and Tamura et al. [94] explore the strategy of making a robot exhibit human-like motion behavior in highly populated environments, in the sense that the robot moves along with the people who are moving in the same direction leading towards the goal of the robot.

One stream of research deals with the challenge of humans being moving obstacles, that can cause the navigation to fail to get to a goal. As an example, several papers put the focus on the robot avoiding to be blocked or delayed by avoiding predicted human paths [21, 65].

2.3. Motion under social constraints

According to our category of sociability, human-aware navigation can be improved adding features to the robot behavior that considering explicit social protocols for behavior in particular situations. Outside navigation, examples are shaking hands, greeting any person once a day, saying goodbye before leaving, knocking at a door before entering a room somebody else may have claimed, and so on. In navigation, there are rules such as standing in a queue, letting people leave an elevator before attempting to enter, giving priority to elderly people at doorways, walking on the right-hand side of corridors, ask for permission to pass, excusing oneself when one has to traverse a personal zone to reach a goal, and so on. We separate this from comfortable and natural motion, because we can construct cases where social rules are violated, but rules for comfort and natural behavior are satisfied. In the example of standing in a queue, a robot could leave the queue on a safe and naturally smooth trajectory, skip ahead respecting the comfort zones of people, and go to the front of the queue. While this behavior respects comfort and naturalness constraints, people may feel aggravated if they assume that robots as intelligent agents are also responsible for following social rules.

Several authors investigate improved approach directions to initiate explicit interaction [1, 10, 17, 43, 46, 63, 93], suggesting that the discomfort is caused by the violation of a social rule, rather than a violation of a comfort distance. Such social rules may involve control of the robot gaze direction in such situations. Similarly, social conventions might exist for making prompting gestures as suggested by Peters et al. [72].

The model for pedestrian behavior by Helbing [36] includes a bias to evade to a preferred side in cases of conflict, breaking symmetry. In a related way, pedestrians can often be observed to walk in virtual lanes in corridors with sufficient density of humans. Which side to prefer is a cultural preference, a norm that varies between cultures.

Kirby et al. [44] modeled the social norm of passing a person by maintaining the right side in the hallway. This has been achieved on a cost grid based framework (COMPANION) by increasing the social costs of moving on the right side of the person. Whereas Pandey et al. [71] incorporated this along with other social norms such as: overtaking a person on his/her right, avoiding to pass through a group of people, proactively moving on the right half portion in a hallway or a narrow passage, in their high-level social path planner framework.

For research on comfort we have seen how a robot can consider discomfort by moving too close to humans. Even outside the immediate vicinity to humans, robots may violate diverse cultural rules by moving in certain areas. This has been researched by Chung et al. [13], building a map of an environment with “spatial effects”. A spatial effect is present when humans can be observed to divert from the shortest or most efficient path to their goal. Crossing the road only at crossings, or standing in queues are such spatial effects, where humans follow conventions. Another kind of spatial effect appears as a consequence of affordances of empty spaces[78]. An empty space in front of a TV set allows people to view the screen, a robot passing there may cause frustration by disturbing the affordance. However if the TV set is off, this might not be a problem at all. Similar in certain rooms like kitchens, almost every spot is an affordance space for access to storage area or a kitchen appliance, so the for households the concept of affordance spaces requires context to be useful.

Similarly the space between interacting groups of people is used for a variety of non-verbal communication cues, such that several authors describe methods to circumvent rather than pass through groups of people [71, 83, 26].

2.4. Other aspects of human friendly navigation

While comfort, naturalness and sociability are the most-mentioned research directions of human-aware navigation, some publications focus on diverse other aspects that have no main category. Often the focus is enabling new forms of interactions with robots, such as moving in formation, or to optimize other aspects for moving among humans, such as time efficiency, or robustness.

Several publications try to follow the idea that people as agents adapt to the environment, so the robot may use that knowledge about humans to pursue its navigation goals. For example, a person who stands in the way of a robot may very well move aside without discomfort if approached by the robot that wants to pass. This knowledge can be used as shown by Kruse et al. [49] and Müller et al. [62]. Similarly, moving humans may themselves adapt their motion to avoid collision with the robot, this can also be exploited [2, 98].

Several works attempt to improve the robot motion with respect to formations with persons, rather than just avoiding the human. There are several methods to research robot behavior for formations. The F-formations concept is mentioned by several authors to classify static human formations like the one in Figure 4a [1, 78]. Studies have shown that explicitly interacting groups of humans chose positions on a virtual circle. This concept can be used to define which areas a robot should avoid, and how to select a spot if the robot’s goal is to start explicit interaction. For dynamic situations, it is useful to classify the situation and behavior qualitatively, and Hanheide et al.[33] have used the Qualitative Trajectory Calculus (QTC) schema to classify 2-agent situations and select appropriate robot behavior. As shown in Figure 4b, the QTC considers the relative motion of agents to each other by comparing the motion vector to the connecting line between the agent positions. For either agent this yields whether the agent is moving towards the other or away from him, and whether it is moving on the left or right hand side. Given neutral states for each aspect, this yields 81 possible states, and the evolution of states in a situation allows to categorize the current




Figure 4: Different approaches to research formations. a.: The F-Formation concept for static situations defines an “O-space” reserved for explicit interaction, and the “P-Space” where agents can place themselves to express a desire to join the interaction. b.: For dynamic situations, the Qualitative Trajectory Calculus allows to qualify situations in an abstract way. c.: For robots moving in formation, virtual targets are used to define the current desired robot position. Virtual target need to be adapted with respect to obstacles.

situation and solutions, leaving it to further research to choose the best solution in this qualitative space of alternatives. Finally for a robot following in formation with a human, several authors define a virtual target for the robot to follow, as sketched in Figure 4c. Some authors focus on defining the ideal relative position of the virtual target [38, 41, 70, 103], while others focus on adapting the virtual target in cases of conflict [27, 37, 77]. The concept can also be extended to moving formations with many participants [16]. The results on virtual targets are however difficult to unify.

When guiding a person to a goal place a robot has to coordinate motion not just to avoid discomfort, but to achieve a joint goal. In this context, the person to be guided may deviate from the path planned by the robot. The guided person could show various natural deviations in her path and behavior. The person can slow down, speed up, deviate or even suspend the process of being guided for various reasons. This poses challenges for developing a robot’s navigation behavior, which is neither over-reactive nor ignorant about the person’s activities. The robot should try to support the person’s activities and re-engage the person if required. In this context [19] adapts a trajectory learning based approach to slowdown or stop when the human is not following by maintaining a proper distance. In [26] a group of robot tries to guide a group of people and adapting their trajectories to avoid humans getting lost from the group. In [70] a framework has been presented to re-engage a guided person, in the case she suspends the guiding process and moves away significantly. In that case the robot re-plans a smooth path to approach the person in a goal-oriented manner so that to exert a kind of social force pulling towards the goal.

2.5. Summary of human-awareness in robot navigation

The properties that a robot needs to exhibit to navigate in a human-aware fashion can mostly be categorized as comfort, naturalness and sociability.

Human comfort is the most mentioned topic in the literature. A moving robot should make a person feel safe not just by moving safely, but by moving in a way that is perceived as safe. But that alone does not avoid all situations of discomfort. Comfortable motion is currently mostly implemented as making the robot prefer additional distances on top of safety distances. However, researchers are beginning to take the issue of comfort beyond the proximity aspect, such as making the robot intelligent to be less noisy, or less interfering by its motion pattern.

Beyond avoiding discomfort, robot acceptance can be increased in several ways. Natural robot motion attempts to increase acceptability by reducing differences between robot and human motion. Such situations where naturalness is important are motion in formation, and motions within limited space with conflict potential, such as moving in a crowded place, following or accompanying a person.

Social motion deals with modeling and respecting cultural norms for specific situations, not only to avoid discomfort, but to avoid resentment and improve explicit interactions. It is implemented sometimes using hand-picked rules such as for moving on the right-hand side of lanes or approaching human from the front for explicit interaction, or it is learned from human behavior as spatial effect.

Further, there are research directions that take into account humans explicitly as intelligent and social agents, whose behavior can not only be predicted but also be influenced and used. This leads to solutions to particular situations of spatial conflicts.

