Toward Seamless Human–Robot Handovers

Kyle Strabala, Min Kyung Lee, Anca Dragan, Jodi Forlizzi,
Siddhartha S. Srinivasa

Carnegie Mellon University

Maya Cakmak

Willow Garage

and

Vincenzo Micelli

Università Degli Studi di Parma

A handover is a complex collaboration, where actors coordinate in time and space to transfer control of an object. This coordination comprises two processes: the physical process of moving to get close enough to transfer the object, and the cognitive process of exchanging information to guide the transfer. Despite this complexity, we humans are capable of performing handovers seamlessly in a wide variety of situations, even when unexpected. This suggests a common procedure that guides all handover interactions. Our goal is to codify that procedure.

To that end, we first study how people hand over objects to each other in order to understand their coordination process and the signals and cues that they use and observe with their partners. Based on these studies, we propose a coordination structure for human–robot handovers that considers the physical and social-cognitive aspects of the interaction separately. This handover structure describes how people approach, reach out their hands, and transfer objects while simultaneously coordinating the *what*, *when*, and *where* of handovers: to agree that the handover will happen (and with what object), to establish the timing of the handover, and to decide the configuration at which the handover will occur. We experimentally evaluate human–robot handover behaviors that exploit this structure and offer design implications for seamless human–robot handover interactions.

Keywords: Physical human–robot interaction, handover, signaling, joint activity

1. Introduction

Handovers are abundant and varied in our daily lives: a care-giver bringing a patient a glass of water, a mechanic receiving a tool from his assistant, someone passing a bucket of water as part of a fire brigade, and a man on the sidewalk handing over a flyer to a busy passer-by. The people in these examples are all participating in complex coordinations to ensure that control of the object

Authors retain copyright and grant the Journal of Human-Robot Interaction right of first publication with the work simultaneously licensed under a Creative Commons Attribution License that allows others to share the work with an acknowledgement of the work's authorship and initial publication in this journal.

is successfully transferred from one person to the other. These four handover examples are used throughout this paper and are explored in detail in Section 4.

People rely on knowledge of context and communication cues for successful handovers. Context gives people knowledge on what to expect from other people and how to interpret their behaviors. Given a context, people use communication cues to establish the *what*, *when*, and *where/how* of the handover. For example, the mechanic establishes *what* by asking for a tool and using context (the assistant is nearby and it is the assistant’s role to provide the tools) to expect a handover. The care-giver and the patient exchange looks to establish *when* they are ready to reach out and transfer the glass. The man handing out flyers holds the flyers close when nobody is near, then establishes *where/how* by reaching out a hand toward passers-by when they approach.

Our key contribution is to extract this crucial handover structure based on our studies of human-to-human handovers (Sections 3 and 4). We present the first holistic representation of handovers, which comprises physical-level and social/cognitive-level coordination behaviors. This handover structure describes how people approach, reach out their hands, and transfer objects while simultaneously coordinating the *what*, *when*, and *where* of handovers. This enables us to study our handover examples through a common unified lens. To accomplish this, our research group performed seven studies encompassing human observation and a laboratory study of both human pairs and robot–human teams to understand, build, and evaluate human–robot handover behavior.

We first study human–human handovers (Section 3) and use them to extract the structure we present in Section 4. Our most unexpected observation is that the intent to do a handover is often established much before the hand even starts moving. According to our studies on human–human interactions, handovers can actually be *predicted* in advance in 89% of the cases (Section 3.2). This is of utmost importance for the robot to be prepared for the handover and also for the robot to take control of the interaction, by signaling its readiness or even its unavailability when it decides to not perform the handover.

By studying this handover coordination, we extract a structure for handovers (Section 4). The handover structure motivates our design of robot behaviors for human–robot handovers (Section 5). We use an iterative, user-centered design process to design human–robot behaviors for handovers. We evaluate what handover configurations a robot should use for the *where* of a handover when taking people’s preferences into account (Studies 1, 2, and 3 in Section 5). We further propose methods for negotiating and adapting the *where* based on the humans’ behavior when the robot is in an active receiver role as opposed to the usual giver role (Study 5 in Section 5). We also show that contrast can be used to establish the *when* of a handover — to signal that the robot is ready by transitioning from a distinct carrying pose to a distinct handover pose (Study 4 in Section 5).

We are excited about pushing robots out of the era of “handing over in the dark” and into an era of personal assistants that coordinate handovers to achieve seamless interaction. We discuss the limitations of our work and our recommendations for key design decisions for robots to perform handovers with people (Section 6).

2. Related Work

2.1 Studies of Human–Human Handovers

Human–human handovers are the ultimate benchmark for seamless handovers. Handovers between two humans have been studied in the literature with an eye toward implications for robot–human handovers. Trajectories and velocity profiles adopted by humans in the role of both giver and receiver are analyzed in Shibata, Tanaka, and Shimizu (1995). Simulation results for a controller that mimics the characteristics of human handovers are presented in Kajikawa, Okino, Ohba, and Inooka (1995). The social modification of pick-and-place movements is demonstrated in Becchio, Sartori,

and Castiello (2010) comparing velocity profiles for placing an object on a container versus another person’s palm. Chan, Parker, Loos, and Croft (2012) investigated the grip forces applied on an object while it is transferred between two humans during a handover. Based on the observed patterns, they proposed a model for a similar handover behavior for robots.