As a summary of the features of human-aware navigation in the literature, here is a list of major human-aware capabilities that a robot could exhibit during navigation:

1. Respect personal zones
2. Respect affordance spaces
3. Avoid culturally scorned upon behaviors
4. Avoid erratic motions or noises that cause distraction
5. Reduce velocity when approaching a person
6. Approach from the front for explicit interaction
7. Modulate gaze direction

The list only represents the research so far, it is possible that in the future further human-aware capabilities are added to that list, as an example additional non-verbal communication cues to emit during motion, and similarly reacting to signals given by humans.

Placing a robot in a human environment also allows to implement human-related features besides moving in human-friendly manner. Figure 5 shows special form of movements and how much research exists for each, showing person-following as the comparatively easiest challenge.

Follow a person	[27, 30, 37, 38, 41, 77, 103, 105]	8
Solve blocked passage and dense crowd	[2, 37, 49, 98]	4
Guide a person	[19, 70, 26]	3
Move in formations	[1, 16, 57]	3

Figure 5: Specific challenges for robots navigating among humans

More details on the citations are provided

Having a good overview of the state-of-the-art research and human-awareness requirements, one may be tempted to create a navigation framework that satisfies all the currently known requirements. The following section maps the required features to common navigation functionalities.

3. Towards a human-aware navigation framework

The previous section defined human-aware navigation based on the existing literature. This section investigates the approaches to satisfy these requirements. A large diversity of navigation systems have been built for the purpose of research focused on particular situations. One challenge is to discern how to integrate the research of several publications into a unified framework. Common to practically all navigation frameworks is a division of the software into reactive and deliberate components. A robot needs to be both reactive to the environment, as well as being able to solve the motion planning problem. The motion planning problem, also called the “piano mover’s problem” is this: Given a workspace, obstacles in the workspace, and a robot, find a solution that allows the robot to move from one configuration to another without colliding with obstacles.

Software modules that plan the whole motion before executing it are called deliberate, modules that only plan a small time ahead during execution are called reactive. Architectures that combine separate reactive and deliberate modules are called “hybrid architectures” [80, ch. 25.7] or “layered architectures” [85, p. 110, 187]. In a layered architecture, the lowest layer is usually called the reactive layer. It senses the environment and controls the motion at high frequency to avoid unwanted situations such as collisions, or large deviations from a plan. Higher layers create plans for the path and behavior of the robot.

We use this distinction between deliberate and reactive modules to categorize work done on human-aware navigation planning.




Figure 6: A possible separation of concerns into modules concerning human aware navigation. Marked rectangles correspond to subsections of this document. Arrows only show main data flows, many more connections are possible.

3.1. General navigation software modules

This section describes basic navigation functionalities in an attempt to generalize existing approaches, and introduce names for different sets of requirements and modules in navigation frameworks.

Figure 6, shows data processing functionalities common to many navigation frameworks. A box in the diagram represents just one functionality. The connections in the diagram show the main data flow, and not the control flow, which may be designed in many different ways. The figure is an informal generalization of several navigation frameworks [73, 86, 87, 90]. Figure 6 does not present a complete software architecture for navigation, it only shows all the functionalities that are relevant as research items in human-aware navigation specifically. The functionalities are split into three groups: for sensing, deliberation, and acting. Simple architectures in the past had a control flow that follows the data flow, giving rise to the notion of a “sense-plan-act” (SPA) paradigm. The state-of-the-art is more complex, with modules having independent control cycles, and reactive components being able to initiate emergency stops without waiting for a high-level planning module to create a long-term plan.

The sense group deals with building up information about the robot’s environment and its own position. For environments with dynamic obstacles, the robot’s navigation behavior can be improved by detecting moving obstacles, tracking them and projecting their independent future motion or intent.

In the deliberation group, modules decide which actions to perform in what order, where to go for each action, how to get there and the behavior while on the way. An exemplary control flow is:

- The robot selects an action (e.g. go to human)
- The robot selects a pose (e.g. in front of human)
- The robot selects a path (e.g. avoiding a group of people talking to each other)
- The robot adapts an aspect of behavior while moving on the path (e.g. reducing speed to avoid noise near persons)

An optimal plan for a problem could only be obtained by making all these decisions in a single planning process. However the size of the search space prohibits such an approach. Instead the decisions are made

in sequence by specialized planners, though different strategies of segmenting the planning process exist. The most common strategy in research is to segment the decisions as in the list above, which works well in mostly static situations and is easy to develop.

When a robot has to act on its environment, it needs to choose a position to do so, and this choice can already impact the comfort of nearby persons. So pose selection generally rates positions in space with respect to the suitability for given actions. Path planning deals with finding a valid solution from a current position to a goal position. It is also called global planning. The main functionality here is to help the robot avoid being stuck in dead ends or taking very long detours when a shorter approach is possible. In general the output of path planning is a set of waypoints that the robot approaches one after the other. Behavior planning deals with modulations of the motions. When a robot moves, it has the freedom to modulate low-level aspects of behavior, by choosing different velocities, moving sideways, pausing, using gestures or sound as signals. Action planning and behavior planning are different as the former deciding on the “whether” and “what” of motions, and the latter on the “how”. As an example for action planning, Tipaldi et al. [97] shows planning for a vacuuming robot choosing an order in which to clean rooms to reduce the annoyance of the inhabitants under assumptions of a daily schedule. On the other hand, when moving from one room to another in an apartment, a robot can choose to keep the vacuum engine on or turn it off, depending on whether there are humans (for example watching TV) in the room the robot passes through. Such decisions belong to behavior planning instead.

The reactive group contains all control algorithms that send commands to the robot motors. This group maintains safety and is responsible for producing smooth motions. The local planning module is called ‘local’ because it only plans a certain distance or time ahead. This is required to ensure that planning time is bounded, so that a result can be guaranteed to be available in time while the robot moves at higher speeds among moving obstacles.




Figure 7: The technologies that the selected papers focused on.

Using this framework of different navigation module functionalities, the papers found in literature research of section 1.1 are classified with respect to the responsibilities they relate to in Figure 7.

Next the different module groups are presented in more detail, showing the different research directions and competing solutions.

3.2. Prediction

The sensing functionalities of a robot take in data from several sensors and produce geometric representations of the environment. Detecting and tracking humans using sensors on a moving robot is particularly difficult. For navigation, prediction is relevant because of the mutual influence of the human navigation processes and the robot navigation processes, and because of the communication effects of navigation as body language. Outside of human-aware navigation research, prediction of human navigation has very little interest apart from crowd motion effects. So researches in human-aware robot navigation also need to do research on human motion prediction. It is beneficial to have a model of what a person will do in the future. That way the robot can make plans that are compliant with the future actions of persons. Figure 6 shows a main data flow from prediction to deliberation. This is a simplification that is also sometimes used for the control flow. A complete approach needs to make plenty of predictions for each alternative decision the robot can make, as the robot’s decision will influence the decisions of humans as well. This complete approach does not scale well however (and is more complex to describe and implement), which is why the




Figure 8: Schematic examples of prediction techniques for moving humans. a: Linear. b: Likely state transitions in a grid. c: Using potential field with walls as repulsors and a perpendicular (dashed) line in front of the person. d: with growing uncertainty. e: Using random samples and stochastic weights. f: Social Force model, person attracted to goal, to walking partner, but also repulsed by partner when too close. g: Using library of observed paths. h: Using closest frontal annotated waypoints on map.

future motion of humans is usually predicted before making planning decisions, as if there were no mutual influences.

Prediction for navigation is closely related to simulating human walking behavior in general. Extensive overviews of related publications are given by Martinez and Furtley [20, 78]. However in robotics, there are additional constraints due to the limited sensor capabilities and real-time requirements for reactivity.

Two approaches to prediction exist, prediction based on reasoning and prediction based on learning. Figure 8 gives an initial overview of some prediction methods, that are explored in detail below. In reasoning-based prediction, the predictions are justified by assumptions of how agents behave in general. In learning-based techniques, predictions are justified by observations of how agents behave, in particular in special circumstances and environments. Both approaches have things in common and can be combined, there is no strict separation, but the publications presented here are divided by the approach they focus on.

In learning based prediction, typical trajectories within a given map are first collected over time. Those then serve to predict new samples of humans moving in that environment. Foka et al. [21] give another dimension along which to classify prediction. They call short-term prediction anything that predicts the path or trajectory, and long-term prediction anything that predicts the next resting position. The latter is more robust against agent-agent influences. As an example, when a robot can predict that a person wants to go to the coffee machine, then the robot crossing the path of the person will typically not change that intention of the person, while the path the person chooses may be altered.

While all prediction approaches we found operate directly on map coordinates, a recent publication by Hanheide et al. [33] opens the possibility to classify and predict human motion on a higher level of abstraction, using the Qualitative Trajectory Calculus, as already introduced in Section 2.4 and Figure 4.

The following subsections first present prediction based on fixed models and assumptions, and then at learning based techniques.

3.2.1. Prediction based on geometric reasoning

In most of the available research a static world is assumed, in which only the robot moves, as an early step in HRI to be extended to dynamic situations later. In a similar way, a lot of publications use simplified models of human motion prediction, such as assuming at any given time that a person will proceed to

walk straight in his current direction at his current speed, such as in Figure 8a. Such obstacle-agnostic prediction methods can be refined to take into account typical human motion properties, such as reduction of speed when walking in arcs, as applied by Granata et al. [30]. But where obstacles are present the projected position may overlap with obstacles, so this method has limited use but may be sufficient for some experiments.

An early model for prediction taking into account the environment is given in Tadokoro et al. [92], using stochastic processes to predict probable state transitions of persons tracked in a grid map of the environment. The idea is depicted in Figure 8b, where grid cells are shown with their likelihood being represented as the darkness of the cell. At crossings in the environment, alternative state transitions become likely, such as moving on in a corridor or entering a room. The idea and technique did not reappear in other publications on human-aware robot navigation.

A similar, but less powerful approach for estimation of the future motion of agents is described in [65], that models agents as obstacles moving in the direction along a corridor (evacuating an area), with the model allowing faster moving agents moving sideways to circumvent slower-moving agents. This also uses a grid model, but no stochastic model of the future human positions.

More recent publications tend to assume more detailed and precise information about the human position as would be obtainable by modern sensors. This leads to models that use pose vectors in continuous space instead of occupancy grids.

As an example, Hoeller et al. [38] demonstrated prediction techniques that take into account that other agents cannot cross obstacles. This approach uses potential fields where obstacles are repulsors and a virtual line in front of the person acts as attractor as shown in Figure 8c.

Both occupancy grid as well as continuous spatial models may benefit from capturing uncertainty of the prediction. In occupancy grids, this is done by assigning probabilities to multiple occupancy cells rather than just marking cells as occupied in the future. In continuous models it is common to represent the uncertainty as growing round or elliptic areas of where the person is suspected to move to, as shown in Figure 8d. More complex approaches use Gaussian distributions [50].

Using growing elliptic areas of uncertainty has inherent problems. On the one hand at crossings there are several distinct paths a human could take, that cannot be adequately represented by a single circle. On the other hand the growing areas lead to the freezing robot problem [98], where in sufficient distance the predicted human motion with uncertainty takes all the available space.

A possible solution to this is presented in Althoff et al. [2]. They integrated local planning and prediction using random sampling over the future path of several agents (including the robot), stochastically maximizing the expectation that no agent will collide with obstacles or one another. An example result is shown in Figure 8e, where two humans are predicted to mutually avoid collision. When considering several humans moving that way, it is also possible to use a model of social forces, as applied by Garcia et al. [57]. In such a model pedestrians walking as a group are supposed to act as if subject to virtual attractive and repulsive effects, shown in Figure 8f.

There are downsides to all reasoning-based predictions. Calculation times increase with the amount of information processed. Also generalized models tend to fail in special cases. As solution to this downsides, the next section lists research on machine-learning approaches to navigation.

3.2.2. Prediction based on machine learning

Machine learning techniques for human prediction have the advantage of improving over time, of adapting to special circumstances, and of yielding prediction results with low online computational effort (most of the effort happens offline).

Foka et al. [21] use a neural network for one-step ahead prediction, but offer no evidence that this yields better results than linear interpolation.

A richer model is presented by Chung et al. [12, 13], using a library of observed short-term trajectories to make short-term and longer-term predictions.

Bennewitz et al. [6, 7] describe how to learn a set of motion patterns for a given map of an environment from observations of human motion, using the Expectation Maximization Algorithm. The prediction here is more of a classification, meaning the technique can only predict one of the motion patterns it has learned.

A similar approach of using observed trajectories is shown in Satake et al. [81]. The idea is visualized in Figure 8g. Bennewitz et al. also extend their approach to predict human motion based on the identity of people, such that for an office worker, it is likely to walk into his own office.

A more high-level prediction using collected trajectories is given by Hansen et al. [34], who shows an approach to predict whether a person intends to interact with a robot, based on the trajectory of the person.

An alternative approach by Thompson et al. [95] uses a mixture of annotated spaces and heuristically determined waypoints as estimates for short-term navigation goals. The idea of waypoints is similar to the one given in Foka et al. [21], sketched in Figure 8h. The prediction quality was shown to be improved for the analyzed environment, compared to predicting motion based on current position and velocity alone.

The variant of Markov decision processes (MDP) used by Ziebart et al. [106] is apparently superior to similar approaches by being feature-based. This allows for the prediction to work in new environments and to immediately adapt to changes in the environment.

A drawback of all prediction techniques based on observations of one agent at a time is that the motion of several agents is not taken into account. As a result, the more crowded an area is, the worse the prediction results are compared to model-based approaches.

3.3. Pose Selection

Pose Selection is the task of identifying good places for the robot to stop and start manipulating. It is also sometimes called “robot placement”. As an example, to talk to a person, a robot should stand where the person can see and hear the robot, the robot should also usually face the person then. This is mostly a task that logically happens prior to navigation in the sense of motion planning. However, as it is closely linked to navigation, publications on that are listed here as well. Human-aware pose selection takes into account the discomfort or comfort a robot can cause when acting at a given spot. Section 2.1 introduced the concept of Proxemics [32] as a model of suitable distances between agents for interaction. Measuring comfort distances is analyzed for that purpose by several research teams [17, 42, 46, 63, 66, 100, 101]. In order to select a spot for the robot to stop and interact, the comfortable distance obviously allows a circle of points around a person. This leaves it to a pose planner to rate points on this circle. Ideally a position should be in front of the person (unless the interaction is to give a back rub). However, most publications do not go into detail on the pose selection. An exception is by Hansen et al. [34] using a potential field for navigation where areas in front of the person were dedicated stopping points.

More complex situations are described for cases where a person is moving, and a pose relative to the moving person is needed. Hoeller et al. [38] describe how their person following robot chooses one of five poses behind the person. Such virtual targets were already introduced in Section 2.4 and Figure 4. Pandey [70] has a similar set of predefined points around a person that the robot traverses to motivate the person to follow the robot as a guide. Cuntoor et al. [16] investigate general interaction scenarios where the robot has to take a position in a formation with several persons. In their experiments the robot knows how many members the formation has, and attempts to identify which formation out of a small set of shapes the group attempts to create. The robot then maintains its own position in the assumed formation.

There is also considerable work on full body motion planning for handover tasks between robots and humans that is excluded from this paper since it belongs to the domain of manipulation motion planning, not navigation. However as an example we mention recent work by Mainprice et al. [56] that plans a handover pose for both the robot and the human, such that both may have to move to a new location for the handover.

3.4. Path planning

Path planning is the task of choosing a list of waypoints to follow on a map in order to optimize the performance of the robot according to some global target function, like the distance traveled. This is often called synonymously “navigation” in the narrow sense, as in “navigation systems” giving travel directions to car drivers. The minimum feature of a path planner is to find an end-to-end path from the current position to the goal. Further optional features are to find the shortest path, or a path satisfying some other global property. Such a path is often called the ‘global plan’ for navigation. Navigation as a reasoning activity is only necessary where a motion towards a goal cannot be achieved sufficiently well using reactive

methods, such as turning towards a goal and moving straight. Reactive methods fail in environments that have sufficiently constraining obstacles like walls that also form dead ends. Navigation algorithms use maps of the environment representing blocked and available space. Humans can be taken into account as moving obstacles, as social agents not to disturb, or as partners in joint motion with particular interests.

The majority of work on human-aware path planning use graph-based search methods, with a graph representing states in a 2D map of the environment. In the graph each node is a point in which the robot can be without collision, and an edge between two nodes means the robot may travel between the nodes without causing collision. The kinds of graph used in the surveyed research are square grids, arbitrary lattices, and expanding random trees with freely chosen edge angles.