Huber, Rickert, Knoll, Brandt, and Glasauer (2008) analyzed the efficiency of handovers in terms of the durations of three phases during a handover and compared human–human handovers with robot–human handovers. Basili, Huber, Brandt, Hirche, and Glasauer (2009) analyzed human approach and handover and observed a preparatory movement of lifting the object before the handover, which might play an important role in signaling the timing of the handover.

Our work contributes to this literature on human–human handovers with the study presented in Section 3. This study is the basis of our holistic handover structure. Although previous work focused on determining parameters of a handover, such as the choice of trajectory profiles or durations of different phases, our emphasis is on the exchange of social cues that play a crucial role in coordinating the handover.

2.2 Handover Behaviors for Robots

It is important for robots to carry out handovers autonomously. This involves a number of steps that each can require certain perception, planning, and decision inference capabilities. Much of the studies on human–robot handovers in the literature has proposed methods for providing these capabilities.

The first step in the handover process is to approach the receiver. Sisbot et al. (2005) proposed a navigation planner that creates safe, legible, and socially acceptable paths for approaching a person for a handover. As a progression from this work, the motion planner was extended to account for the amount of human motion required to perform a handover, allowing the robot to choose the best handover location based on context (Mainprice, Gharbi, Simeon, & Alami, 2012).

An important decision for the robot is how the object will be presented to the receiver. For this, Sisbot and Alami (2012) developed a manipulation planning framework that chooses handover locations based on the human’s safety, accessibility, field of view, posture, and preferences. A handover motion controller that adapts to unexpected arm movements of a simulated human is presented in Agah and Tanie (1997). Glasauer, Huber, Basili, Knoll, and Brandt (2010) investigated how a robot can use human-like reaching gestures to convey the intent to do a handover and signal readiness. Kim, Park, Hwang, and Lee (2004) investigated the mechanism of how a robot can grasp an object before handing it over to a human that incorporates the object’s shape, the object’s function, and the safety of both the robot and the human. Similarly, Lopez-Damian, Sidobre, DeLaTour, and Alami (2006) presented a planner to grasp unknown arbitrary objects for interactive manipulation tasks.

Although certain handovers can be planned in advance, the robot needs to be responsive and dynamic during the interaction, particularly during the transfer of the object. Nagata, Oosaki, Kakikura, and Tsukune (1998) presented a grasping system based on force and torque feedback that senses when the humans has a stable grasp on the object, after which the robot can release the object. Sadigh and Ahmadi (2009) presented a robotic grasping controller inspired by human grasping to grasp an object with minimal normal forces while ensuring the object does not slip.

Our work on human–robot handovers also contributes to the design of autonomous handover behaviors. Our findings about how an object should be presented in order to convey intent (Section 5.1.1) and be preferable by humans (Section 5.1.3) can be incorporated into the objective functions of the motion planners proposed in the literature. Other work focusing on navigation and grip forces is complementary to this work. Also, all of the works mentioned above consider the transfer

of objects from robots to humans. One of our studies (Section 5.3) is the first, to our knowledge, to focus completely on the transfer of objects from a human to an adaptive robot.

2.3 Studies of Human–Robot Handovers

We believe that user studies involving actual human–robot handovers are invaluable in guiding the design of seamless handover interactions. One benefit of user studies is to establish human preferences. A study presented in Koay et al. (2007) analyzes human preferences on the robot’s handover behaviors in terms of the approach direction and of the height and distance of the object. User preferences between two velocity profiles for handing over are analyzed in Huber et al. (2008) in terms of participants’ ratings of human-likeness and feelings of safety. Our studies, presented in Section 5, complement these by focusing on human preference for how the object and the robot are configured during the object transfer.

In addition, user studies can provide important observations that guide the design of handover behaviors. Edsinger and Kemp (2007) presented a study that demonstrated the effectiveness of a simple handover mechanism that automatically drops the object without any sensing. The authors found that, during a handover, humans will pose an object in the robot’s stationary hand regardless of the robot’s hand pose, demonstrating that humans adapt to the robot’s hand pose. Aleotti, Micelli, and Caselli (2012) found that robots should take into consideration how the human will grasp the object and that robots should present the human with the most appropriate part of the object (e.g., a handle). Pandey, Ali, Warnier, and Alami (2011) investigated how a robot can predict where the human will hand over and then proactively move to this location.

Although handovers are a unique type of interaction in several aspects, certain findings from other fields of human–robot interaction (HRI) research, such as proxemics or social communication, have strong implications on how handover interactions should be designed. For instance, Satake et al. (2009) proposed a navigation model for how a robot should approach a human to make initiating an interaction easier. Although this was specific to starting a conversation, it could easily be applied to handover interactions. Takayama and Pantofaru (2009) explored how personal space varies when approaching and being approached by a robot based on the human’s experience with robots and where the robot looks during the approach. Another study explored how proxemics varies with a robot’s likability and eye gaze (Mumm & Mutlu, 2011). Such factors should be considered when designing handover behaviors in order to respect the human’s personal space.