Expanding random trees are usually used for obstacle avoidance search with higher dimensionality. They can have less memory consumption in average cases and may find non-optimal results in less time. It is possible to also consider kinematic constraints in the construction of such trees. The main disadvantage for human-aware navigation is the difficulty to make good use of cost functions in the absence of optimality properties. While two publications use them [83, 91], their usage there for human-aware navigation is not proven to be superior to grids or lattices.

Grids avoid duplicating states and have a fixed memory size, they are also simpler to program. Given that optimality of results is proven, cost-function based approaches usually prefer grid-based approaches for the sake of provably optimal results, in particular for static situations. The disadvantage of grids is that the paths cannot take into account robot kinematic constraints. Kinematic lattices such as suggested in [50] are more suitable to represent kinematic constraints. A kinematic lattice here means a graph that is built from motion primitives based on robot properties, such that an edge of the graph is a so-called motion-primitive. To prevent the graph to grow too quickly, a special set of motion primitives is selected to build a regular structure, such that expanding a node often produces edges leading to existing nodes. However, kinematic lattices are more complex to construct than grids, so HRI publications have not adopted them yet.

Regardless of grids, kinematic lattices or trees, global path planning uses a search method to find a sequence of states from the robot position to the goal that is considered optimal with regard to a cost function. A major improvement to plan results is also possible when planning with time. This allows to find solutions that static representations cannot find. On the other hand, adding the dimension of time exponentially increases the search space and can cause the robot to respond too slowly to changes.

An alternative to graph-based methods is navigation based on partially observable Markov Decision Processes (POMDP) [21]. This approach focuses on dealing with uncertainty, and is difficult to adapt to global quality criteria. Also the approach does not scale well with the resolution of states. While POMDP-based navigation may generally have advantages, the technical effort to maintain such frameworks seems to have made it impracticable for research on human-aware navigation so far.

The next two subsections look at cost functions in detail, and then at temporal planning.

3.4.1. Global cost functions

Human-agnostic navigation will commonly choose the shortest or most energy efficient path in a graph, using a variant of the A* search algorithm. In HRI, the assumption is that the shortest or most energy efficient path is not necessarily the most desirable.

Instead in HRI the intention is to find a path that is also sufficiently safe, comfortable, natural, legible, etc. to persons in the area. Such desirable path attributes can be represented in global path planning as cost functions.

The purpose of a cost function is to assign a cost value to robot actions. Those costs depend on the context, notably the robot environment and its state. Individual researchers use varying details about the environment as state, the objects in the environment (their type, geometry, and state), the persons (their age and gender), their current activity, the current ongoing interactions among people, and interactions between people and objects. Also past data and predictions about the future may be encoded as part of the context. Basically all the knowledge a robot may have about its environment.

Definition 3.1 A cost function f_{cost} maps a cost value to all combinations of robot actions \mathbb{A}_r and action contexts \mathbb{C}_r :




Figure 9: Visualizations of separate costmaps, thick black areas are obstacles, the square in b is a robot, human in h is moving. Areas shaded in grey have costs, meaning the robot should prefer to avoid those if possible. Shapes and sizes of cost function may vary and be context-dependent. All cost-functions can be combined.

$$f_{cost}(\mathbb{A}_r \times \mathbb{C}_r) \rightarrow \mathbb{R}$$

In navigation, it is common to use the “action” of temporary occupancy at an (x, y) coordinate on a map, and the static map and 2d-pose of humans as a navigation cost function. Since costs are then assigned to (x, y) coordinates, some authors also call this a costmap:

Definition 3.2 A costmap for navigation f_{nav} maps a cost value to 2d coordinates and action contexts \mathbb{C}_r :

$$f_{nav}(\mathbb{R} \times \mathbb{R} \times \mathbb{C}_r) \rightarrow \mathbb{R}$$

It is also common to construct such a costmap by combining several cost functions capturing distinct areas for the robot to avoid. This allows for quick path searches and simple visualizations, but also significantly oversimplifies the problem. A more fitting cost function for navigation may also consider direction of travel or robot velocities. As an example, a cost function could model that moving closer to humans has less costs if done at low speed than at high speed. Such extensions however increase the search space, often leading to combinatorial explosion, causing the robot to be less reactive.

This is a list of cost functions used in the publications, explained in more detail below:

- object padding [44, 91], Figure 9a
- object occlusion [14], Figure 9b
- hidden zones [89], Figure 9c
- zones of high/low noise [58, 59, 60]
- basic comfort distance [43, 44, 54, 83, 89], Figure 9d
- visibility [43, 83, 89], Figure 9e

- interaction regions [79, 83], Figure 9f
- pass on their left [62], Figure 9g
- space ahead for moving [48, 83, 106], Figure 9h
- discard compatible [47, 62]
- crowd density and velocity similarity [37]
- inertia [44]

Object padding [44, 91] as in Figure 9a describes cost regions around static obstacles so that a robot prefers to move at a greater distance from objects than what is strictly necessary for collision avoidance, unless the space is so limited it has to get close. The resulting behavior allows people in the environment to be less worried about the safety of the robot and obstacles.

Object occlusion costs [14] as in Figure 9b try to make the robot avoid motions along paths where the sensors of the robot cannot well perceive the area ahead. This relates in particular to convex corners of rooms, because the robot cannot sense humans behind the corner, and vice versa. Moving around the corner at a greater radius allows earlier coverage of the space behind the corner, which allows for the robot to exhibit human-friendly behavior earlier.

The social costs of hidden zones were described by Sisbot et al. [89] and this refers to the robot avoiding to appear surprisingly from behind objects close to a human, as shown in Figure 9c. This is similar in its intention to object occlusion but requires the robot to know where a present human is. The costs can be discarded if the obstacle does not hide the robot body, even if the obstacle causes robot sensor occlusion.

Costs for object padding, object occlusion and hidden zones cannot replace safety features of the reactive layer of the robot. For purposes of safety, the robot velocity has to be adapted to be safe even when moving tightly around corners or close to obstacles. This velocity adaption remains a necessary feature of the reactive software components (not path planning). The cost functions can only reduce the likelihood of collisions caused by humans, where enough space is available.

The basic comfort distance costs [43, 44, 83, 89] cause a robot to avoid moving closely to humans where that is avoidable, see Figure 9d. The shape and size of such comfort costs can vary according to the context, and Luber et al. [54] even suggest to dynamically change those during a crossing situation.

Visibility costs [43, 83, 89] as in Figure 9e assume that humans prefer a robot to move where they can see the robot when the robot has to move close. This is similar to costs that make the robot approach a human from the front [43]. An opposing force are interaction region costs [79, 83], which assume a human prefers a robot not to cross the space between the human and an object the human interacts with, such as another human or a TV set (Figure 9f). This can be called territoriality, where persons claim free space for themselves. This pair of contradicting forces, visibility and territoriality, illustrates well how difficult it is to tweak and compare cost functions.

To pass humans on their left [62] is a social convention that may be tied to cultural traffic rules. Adding costs on the right of humans as in Figure 9g achieves this behavior in several circumstances like corridors. But it is not clear how these costs influence robot behavior in general.

The other cost functions presented here relate to robot motion among moving humans.

Several publications investigate the effects of defining costs in front of moving humans [48, 83] as in Figure 9h, in particular for non-temporal planning. At the same time, it is also apparent that for non-temporal planning, for humans moving in the same direction as the robot no costs should be applied in path planning [47, 62]. Instead, the robot may benefit from planning a path that follows a person if the robot wants to move in the same direction, in particular for moving in crowds. This relates most closely to natural motion rather than comfort or safety, as certain path deviations that emerge as a result of cost functions applied to an unsuitable context make the robot do unnatural motions. The approach by Ziebart et al. [106] uses an iterative planning approach that may eventually find paths being hindrance-free in time-space. Since the number of iterations is not bounded, this approach must sacrifice either correctness or

reactivity in crowded situations. Crowd motion is also the focus of Henry et al. [37], who model a preference of path steps depending on crowd density and crowd motion.

The actual shape of individual cost functions is subject to tweaking. Most of them have growing costs as the distance to some area decreases. The growth can be modeled linearly, exponentially, sinusoid etc. Also proxemics research has shown that preferred distances between humans and robots depend on many context factors, and it is therefore a valid assumption that this equally holds for other behaviors captured by the social cost functions above.

Similarly, the combination of cost functions is subject to tweaking. Combinations used in publications are weighted sums [44, 83, 89], the maximum [83, 89], and context dependent enabling of costs [44, 47]. In publications cost function shapes and combinations were generally selected and tweaked manually, only few approaches used machine learning techniques [13, 37] for this purpose.