As illustrated in our human–human handover studies (Section 3), gaze and eye contact are important cues for coordinating a handover. On human-like robots, such cues can be exploited to allow seamless handovers. Previous work on social gaze for robots provides guidance on how this could be achieved (Mutlu, Yamaoka, Kanda, Ishiguro, & Hagita, 2009). Another work on human–robot collaboration is relevant for our work as it establishes metrics of interaction fluency, which can be used in measuring the success of handovers (Hoffman & Breazeal, 2007). This work reveals that anticipatory agents provide more fluent interactions than reactive agents, highlighting the importance of being able to accurately predict human actions.

3. Human–Human Handovers Exhibit Structure and Communication

Handovers are complex interactions, yet humans are capable of performing handovers seamlessly and without conscious thought. This suggests people share a common procedure that guides the handover interaction. To learn about this possible structure, we analyzed how people hand over objects to each other. In this section, we summarize two studies from our previous work. In the first study, we found a structure consisting of carrying (approaching with the object), signaling readiness to do a handover, and transferring the object. Although all three components happen in the physical channel, signaling has a complex social-cognitive channel as well. To explore this complexity, we

ran a second study to analyze the cues used to signal readiness. We found that, in many cases, the exact time when an actor starts reaching can be predicted (in 89% of the cases) from communication cues that the actor uses right before the act, meaning the communication between humans is so rich that signaling readiness to do a handover can happen *before* either actor starts reaching out.

3.1 Observation of Human–Human Handover Interactions

In our first study (Lee, Forlizzi, Kiesler, Cakmak, & Srinivasa, 2011), we observed five pairs of participants handing objects to one another, as shown in Figure 1. The participants took turns being a care-giver and care-receiver, and the care-giver brought objects to the care-receiver. The care-receiver was either sitting on a chair and reading a magazine, or standing and packing a box on a table. For each of two scenarios, the participants transferred 10 objects, resulting in 40 trials for each pair. The session took about 30 min to complete and was videotaped.

Participants. We recruited participants on a university study participant recruiting website. Participants were paid \$10.00. Participants were run in pairs; there were two male–male pairs, one female–female pair, and two male–female pairs. None of the participants knew the partners they were paired with. All but one of the participants were right-handed.

Procedure. Participants were brought into a room furnished as a kitchen with storage, appliances, cooking tools, and a table with chairs. We created two scenarios: (1) the care-receiver is seated and reading a magazine, or (2) the care-receiver is standing at a table and organizing items in a large box. We gave each care-receiver a task because we imagined it is more common for a care-receiver to be occupied with a task when he/she asked for a care-giver’s help, instead of being focused on the care-giver. We also posed each care-receiver as either sitting or standing, to see whether this body posture led to a difference in the handover coordination.

In the first scenario, the care-receiver is seated at one end of the room at the table and is reading a magazine; the care-receiver would receive an object name from the experimenter and then ask the care-giver to get it. The care-giver would take the specified item from a stack of items at the end of the room and bring it to the care-receiver at the other end of the room. This procedure was performed for 10 everyday household items: water bottle, tray with a glass of water on it, cup, apple, plate, book, pens, coins, pot, and newspaper. In the second scenario, the same procedure was followed, except that the receiver was given the task of organizing miscellaneous items into a large box while standing at a table.

Measures. Our measures were constructed from the video behavioral data. We analyzed the trials by coding the videos for the physical activities that the participants were engaged in (carrying, reaching out an arm, and transferring). We also noted who reached out first — giver or receiver — in order to extract coordination patterns. The three activities are shown in Figure 1. To analyze variations in these data, we developed the following unitization and coding scheme:

- *Carrying*: Height and number of hands holding an object (one or two); holding fist orientation (toward the receiver or up).
- *Signaling*: Timing and method of signaling (looking, reaching out with hand or object); direction of signaling (for cases where an arm is reached out); number of hands used to signal; holding and grabbing hand orientation of giver and receiver; whether giver adjusted the object for handover (e.g., turn handle toward the receiver).
- *Transfer*: Height of the object at the moment of handover; duration of time the giver and receiver held the object together; where the giver put the object (hand vs. table); distance between the giver and the receiver.

Approaching/carrying. When carrying objects and approaching the receivers, the givers were carrying objects in distinct postures. Sixty-six percent of the time, the givers used both hands when carrying objects, even though the objects used in our experiment were not heavy. This suggests that,

Figure 1. Handover activities observed in a human–human handover study (Lee et al., 2011).

when approaching, actors are not optimizing for comfort but for expressing the intent of not wanting to hand over the object. This is validated in our human–robot study from Section 5.2.

Signaling. All givers and receivers coordinated when and where the handover should happen by communicating their readiness through various communication cues:

Giver signaling readiness. Givers who were carrying an object with two hands, just prior to coming to a stop in front of the receiver, signaled a handover by dropping a hand and reaching out with the object. Givers reached out with the object using one hand. Givers typically started reaching out before they came to a stop near the receiver. However, they did not perform this early reaching behavior if the receiver was not paying attention, which leads us to believe that reaching out was used as a signal.