3.4.2. Temporal planning

Apart from work introducing cost functions for path finding algorithms, several authors investigated fundamental changes to path search in order to improve path planning of robots among humans tackling mainly the challenge of avoiding blocking each others path. Temporal planning has two main challenges: 1) predictions of future human motions decrease rapidly in accuracy, and 2) the added dimension of time exponentially increases planning times.

The early paper on human-aware navigation by Tadokoro et al. [92] describes an exploration of the search space using a genetic algorithm rather than classical search. However, they do not motivate that choice nor compare it.

The evacuation simulation planner used by Ohki et al. [65] uses the wavefront algorithms with addition of the dimension of time. The authors do not describe a means to avoid the scalability issues of adding a dimension to the algorithm.

An attempt at temporal global planning avoiding scalability issues is presented by Kushleyev et al. [50]. The idea here is to perform temporal planning for n steps ahead, where the authors suggest n to be determined by the time the prediction algorithm for moving obstacles can return meaningful values. A global planner can then perform static planning for the rest of the way up to the goal until the next replanning cycle.

Because the temporal planning challenge of scalability is so dominant, temporal planning is applied in local planning described in section 3.6 more often than in global planning. The local scope of local planning allows temporal planning to remain tractable.

An attempt to reduce the search space even with temporal planning is presented by Phillips and Likhachev in [75]. The temporal dimension in that work is segmented into safe and unsafe intervals instead of regularly discretized time, reducing the search space expansion. Gonzalez et al. [28] enhance that idea allowing for continous cost functions as well, a prerequisite to apply social cost functions to search.

Before that, the next section describes methods to generate robot behavior considering people as interaction partners, not mere social obstacles.

3.5. Behavior Selection

Behavior is the result of planning and execution. General navigation behavior of a moving robot can be distinguished from the specific behavior of a robot when performing a task. The previous sections described pose selection and path planning. Pose selection and path planning decide ways to move past humans avoiding unnecessary discomfort. Humans are thus treated as social obstacles, having properties to respect beyond those of non-human obstacles. So the interaction between robots and humans is regarded as a necessary evil if the robot has to perform a task, and the aim is to reduce the impact of interaction.

In behavior selection, the problems to be solved more strongly involve a person as interaction partners, relying more on explicit interaction. Therefore interaction is not a necessary evil here, but interaction is the purpose of the robot. Research on navigation then focuses more on how to modulate robot base motion to best serve the interaction goal, rather than merely avoiding discomfort. Partly this reflects the aspect of sociability, partly this reflects the practical needs arising from interaction with a human.

On the more practical side, several authors use state machines to switch between behaviors such as approaching a human, following the human, patrolling, searching, standing in line, and so on [1, 30, 35, 70, 81].

Several authors address adequate approaching behavior for moving humans using a potential field [1, 34]. Althaus et al. [1] define three phases: approaching, keeping the distance, and moving away, for a lengthy situation of interaction. Using potential fields offers some advantages and some disadvantages over search based planning. Advantages over search based planning are that reactions can be found in a more timely fashion and there is no need for a fixed goal pose. Disadvantages are the problem of local minima and issues related to that.

Satake et al. [81] use a custom estimator of whether the robot will be able to approach a walking human from the front. The estimator categorizes the observed motion into one of four categories based on trajectory shape and velocity: fast-walk, idle-walk, wandering and stopping. This allows to estimate whether a suitable meeting point for the robot and the human exists.

The concept of velocity constraint modulation can also be transferred to other areas where velocity limits can be changed, like for allowing increased velocities in corridors. This is done by Zender et al. [105] who use a sparse graph of the environment where navigation points are 1m apart and classified according to laser range data into rooms or corridors.

Similarly for a robot to guide a human, the local planner can be constrained by a model of the joint motion. The work by Feil-Seifer et al. [19] achieves more robust guiding velocities by selecting velocities on the path towards the goal depending on the distance to the goal and the person following. That way, when subjects in experiments were instructed to wait or follow in a slow fashion, the robot adjusted by reducing its own velocity.

With path and behavior selection in place, a robot has a high-level plan of where to move next. What remains is the task to control the robot motors and monitor all sensors to follow this plan while preventing collisions, as described in the next section.

3.6. Local Planning

Local planning is the task of finding commands for the robot actuators for the immediate future. It is also called “reactive planning” and “collision avoidance”. In its usual application, there is nothing particularly human-aware about local-planning, as it merely defines motor responses to perceived static or moving obstacles, regardless of the nature of the obstacle. However some authors use the knowledge about the nature of the obstacle, and derive special behaviors when an obstacle is known to be human, and this allows for human-aware aspects in local planning.

The key responsibility of local planners is to ensure the safety of the robot and the environment. A recent publication by Fraichard [24] defines criteria for safe robot motion, in the sense of collision avoidance. In theory, a navigation framework needs to reason about:

1. its own dynamics,
2. the future dynamics of moving obstacles,
3. this with respect to an infinite time horizon.

For the real world, the time horizon does not need to be infinite, as the value of reasoning about the future quickly deteriorates due to uncertainty and sensor limitations. Hence the requirement can be relaxed for real-world navigation to “with an appropriate lookahead” [5]. What minimum time horizon is appropriate depends on the field-of-view of the robot and the maximally expected dynamics of all objects in the environment.

If one of those criteria is violated, so-called inevitable collision states (ICS) can occur, where a robot will provably enter collisions that it could have avoided by following the criteria. So these criteria form hard constraints to robot navigation. Obeying those alone will not make a robot behave acceptably or efficiently, only safe. Local planners can thus be extended for human environments to create motions that are not just safe, but also adequate for human stakeholders.

There are several kinds of local planners as summarized by Fraichard [24]. The most used ones in our surveyed literature are sampling algorithms and the velocity obstacle approach. In sampling algorithms the

local planner calculates the future positions of the robot under different velocity commands, prunes those that are in collision, and selects the one of the remaining trajectories that best fits a goal function.

The Velocity Obstacle (VO) approach assumes perfect knowledge about the shape of obstacles and other agents' future motions and finds an analytic solution that prevents collisions. It thus works for multi-robot navigation, see [11, 25] for examples. On the other hand it is more difficult to implement analytical solutions with nonlinear trajectories, noisy obstacle data, uncertainty, and cost functions. Therefore, in human-aware navigation research, the sampling approach is more common. Sampling algorithms are easy to extend with social cost functions and additional sensor inputs, but representation of dynamic obstacles is often not part of the solution, as an example the popular Dynamic Window approach [23] considers all obstacles to be static, though it can be extended with object motion predictions.

It is useful to consider that persons in the environment cooperate, and will help avoiding collisions. This can be used for finding solutions with reflective or reciprocal collision avoidance [45, 99]. Such reciprocal schemes assume that the human is aware of the robot and chooses his own trajectory accordingly, instead of blindly or randomly moving on. So in many situations the robot can make smaller evading motions than if it had to prepare for humans moving on carelessly. Making smaller evading motions then makes the robot move more efficiently and more goal directed.

The same basic idea that humans adapt their path as well has been realized by Trautman et al. [98], who suggest a solution using an Interaction Gaussian Process Model. If the prediction of the future human path is represented by growing regions such as shown in Figure 8d, and the planner is forbidden to enter such regions, then the robot may be unable to find any allowed space left, and not move on. Using Interaction Gaussian Process Model to model human motion, the predictions take into account that the robot is being perceived, thus possibly opening spaces in human formations for the robot to move through.

A similar attempt is made by Althoff et al. [2] to consider uncertainty in bilateral sampling and selection of most likely collision free trajectories. While the concept is promising, the approach requires a lot of processing and was not achieved in realtime.

So human-awareness is represented in localplanners mostly by interleaving prediction of human motion and planning or robot motion to identify potentials for joint efforts in collision avoidance.

3.7. Feature integration and Synthesis

Designing a suitable control flow for the functionalities of Figure 6 is not straightforward. The main challenge is that when humans are present, the base assumption needs to be that the world is changing all the time unpredictably. So any plan under execution is outdated immediately after any unpredicted change in the world. Several of the modules in our generalized navigation framework require significant planning time on their own, let alone when processed in sequence. This performance challenge is not specific to human-aware navigation, and outside the scope of this paper. The publications on human-aware navigation do not go into detail on this aspect.