Receiver signaling readiness. The receiver often signaled receptivity by making grabbing gestures with one or both hands. We saw this behavior in receivers significantly more often when givers were carrying a cup, a pen, or a tray with a glass of water on it. These objects are more likely to be problematic if dropped (as compared with a newspaper or book, for example), so it makes sense that receivers should nonverbally reassure givers they are ready to receive the handover.

Coordination patterns. The most common coordination pattern (58% of trials) was givers communicating a desire to hand over an object by coming close to the receiver. The giver moved the hand holding the object toward the receiver’s hand, and the receiver then would take the object. The second most common coordination pattern (34% of trials) happened when givers reached out the hand with the object at a point where the distance between the two participants was further apart than the sum of their two arm lengths. In these situations, the participants closed the gap somewhat but were further apart when the object was actually transferred. In those cases, the receiver also reached out an arm to grab the object. The giver would then move his or her hand toward the receiver’s hand. Some receivers exhibited very cooperative behavior by leaning their bodies forward while reaching out their arms. The third pattern, although less common (7%), happened when the receiver waited with a grabbing hand gesture but was not looking at the giver. The givers came close to the receivers who did this, and put the objects into the receivers’ hands. The two less common patterns were more frequent when receivers were standing ($\chi^2[2, 158] = 5.7, p = .05$), suggesting that either the receiver’s standing position or the busyness of the receiver (sorting items into a box) led to more signaling and intricate coordination between the givers and the receivers.

Transfer. In the majority of the experiments, the distance between the giver and the receiver did not vary across objects. Also, all the objects were transferred at a height that was below the receiver’s neck, chest level, or even lower. A majority of the object handovers were above the waist. In 24 turns, givers turned a newspaper, book, cup, or pot so that receivers could more easily receive the object. For example, the giver would rotate the cup so the receiver could grab the handle. This

phenomenon occurred in 30% of the handovers with those four objects.

We observed very consistent behavior patterns in givers that replicate the conclusion in Basili et al. (2009), i.e., that handovers have consistency. The results of the study show three main activities that happen during human–human handovers: (1) carrying, (2) coordinating, and (3) object transfer. However, we also observed more coordination with receivers and consideration of their context than what was observed in Basili et al. (2009). When givers were carrying an object, they held it with two hands, exhibiting a very distinct posture when compared to extended arms. As givers approached receivers, givers or receivers indicated whether they were ready by reaching out their arms, and their partners responded by moving their hands toward their partners' hands. When the giver signaled readiness, the giver seemed to take into consideration the receiver's attention and interruptibility (e.g., looking at giver vs. task at hand) and social norms (e.g., being polite by rotating a cup so that the handle faces the receiver).

3.2 Learning the Communication of Intent

The results of our first study suggest that people use reaching actions to communicate their intent to do a handover and help coordinate the timing of the handover. As soon as the givers reached out their arms, the receivers responded by reaching out their arms toward the objects that the givers were holding. The varying coordination patterns between givers and receivers suggest givers intentionally time when to signal. For example, when receivers were looking at magazines, givers did not reach out their arms until they got close to the receivers. On the other hand, when receivers were looking at the givers, givers reached out their arms while still moving toward the receivers. In this section, we analyze communication prior to the givers' reaching actions in order to understand how givers decide when is the right time to signal their readiness to do a handover, i.e., agree on the *where* of the handover.

In our second study (Strabala, Lee, Dragan, Forlizzi, & Srinivasa, 2012), we observed 27 human pairs performing a task that required handovers. The participants were placed in a kitchen environment and tasked with putting away a bag of groceries and packing a picnic basket. Each experiment lasted an average of 250 s during which the participants interacted with the bag of groceries, picnic basket, kitchen cabinets, refrigerator, and each other, performing an average of 9.2 handovers per experiment. We recorded the experiments using three color cameras, four depth cameras, and two micro-phones. From these data, we manually annotated the eye gaze and two-dimensional position of each participant, object locations, and handover actions at a frequency of 10 Hz. These data capture important communication cues, and they are primitive, so a robot can perceive them when we implement our results on a robot in future work. We chose to annotate the data at 10 Hz, a higher rate than that used in our previous study, because we wanted to capture communication cues that cannot be detected with a smaller data frequency, such as glancing at a person to see what they are doing.

We used a machine learning technique called feature selection (Chakraborty, 2007; Guyon & Elisseeff, 2003) to automatically extract sequences of events that are predictive of the physical reaching actions. Then, we used these *sequence features* in a variety of machine learning classifiers (Mitchell, 1997) and found that a decision tree classifier (Figure 2) performs best on the evaluation data set. Finally, this decision tree was validated on a test data set where it accurately predicted 82% of the reaching actions.

The tree is composed of decision points based on the values of the features of the interaction. The decision tree algorithm picks critical values at which the tree splits (e.g., in Figure 2, the first split is at > 0.25 vs. ≤ 0.25). A high value (closer to 1.0) indicates a higher match between the interaction and the feature (details in Strabala et al. (2012)). The figure also shows, for each branch, the percentage of the total data for signals (green) and non-signals (red). The tree misclassified a

Figure 2. Decision tree classifier used to predict reaching actions in human-human handovers (Strabala et al., 2012).

total of 11% of the training data (7% were false positives and 4% were false negatives).

We interpreted the decision tree (Figure 2, right) and determined how to predict the intent to do a handover. When the following four features were all true for one of the participants' hands in our data set, there was an 89% probability that a handover with that hand would directly follow.

No previous signals: Within the previous 3 s the hand did not receive an object and neither participant had performed indirect handovers (placing an object on the table instead of directly handing it over).

Giver orientation and hand occupancy: At the end of the sequence, the giver must be holding an object and must not be facing away from the receiver.

Giver orientation and receiver gaze: At the end of the sequence, the giver must turn to face the receiver, and the receiver must be looking toward the giver.

Giver gaze: At the end of the sequence, the giver is looking either at the object in hand or at the receiver.

The classifier misclassified 11% of the examples. The majority of misclassified examples corresponded to handovers where there was no communication of intent prior to the reaching actions. In these cases, the giver communicated the intent to do a handover by reaching out and expected the receiver to take the object when able.

We interpreted these four features to make the following claims. Participants do not perform both indirect and direct handovers at the same time. Joint attention (i.e., attending to the same thing), and not mutual eye gaze, is a major signal when communicating intent and coordinating reaching

actions. Distance between participants is not a discriminative feature, meaning that reaching actions can be started when the participants are not near each other, up to 3.5 m in our experimental setup. These distances fall within the social distance defined by proxemics.

These results suggest that humans often implicitly signal to each other prior to reaching out, communicating their intent to do a handover, and coordinating the start of reaching actions.

4. The Handover Structure

Our studies on human-human handovers show that people coordinate their behaviors at both the physical and the social/cognitive levels. Physical coordination involves actions that enable the object handover, such as approaching/carrying, reaching, and transferring control. Social-cognitive coordination includes activities that establish an agreement on the *what*, *when/timing*, and *where/location* of the handovers between two individuals. For example, our studies show that people signal their readiness to start a handover to their partner through non-verbal cues such as eye gaze, body orientation, and starting to reach out an arm. In this section, we generalize this notion to three coordination problems at the social/cognitive level: the *what*, *when*, and *where/how* of the handover.

The physical and social/cognitive level coordinations are closely intertwined. For example, the action of reaching out an arm serves both to move the object closer to the receiver and to communicate the intent to start the handover. For descriptive purposes, however, we explain separately the coordination actions that occur at the physical and social/cognitive levels.

In the following section, we describe the physical and social/cognitive level coordination activities involved in the handover process. Informed with theories of common ground and joint activity (Clark, 1996), we added context (common ground) and joint commitment prior to physical and social level coordination.

We describe this handover structure model with four exemplary situations. We chose situations diverse in the means through which and the timings at which the *what*, *when*, and *where* are agreed upon:

Care-giver: a *care-giver* handing over a glass of water at the patient’s request.

Mechanic: a car *mechanic* reaching out while working and asking the assistant for a wrench.

Fire brigade: a *fire brigade* line in which a group of citizens is passing water buckets from a water source to a fire scene.

Flyer: an employee handing out concert *flyers* on a busy university sidewalk.

Our care-giver example follows a typical handover: first, the *what* is established, then the care-giver approaches, the *when* is established before reaching starts, and the *where* is established during reaching — just like in Figure 3, which shows the preferred interaction for a typical fetching task. However, as this section will reveal, not all handovers follow this timeline, but as long as the handover structure coordinates the *what*, *when*, and *where*, handovers will be seamless. Our examples are chosen to emphasize this richness of handovers — note that, although diverse, they do not span the rich space of handover patterns.

This handover structure resonates with the interaction theory in Clark (1996) by representing the handover as a sequence of phases within a joint activity (approach, reach, transfer), requiring joint commitment (the *what* — the handover agreement), and using common ground (the information the actors in a handover believe they share). The actors in the joint activity synchronize with the transitions between these phases, e.g., establishing the *when* nominally transitions from a preparatory phase (approach) to a reaching phase. Once committed, the failure of an actor to perform affects the public perception of every actor’s self-worth and autonomy (Clark, 1996), thus creating a social consequence tied to the success or failure of the handover.

Figure 3. The canonical handover process (physical and social-cognitive channel) for an assistant fetching an object for a requester. The actors first agree that the handover will happen and what object will be transferred. For example, the requester could ask for the object and the assistant could verbally agree. Next, after retrieving the object, the assistant starts approaching the requester while carrying the object. The two actors now exchange communication cues to establish when the handover will happen. For example, as the assistant approaches the requester, they can exchange looks, establishing that they are both ready and the handover can begin as soon as they are close enough. They start reaching at the same time, establishing where the handover will happen based on their motion and their common ground. They then transfer the object and exit the joint activity.