Apart from the performance challenge, the semantic integration of modules is also not straightforward. As an example, one can imagine a pose planner that finds a very comfortable pose for the robot to hand over an item to a person, and a path planner that finds a very comfortable path to that pose. When the robot executes the plan, it is possible that a behavior results that is confusing to the human. That's because it is natural to select a pose for handover in accordance with the shortest path to that pose. Humans may sacrifice some pose comfort to reduce the path length and reduce the duration of the interaction. This split between pose and path selection can be avoided in different ways, for example by using potential fields [1, 34], but the underlying problem of suboptimal global solutions with separate planners surfaces in many ways.

Another separation to be merged is between prediction and planning (both global and local), as integrating predictions in the planning allows to take into account that a person may adapt to the robot and autonomously avoid collisions [2, 45, 99], or may cooperatively help the robot.

This concludes the examination of functionalities to solve challenges of human-aware navigation. The remainder of the paper investigates evaluation methods before ending with a conclusion.

4. Evaluation Methods for Human-aware Navigation

The focus of research in human-aware navigation are the algorithms that produce the decisions that produce the motions. In order to scientifically compare algorithms, measures are required that describe properties of robot motion.

It is comparatively easy to define measures of efficiency and safety, measuring time, energy consumption, path lengths and collisions. This is mostly similar outside HRI, so this paper won't go into detail about it. The difficult kind of evaluation is to analyze the impacts of a changed algorithm of an autonomous robot on human comfort, naturalness and sociability.

All three notions are very subjective and context dependent. Methods from neuroscience and psychology provide direct measures such as gaze behavior or skin conductance and indirect feedback in the form of questionnaires or interviews. By nature of the intrinsic variance and uncertainty, experiments with humans are restricted to well-defined, repeatable contexts. If an improvement is shown in one context, that is already a small scientific gain. But this does not allow inferring a blueprint for a robot, in particular since improvements of different research efforts cannot well be compared to each other, as there is no scale for comparison.

Another reason for the difficulty of evaluating the impacts of different algorithms on comfort is more practical: The time and preparation required to prepare a lab and a robot as well as human participants to perform experiments giving meaningful results. Getting meaningful results is not just a matter of taking great care in setting up the experiment, it also involves considering effects of adaptation and fatigue of human participants.

As a result, there are two main evaluation strategies. One strategy is to put a real robot in a human environment and create situations of interaction. This approach provides empirical data from robot measurements and ratings of participants or observers. The other strategy is to run automated experiments in a simulator using measures assumed to influence comfort, showing that if that assumption were true, the algorithm performs better. Both approaches have their merits and caveats.

4.1. Simulation based evaluation

Using simulation for evaluation, the common property used to approximate discomfort is the distance a simulated robot kept from a simulated human [44, 65, 92] for running simulations in special situations.

Some research also focuses on Wizard of Oz studies to provide descriptions of robot behavior that cause more or less discomfort [10, 46, 66], though the results may still be too few and specialized to predict human discomfort in general situations.

Several authors chose to perform an evaluation in simulation [38, 43, 62], along with a demonstration on a real robot.

4.2. User study based evaluation

In life experiments, a robot with given algorithms performs tasks, and human participants experiencing the robot rate the performance of the robot on scales, such as a Likert scale expressing how comfortable they felt during interaction.

A good example for research using direct interaction is the work by Pacchierotti et al. [67, 68]. In the studies a robot moved along a hallway with participants moving in the other direction. The robot behavior could be modified by parameters, and the study found correlations between parameters and user ratings.

Another example of this approach is shown in Young et al. [103]. While evaluating a dogleash system for robots, the robot behavior deviated between following the human or walking in front.

The previous studies have been performed in a lab environment, a similar study in a mall is presented by Hayashi et al. [35]. In that work, the robot behavior was based on two parameters, and again participant ratings were compared to parameter settings after a series of trials.

Typically the participants need to be informed about the experiment before it starts, and the participants are aware of the experimenter observing them. This can strongly bias the questionnaire. In the work by Satake et al. [81] the participant behavior in public places was monitored by TV cameras, so that the participants were unaware of being part of an experiment. The task of the robot was to engage participants

in a trivial conversation, and the evaluation was based on the situations in the video data being labeled by humans as success or failure of the robot to engage in a conversation. Video-based approaches often produce similar results to live experiments, as found by Woods et al. [102]. The video-based approach was also used by Lichtenthahler et al. [53] to show an improvement of robot legibility using human-aware motion-planning.

In contrast to the previous works, Gockley et al. [27] use participants who observe the robot interacting with an experimenter.

A look at long-term changes of human preferences with a robot is provided by Walters et al. [100]. It shows an adaptation effect, as over time shorter distances become more acceptable to the subjects.

4.3. Summary of evaluation methods

The most problematic area of evaluations in the context of human-aware navigation is the evaluation of comfort improvements due to changes in navigational decision making.

The literature shows two main strategies in doing so, using a simulation based on a model of what causes human-discomfort, or presenting participants with a robot and asking them to rate robot qualities via a questionnaire.

5. Conclusion

This survey presented the state-of-the-art in human-aware navigation for wheeled robots in assistive scenarios. The features that allow for human-aware navigation can mostly be categorized into the topics of comfort, naturalness and sociability of motions.

We have shown that for each feature category, a diversity of methods were explored for a variety of problems. The progress made so far allows building robot navigation systems that demonstrably avoids several behavior flaws that non-human-aware navigation exhibits.

The publications we found deal with individual domains and challenges, a holistic theory of human-aware navigation was not attempted yet. So one main challenge in the research area is to unify different methods and solutions technically and semantically. This unification requires a grounding of methods in a semantic context of behavior. Such a context needs to define each spatial human-robot encounter as an interaction following social rules. To find and evaluate such rules, and to map them to suitable software processes remains the main challenge of human-aware navigation planning.

Acknowledgement

With the support of the Technische Universitat Munchen - Institute for Advanced Study and the Bavarian Academy of Sciences and Humanities. Funded by the German Excellence Initiative and the CoTeSys (Cognition for Technical Systems) cluster of excellence. This work has been conducted within the EU SAPHARI project (<http://www.saphari.eu/>) funded by the E.C. Division FP7-IST under Contract ICT-287513.

References

- [1] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. Christensen. Navigation for human-robot interaction tasks. In *ICRA, IEEE*, 2004.
- [2] D. Althoff, D. Wollherr, and M. Buss. Safety assessment of trajectories for navigation in uncertain and dynamic environments. In *ICRA, IEEE*, 2011.
- [3] G. Arechavaleta, J.-P. Laumond, H. Hicheur, and A. Berthoz. On the nonholonomic nature of human locomotion. *Autonomous Robots*, 25(1-2):25–35, 2008.
- [4] K. O. Arras, N. Tomatis, and R. Siegwart. Robox, a remarkable mobile robot for the real world. In *ISER*, pages 178–187, 2002.
- [5] A. Bautin, L. Martinez-Gomez, and T. Fraichard. Inevitable collision state: A probabilistic perspective. In *ICRA, IEEE*, pages 4022–4027, 2010.
- [6] M. Bennewitz. *Mobile Robot Navigation in Dynamic Environments*. PhD thesis, Albert-Ludwigs-Universitat Freiburg, 2004.
- [7] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *International Journal of Robotics Research*, 24(1), 2005.