4.1 Context

Context — the state of the world before entering the handover activity — is very important for handovers. Social contexts such as norms or roles (Goffman, 1959) influence how people behave and what they expect from other people. The handover process will be different depending on what context the handover is happening in. Our examples illustrate very different contexts and their impacts on the handover process.

Examples:

Care-giver: The context contains the roles of the patient and care-giver (e.g., the care-giver is supposed to fulfill the patients' requests), and previous handover experiences (e.g., the patient has limited reaching capabilities).

Mechanic: The context contains the roles of the actors, as well as the fact that the mechanic is working underneath a car and cannot see the assistant.

Fire Brigade: The context here contains an established procedure of swinging buckets from one person to another, and the fact that the state of emergency has eradicated many of the usual social guidelines (e.g., personal space).

Flyer: The employee has no prior relationship with the people on the sidewalk, so established social norms shape the behavior that occurs.

4.2 The Physical Channel

The physical channel is strongly tied to the social-cognitive channel: the physical channel implements what the social-cognitive channel decide, e.g., when to do a handover, but also what cues to use for communication.

Carrying/approach: The carrying pose during approach conveys information about the object (e.g., weight, fragility, importance). Depending on context/common ground, this plays a role in

coordinating the *what*, *when*, and *where* of handovers. The social-cognitive channel might also dictate additional communication cues during approach, like eye gaze and body orientation.

Reaching: Flash and Hogan (1985) found that human hand trajectories often follow a minimum-jerk profile. Shibata et al. (1995) analyzed the hand trajectories of givers and receivers during handovers. Huber et al. (2009) observed seated humans handing over objects to one another and came up with a novel trajectory generator based on a decoupled minimum-jerk profile that reproduces the reaching motions of the humans and performs similarly to the minimum-jerk profile. Reaching can be used to establish *what* and plays a role in coordinating both *when* and *where*.

Transfer: In the majority of everyday handovers, both actors are in direct contact with the object and the object, and actors are stationary with respect to one another. In these situations, the actors transfer control by the giver and the receiver exchanging the object load due to external forces such as gravity and wind. After transferring the entire object load, often the receiver will retract or otherwise move the object to signal the giver that the handover is complete. Then the giver will retract his/her arm signaling the same, thus ending this phase and handover interaction. Chan et al. (2012) found a linear relationship between grip force and load force except when either actor is supporting very little of the object load. Analysis of these grip forces suggests that the giver is responsible for the care of the object during the transfer, while the receiver is responsible for the timing of the transfer.

4.3 The Joint Commitment on What — Agreeing to Do a Handover

Before handing over, the giver and receiver must both agree to do a handover, establishing that they are both willing and able to perform the handover. People come to this agreement after one actor proposes the handover and the other actor accepts it. People signal these proposals and acceptances using both actions (verbal and nonverbal) and context, and use the current common ground to decide on what is appropriate. For example, people with little common ground may need to rely on speech to propose and accept the handover, whereas people who handover with each other frequently have more common ground and may use more subtle and efficient signals such as gestures to propose and accept the handover. In our studies, this agreement was either implicit in the task description, or the participants asked for a particular object. Once the agreement to do the handover is established, it enters the common ground.

Examples:

Care-giver: The patient verbally proposes the handover by asking for a glass of water, and the care-giver verbally agrees.

Mechanic: The mechanic reaches out, as well as asks for a wrench. This, together with context (the assistant is around and is tasked to fulfill the mechanic's requests), establishes joint commitment.

Fire brigade: The agreement is assumed based on context alone, i.e., their mutual participation in the task.

Flyer: The employee expresses the desire to interact by facing and approaching a passer-by and reaching out with the flyer. The joint commitment, or agreement that a handover will happen, however, only forms when the passer-by confirms by reaching out or by establishing mutual gaze with the employee. This is an example in which joint commitment is established very late in the handover process, after the giver has finished reaching.

4.4 Coordinating When — Signaling Readiness to Do a Handover

In our previous studies, we found that a way to establish the handover timing is to start reaching out. However, we have found that people also use communication cues (e.g., gaze, body orientation) before reaching that dictate the exact moment of the reaching and establish when the handover will occur. Castiello (2003) showed that eye gaze can be used to infer action intention. Sebanz, Bekkering, and Knoblich (2006) reported that joint attention helps coordinate the initiation and

progression of joint actions. Indeed, in our work in Section 3.2, we found that readiness to do a handover is sometimes established by turning toward and focusing on the other actor or the item to hand over. Furthermore, our work in Section 3.1 indicated that the giver holds the object in a carrying pose when preparing for a handover. The carrying pose conveys information about the object (e.g., weight, fragility, importance). Depending on common ground, the carrying pose can immediately convey to the receiver the desire and readiness to do a handover. The giver may also grasp the object in a way that will facilitate the physical transfer of the object (e.g., allow the giver to present the mug’s handle to the receiver), which also contributes to signaling readiness.

Examples:

Care-giver: The care-giver focuses on the patient. When the care-giver is in view, the patient looks up to the care-giver and to the glass of water. Based on their common ground, this joint attention on the handover-enabled scene signals that both actors are ready to reach, establishing the time of the handover before either party reaches.

Mechanic: The mechanic continuously signals readiness by holding a hand out. However, the timing is only set when the assistant signals readiness by placing the wrench in the mechanic’s hand.

Fire brigade: The rhythm of the task (context) determines the timing of the handover, with no explicit actions required.

Flyer: The timing in this example is established at the same time as the agreement to do a handover, once the passer-by starts reaching back.

4.5 Coordinating Where — Establishing the Configuration of the Handover

The handover configuration is the pose the actors have when they start transferring control of the object (e.g., arms extended and hands grasping the object with some orientation).

In most cases, the giver and receiver negotiate the handover configuration as they reach toward each other. During the reaching, the giver and receiver will pose their hands and their grasp on the object to communicate how they wish to transfer the object (e.g., one actor holds one end of a rod so the other actor can grasp the other end). The reaching communicates some information about the actors and the object (e.g., the object is heavy, one actor cannot reach any further, or one actor is reaching more slowly than the other so the other needs to move closer). For problematic objects (e.g., glass of water), the receiver may reach out more to presumably ensure communication of readiness and speed up the handover.

From our work in Section 3.1, we found that when a care-giver is fetching an object, the physical handover location is at the torso level (between the waist and the neck) and the object is presented to allow for easy grasping.

Examples:

Care-giver: The care-giver and patient reach toward each other based on their previous experiences with each other and meet somewhere in between. The care-giver enforces that the handover occurs with the glass upright.

Mechanic: The mechanic’s hand pose specifies the handover configuration and the assistant complies. From experience working with the mechanic, the assistant will pose the wrench in the mechanic’s hand so the mechanic can immediately use the wrench.

Fire brigade: The giver and receiver have established through routine an agreed upon orientation and location for the handover, including how they each grasp the bucket for the transfer.

Flyer: The passer-by reaches out to where the employee is offering the flyer: the handover configuration is established by the employee. The employee should be considering the passer-by’s preferences when reaching out, so the passer-by can easily receive the flyer.

5. Human–Robot Handovers

We port this knowledge of handover structure to robots and examine some specific aspects of human–robot handovers. In this section, we summarize five studies from our previous work that implement various aspects of human–robot handovers and evaluate their performance. In Sections 5.1 and 5.2, we concentrate on robot-to-human handovers and investigate the questions of *where* (Studies 1, 2, and 3) and *when* (Study 4) the handover should occur in robot-initiated handovers. Next, in Section 5.3, we turn to human-to-robot handovers in which the robot is the receiver as opposed to the giver and study the negotiation of these questions in human-initiated handovers (Study 5).

For the experiments presented in this section, we implemented and evaluated several versions of handover behaviors on our robot HERB (Home Exploring Robot Butler) Srinivasa et al. (2012). HERB was specifically developed for personal assistance in domestic tasks. It has two 7 *df* Barrett WAM arms, each with a 4 *df* Barrett hand with three fingers. The base is a Segway with an additional wheel for safety.

The handover behaviors, implemented for the different experiments presented in this section, ran autonomously during the interactions with participants. However, the sequence of events was scripted for the particular study and an experimenter was in the loop to provide objects for the robot or to put away objects collected by the robot. At the end of this section, we review the existing literature on human–robot handovers, discussing how they relate to our handover structure and how they complement our work on human–robot handovers.

5.1 Choosing Robot–Human Handover Configurations

In robot-to-human handovers, the robot primarily controls where the handover will occur. For the handover to be seamless, the robot must choose the best handover configuration considering the human’s comfort and preference in taking the handed object. This involves choosing the pose (position and orientation) at which the object should be presented and how the robot is configured around the object (its grasp of the object, its base position, and its arm configuration).

5.1.1 Study 1: Handover Configurations for Conveying Handover Intent. As discussed in Section 4, the precursor of the physical handover is the agreement between the two actors that a handover will happen. It is therefore crucial for a robot to communicate its intent when it initiates a handover. This is particularly important in scenarios where the human might not expect the handover, as in our example of handing flyers.

As there is no direct mapping between human and robot poses, we tried to discover robot and object configurations that convey handing intent, by exploring different variables of these configurations. In a study presented in Cakmak, Srinivasa, Lee, Kiesler, and Forlizzi (2011), we created still images of 15 different configurations in which a robot is holding a bottled drink. These images differed in the following dimensions: whether the robot’s arm was extended or not, whether the object was upright or tilted, and whether the robot was holding the object from the front, side, or back. These variables were hypothesized to be potentially important in conveying the intent of doing a hand over.

We conducted an online survey ($N = 50$), asking participants to label each pose as one of six different actions, one of which was handing over. We then analyzed the common features of the four poses that were labeled as handing over, more than any of the other actions (labeled as handing by 92%, 81%, 40%, and 37% of the participants). We observed that a fully extended arm is most effective in communicating the handover intention, presumably because it results in poses that are inefficient for other purposes. All four poses had an extended arm, whereas no other pose had an extended arm. In other words, an extended arm is a distinctive feature of handing configurations.

In addition, we saw that holding the object from the side opposite to the receiver and tilting the object toward the receiver increased the likelihood of a pose being labeled as handing. The pose that was labeled as handing by 92% of the participants had both of these features, in addition to an extended arm. Both of these features result in exposing more object surface to a potential receiver, hence communicating the intent of handing.