- [8] O. Brock and O. Khatib. Elastic strips: A framework for motion generation in human environments. *International Journal of Robotics Research*, 21(12):1031–1052, 2002.
- [9] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2):3–55, 1999.
- [10] J. T. Butler and A. Agah. Psychological effects of behavior patterns of a mobile personal robot. *Autonomous Robots*, 10(2):185–202, 2001.
- [11] N. Chan, J. Kuffner, and M. Zucker. Improved motion planning speed and safety using regions of inevitable collision. In *17th CISM-IFToMM Symposium on Robot Design, Dynamics, and Control (RoManSy'08)*, July 2008.
- [12] S.-Y. Chung and H.-P. Huang. A mobile robot that understands pedestrian spatial behaviors. In *IROS, IEEE/RSJ*, 2010.
- [13] S.-Y. Chung and H.-P. Huang. Incremental learning of human social behaviors with feature-based spatial effects. In *IROS, IEEE/RSJ*, 2012.
- [14] W. Chung, S. Kim, M. Choi, J. Choi, H. Kim, C. bae Moon, and J.-B. Song. Safe navigation of a mobile robot considering visibility of environment. *Industrial Electronics, IEEE Transactions on*, 56(10):3941–3950, October 2009.
- [15] A. Clodic, S. Fleury, R. Alami, R. Chatila, G. Bailly, L. Brthes, M. Cottret, P. Dans, X. Dollat, F. Elise, I. Ferran, M. Herrb, G. Infantes, C. Lemaire, F. Lerasle, J. Manhes, P. Marcoul, P. Menezes, and U. P. Sabatier. Rackham: An interactive robot-guide. In *RO-MAN, IEEE*, 2006.
- [16] N. P. Cuntoor, R. Collins, and A. J. Hoogs. Human-robot teamwork using activity recognition and human instruction. In *IROS, IEEE/RSJ*, 2012.
- [17] K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, E. A. Sisbot, R. Alami, and T. Siméon. How may i serve you?: a robot companion approaching a seated person in a helping context. In *HRI, ACM/IEEE*, Utah, USA, 2006.
- [18] D. Feil-Seifer and M. J. Matarić. Human robot interaction. In R. A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 4643–4659. Springer New York, 2009.
- [19] D. J. Feil-Seifer and M. J. Matarić. People-aware navigation for goal-oriented behavior involving a human partner. In *Proceedings of the International Conference on Development and Learning*, Frankfurt am Main, Germany, August 2011.
- [20] F. Feurtey. Simulating the collision avoidance behavior of pedestrians. Master’s thesis, University of Tokyo, Department of Electronic Engineering, February 2000.
- [21] A. Foka and P. Trahanias. Predictive autonomous robot navigation. In *IROS, IEEE/RSJ*, 2002.
- [22] T. Fong, I. R. Nourbakhsh, and K. Dautenhahn. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42, 2003.
- [23] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 1997.
- [24] T. Fraichard. A short paper about motion safety, 2007.
- [25] O. Gal, Z. Shiller, and E. Rimon. Efficient and safe on-line motion planning in dynamic environments. In *ICRA, IEEE*, 2009.
- [26] A. Garrell and A. Sanfeliu. Local optimization of cooperative robot movements for guiding and regrouping people in a guiding mission. In *IROS, IEEE/RSJ*, 2010.
- [27] R. Gockley, J. Forlizzi, and R. Simmons. Natural person-following behavior for social robots. In *HRI, ACM/IEEE*, 2007.
- [28] J. P. Gonzalez, A. Dornbush, and M. Likhachev. Using state dominance for path planning in dynamic environments with moving obstacles. In *ICRA, IEEE*, 2012.
- [29] M. A. Goodrich and A. C. Schultz. Human-robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction*, 1(3):203–275, 2007.
- [30] C. Granata and P. Bidaud. A framework for the design of person following behaviours for social mobile robots. pages 4652–4659, 2012.
- [31] H.-M. Gross, H. Boehme, C. Schröter, S. Müller, A. Koenig, E. Einhorn, C. Martin, M. Merten, and A. Bley. Toomas: Interactive shopping guide robots in everyday use - final implementation and experiences from long-term field trials. In *IROS, IEEE/RSJ*, 2009.
- [32] E. Hall. *The hidden dimension*. Anchor Books, 1966.
- [33] M. Hanheide, A. Peters, and N. Bellotto. Analysis of human-robot spatial behaviour applying a qualitative trajectory calculus. In *RO-MAN, IEEE*, 2012.
- [34] S. T. Hansen, M. Svenstrup, H. J. Andersen, and T. Bak. Adaptive human aware navigation based on motion pattern analysis. In *Robot and Human Interactive Communication, 2009.*, Toyama, Japan, Sept.-Oct. 2009.
- [35] K. Hayashi, M. Shiomi, T. Kanda, and N. Hagita. Friendly patrolling: A model of natural encounters. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2011.
- [36] D. Helbing. A mathematical model for the behavior of pedestrians. *BEHAVIORAL SCIENCE*, 36:298–310, 1991.
- [37] P. Henry, C. Vollmer, B. Ferris, and D. Fox. Learning to navigate through crowded environments. In *ICRA, IEEE*, 2010.
- [38] F. Hoeller, D. Schulz, M. Moors, and F. E. Schneider. Accompanying persons with a mobile robot using motion prediction and probabilistic roadmaps. In *IROS, IEEE/RSJ*, 2007.
- [39] K.-C. Huang, J.-Y. Li, and L.-C. Fu. Human-oriented navigation for service providing in home environment. In *SICE Annual Conference 2010, Proceedings of*, August 2010.
- [40] B. Jensen, N. Tomatis, L. Mayor, A. Drygajlo, and R. Siegwart. Robots meet humans interaction in public spaces. *IEEE Transactions on Industrial Electronics*, 52(6), December 2005.
- [41] E.-J. Jung, B.-J. Yi, and S. Yuta. Control algorithms for a mobile robot tracking a human in front. In *IROS, IEEE/RSJ*, 2012.
- [42] J. Kessler, D. Iser, and H.-M. Gross. I’ll keep you in sight: Finding a good position to observe a person. In *IROS*,

- IEEE/RSJ, 2012.
- [43] J. Kessler, C. Schröter, and H.-M. Gross. Approaching a person in a socially acceptable manner using a fast marching planner. In *ICIRA (2)*, pages 368–377, 2011.
- [44] R. Kirby, R. Simmons, and J. Forlizzi. Companion: A constraint optimizing method for person-acceptable navigation. In *RO-MAN, IEEE*, 2009.
- [45] B. Kluge and E. Prassler. Reflective navigation: Individual behaviors and group behaviors. In *ICRA*, pages 4172–4177, 2004.
- [46] K. L. Koay, E. A. Sisbot, D. S. Syrdal, M. L. Walters, K. Dautenhahn, and R. Alami. Exploratory study of a robot approaching a person in the context of handing over an object. In *AAAI Spring Symposium: Multidisciplinary Collaboration for Socially Assistive Robotics*, pages 18–24, 2007.
- [47] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch. Legible robot navigation in the proximity of moving humans. In *Advanced Robotics and its Social Impacts (ARSO), 2012 IEEE Workshop on*, pages 83–88, May 2012.
- [48] T. Kruse, A. Kirsch, E. A. Sisbot, and R. Alami. Dynamic generation and execution of human aware navigation plans. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2010.
- [49] T. Kruse, A. Kirsch, E. A. Sisbot, and R. Alami. Exploiting human cooperation in human-centered robot navigation. In *RO-MAN, IEEE*, 2010.
- [50] A. Kushleyev and M. Likhachev. Time-bounded lattice for efficient planning in dynamic environments. In *ICRA, IEEE*, 2009.
- [51] C.-P. Lam, C.-T. Chou, K.-H. Chiang, and L.-C. Fu. Human-centered robot navigation - towards a harmoniously human-robot coexisting environment. *Robotics, IEEE Transactions on*, 27(1):99–112, February 2011.
- [52] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, MA, 1991.
- [53] C. Lichtenthaler, T. Lorenz, and A. Kirsch. Influence of legibility on perceived safety in a virtual human-robot path crossing task. In *RO-MAN, IEEE*, 2012.
- [54] M. Luber, L. Spinello, J. Silva, and K. O. Arras. Socially-aware robot navigation: A learning approach. In *IROS, IEEE/RSJ*, 2012.
- [55] I. Macaluso, E. Ardizzone, A. Chella, M. Cossentino, A. Gentile, R. Gradino, I. Infantino, M. Liotta, R. Rizzo, and G. Scardino. Experiences with cicerobot, a museum guide cognitive robot. In *AI*IA 2005*, S. Bandini and S. Manzoni Eds, pages 474–482. Springer-Verlag, 2005.
- [56] J. Mainprice, M. Gharbi, T. Siméon, and R. Alami. Sharing effort in planning human-robot handover tasks. In *RO-MAN, IEEE*, 2012.
- [57] E. A. Martinez-Garcia, O. Akihisa, and S. Yuta. *Crowding and guiding groups of humans by teams of mobile robots*, pages 91–96. 2005.
- [58] E. Martinson. *Acoustical awareness for intelligent robotic action*. PhD thesis, Atlanta, GA, USA, 2007. AAI3294521.
- [59] E. Martinson. Hiding the acoustic signature of a mobile robot. In *IROS, IEEE/RSJ*, 2007.
- [60] E. Martinson and D. Brock. Improving human-robot interaction through adaptation to the auditory scene. In *HRI, ACM/IEEE*, New York, NY, USA, 2007. ACM.
- [61] M. Mori. Bukimi no tani [the uncanny valley]. *Energy*, 7(4):33–35, 1970.
- [62] J. Müller, C. Stachniss, K. O. Arras, and W. Burgard. Socially inspired motion planning for mobile robots in populated environments. In *International Conference on Cognitive Systems (CogSys'08)*, Karlsruhe, Germany, 2008.
- [63] J. Mumm and B. Mutlu. *Human-robot proxemics: Physical and psychological distancing in human-robot interaction*. In *HRI, ACM/IEEE*. ACM, ACM Press, 2011.
- [64] I. Nourbakhsh, C. Kunz, and T. Willeke. The mobot museum robot installations: A five year experiment, 2003.
- [65] T. Ohki, K. Nagatani, and K. Yoshida. Collision avoidance method for mobile robot considering motion and personal spaces of evacuees. In *IROS, IEEE/RSJ*, 2010.
- [66] A. Oosterhout, van T.; Visser. A visual method for robot proxemics measurements. In *Proceedings of Metrics for Human-Robot Interaction: A workshop at HRI, ACM/IEEE*, p.61-68., 2008.
- [67] E. Pacchierotti, H. Christensen, and P. Jensfelt. Embodied social interaction in hallway settings: a user study. In *RO-MAN, IEEE*, 2005.
- [68] E. Pacchierotti, H. Christensen, and P. Jensfelt. Evaluation of passing distance for social robots. In *RO-MAN, IEEE*, September 2006.
- [69] A. Pandey and R. Alami. A framework for adapting social conventions in a mobile robot motion in human-centered environment. In *Advanced Robotics, 2009. (ICAR)*, 2009.
- [70] A. Pandey and R. Alami. A step towards a sociable robot guide which monitors and adapts to the person’s activities. In *Advanced Robotics, 2009. (ICAR)*, pages 1–8, June 2009.
- [71] A. Pandey and R. Alami. A framework towards a socially aware mobile robot motion in human-centered dynamic environment. In *IROS, IEEE/RSJ*, October 2010.
- [72] A. Peters, T. P. Spexard, P. Weiß, and M. Hanheide. Make room for me - a spatial and situational movement concept in hri. *Workshop on Behavior Monitoring and Interpretation*, September 2009.
- [73] R. Philippsen. *Motion Planning and Obstacle Avoidance for Mobile Robots in Highly Cluttered Dynamic Environments*. PhD thesis, Ecole Polytechnique Fédérale de Lausanne, 2004.
- [74] R. Philippsen and R. Siegwart. Smooth and efficient obstacle avoidance for a tour guide robot. In *ICRA, IEEE*, 2003.
- [75] M. Phillips and M. Likhachev. Sipp: Safe interval path planning for dynamic environments. In *ICRA, IEEE*, 2011.
- [76] E. Prassler, D. Bank, and B. Kluge. Key technologies in robot assistants: Motion coordination between a human and a mobile robot. *Transactions on Control, Automation and Systems Engineering*, 4, 2002.
- [77] E. Prassler, J. Scholz, and P. Fiorini. Navigating a robotic wheelchair in a railway station during rush hour. *I. J. Robotic*

- Res., 18(7):711–727, 1999.
- [78] J. Rios-Martinez. Socially-Aware Robot Navigation: combining Risk Assessment and Social Conventions. PhD thesis, Université de Grenoble, January 2013.
 - [79] J. Rios-Martinez, A. Renzaglia, A. Spalanzani, A. Martinelli, and C. Laugier. Navigating between people: A stochastic optimization approach. In ICRA, IEEE, 2012.
 - [80] S. Russell and P. Norvig. Artificial Intelligence - A Modern Approach. Prentice Hall, Upper Saddle River, New Jersey, 2003.
 - [81] D. F. G. M. I. H. I. N. H. S. Satake, T. Kanda. How to approach humans? strategies for social robots to initiate interaction. In HRI, ACM/IEEE, 2009.
 - [82] P. Saulnier, E. Sharlin, and S. Greenberg. Exploring minimal nonverbal interruption in social hri. In RO-MAN, IEEE, August 2011.
 - [83] L. Scandolo and T. Fraichard. An anthropomorphic navigation scheme for dynamic scenarios. In ICRA, IEEE, February 2011.
 - [84] D. Shi, E. G. Collins Jr., B. Goldiez, X. Liu, A. Donate, and D. Dunlap. Human-aware robot motion planning with velocity constraints. 2008 International Symposium on Collaborative Technologies and Systems, pages 490–497, 2008.
 - [85] B. Siciliano and O. Khatib, editors. Springer Handbook of Robotics, chapter 57, 58. Springer, Berlin, Heidelberg, 2008.
 - [86] R. Siegwart, K. O. Arras, B. Jensen, R. Philippsen, and N. Tomatis. Design, implementation and exploitation of a new fully autonomous tour guide robot. In Proc. 1st International Workshop on Advances in Service Robotics (ASER'03), Bardolino, Italy, 2003.
 - [87] E. A. Sisbot. Towards Human-Aware robot Motions. PhD thesis, LAAS/CNRS, Universite Paul Sabatier, October 2008.
 - [88] E. A. Sisbot, R. Alami, K. D. T. Simeon, M. Walters, S. Woods, K. L. Koay, and C. Nehaniv. Navigation in the presence of humans. In IEEE-RAS International Conference on Humanoid Robots (Humanoids 05), 2005.
 - [89] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon. A human aware mobile robot motion planner. IEEE Transactions on Robotics, 23, 2007.
 - [90] M. Svenstrup. Advances in Robot Navigation, chapter Navigation Among Humans, pages 133–158. InTech, 2011.
 - [91] M. Svenstrup, T. Bak, and H. J. Andersen. Trajectory planning for robots in dynamic human environments. In IROS, IEEE/RSJ, October 2010.
 - [92] S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which coexist and cooperate with human. IROS, IEEE/RSJ, 1995.
 - [93] L. Takayama and C. Pantofaru. Influences on proxemic behaviors in human-robot interaction. In IROS, IEEE/RSJ, 2009.
 - [94] Y. Tamura, P. D. L. abd Kentarou Hitomi, N. P. Chandrasiri, T. Bando, A. Yamashita, and H. Asama. Development of pedestrian behavior model taking account of intention. In IROS, IEEE/RSJ, 2012.
 - [95] S. Thompson, T. Horiuchi, and S. Kagami. A probabilistic model of human motion and navigation intent for mobile robot path planning. 4th International Conference on Autonomous Robots and Agents, 2009.
 - [96] S. Thrun, M. Beetz, M. Bennis, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. R. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. I. J. Robotic Res., 19(11):972–999, 2000.
 - [97] G. D. Tipaldi and K. O. Arras. Please do not disturb! minimum interference coverage for social robots. In IROS, IEEE/RSJ, 2011.
 - [98] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In IROS, IEEE/RSJ, October 2010.
 - [99] J. P. van den Berg, M. C. Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In ICRA, IEEE, 2008.
 - [100] M. Walters, M. Oskoei, D. Syrdal, and K. Dautenhahn. A long-term human-robot proxemic study. In RO-MAN, IEEE, 2011.
 - [101] M. L. Walters, K. Dautenhahn, R. Te Boekhorst, K. L. Koay, D. S. Syrdal, and C. L. Nehaniv. An empirical framework for human-robot proxemics. Human-Computer Interaction, 2009.
 - [102] S. N. Woods, M. L. Walters, K. L. Koay, and K. Dautenhahn. Methodological issues in hri: A comparison of live and video-based methods in robot to human approach direction trials. In RO-MAN, IEEE, 2006.
 - [103] J. Young, Y. Kamiyama, J. Reichenbach, T. Igarashi, and E. Sharlin. How to walk a robot: A dog-leash human-robot interface. In RO-MAN, IEEE, 2011.
 - [104] F. Yuan, L. Twardon, and M. Hanheide. Dynamic path planning adopting human navigation strategies for a domestic mobile robot. In IROS, IEEE/RSJ, October 2010.
 - [105] H. Zender, P. Jensfelt, and G.-J. M. Kruijff. Human- and situation-aware people following. In RO-MAN, IEEE, 2007.
 - [106] B. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. D. Bagnell, M. Hebert, A. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In IROS, IEEE/RSJ, October 2009.