5.1.2 Study 2: Handover Configurations From Human Demonstrations. Besides conveying the intent of handing, a robot handover configuration should be preferable from the human receiver’s perspective. To capture these preferences, in a study presented in Cakmak, Srinivasa, Lee, Forlizzi, and Kiesler (2011), we asked participants to configure a simulated version of the robot HERB using sliders for the different parameters of the handover configuration ($N = 10$). Participants labeled five poses as *good* handover configurations and five as *bad* ones, for five different objects. Afterwards, they were shown pairs of images, across which one parameter of the handover configuration was varied, and they were asked to pick the configuration they preferred.

We observed several trends when analyzing the data from this study. The variation in the good handover configurations demonstrated by participants was small, indicating that they shared a *common understanding* of good handovers. Good configurations are *more reachable* than bad ones, allowing at least one configuration where the human model is able to grasp and take the object. For instance, 93% of the poses demonstrated by users as a good way of handing over a plate were reachable for the human, whereas only 53% of the bad examples were reachable.

Good configurations often hold the object in a *default configuration*, meaning the object is oriented as it is seen most in everyday life, which is often also the most stable orientation for the object. For example, we observed that a plate was presented in its upright pose (as if it had food on it) in 85% of the good handover configurations, whereas it was in this default orientation only in 20% of the bad handover configurations.

In the paired comparison questions, participants had significantly higher preference for configurations that have a larger *arm extension* ($p < .001$), supporting findings from Section 5.1.1, and that are *consistent* ($p < .001$), meaning the elbow joint, wrist joint, and the point on the object furthest from the finger contacts are positioned monotonically along all three dimensions. In addition, configurations that look *natural* or *human-like* are preferred by humans ($p < .001$).¹

5.1.3 Study 3: Human Preferences in Handover Configurations. The first two studies provide important guidelines as to how the robot and object should be configured in order to achieve handovers that are preferable for humans. However, ultimately, the robot’s handover configuration needs to be autonomously computed by the robot in novel situations. There are existing planning techniques that can compute possible handover configurations given a kinematic model of the human. However, these methods do not currently take into account the human preference or the interpretability of the robot’s intent. To investigate the gap between handover configurations autonomously computed with these methods, and handover configurations demonstrated by humans, we conducted a study that compares the two approaches in Cakmak, Srinivasa, Lee, Forlizzi, and Kiesler (2011). The *planned* configurations were computed by maximizing the number of available ways to take the presented object, given a kinematic model of the human. The *learned* configurations were obtained from the demonstrated *good* handover configurations, using the median of the demonstrated values for each parameter. The handover configurations obtained with these two methods for five different objects are shown in Figure 4.

We conducted a within-subjects study to compare these two methods ($N = 10$). The study involved our robot HERB handing each object twice, once with each of the configurations obtained

¹ All forced choice paired comparisons were verified with a χ^2 test.

Figure 4. The handover configurations for five objects obtained through two different methods: 1) planned by optimizing ease of taking (top) and 2) learned from human demonstrations (bottom) (Cakmak, Srinivasa, Lee, Forlizzi, & Kiesler, 2011).

through the two methods. The human stood still at a given location, while the robot took the object from an experimenter, turned toward the participant, approached the participant carrying the object, stopped near the participant, and moved to the handing configuration. The robot then said, “Please take the object” and the participant took the object by pulling it. After each pair of handovers, we asked the participants to compare the two handovers in several dimensions.

From this study, we observed that subjects tend to prefer the learned configurations (62%, $p=.09$) and thought they were more natural (65%, $p=.05$) and appropriate (62%, $p=.09$).² On the other hand, from our analysis of handover recordings, we saw that planned configurations provided better reachability. For instance, the participants needed to *step forward*, *bend*, or *fully extend their arm* more times when taking objects that were handed with learned configurations than with planned ones (36 total occurrences for learned handovers, as compared to 27 for planned ones). This demonstrates that existing methods for autonomously generating handovers can provide practical solutions; however, they need to be improved based on the findings from our first two studies (Sections 5.1.1 and 5.1.2) to better fit the humans’ preferences and expectations.

5.2 Communicating Timing of Robot–Human Handovers

As discussed in Section 3, humans communicate and negotiate *when* to do handovers through a subtle exchange of nonverbal signals. Next, we explore how the cues used by humans can be translated to robots in order to make robot-to-human handovers more seamless.

5.2.1 Study 4: Signaling Handover Timing Through High-Contrast Trajectories. In order to signal readiness to hand over an object, humans transition from holding the object close to their torso to presenting the object to the receiver with an extended arm. This distinct transition allows time for the receiver to perceive the intent and predict the trajectory of the arm, resulting in more seamless transitions of the object between the two parties. Similarly, we hypothesize that the handover will be most seamless when the robot transitions to the handing configuration from a contrasting carrying configuration.

²The significance of the preference was measured with χ^2 tests.

Figure 5. On the left: Four conditions for testing the impact of contrast between the carrying and the handing configurations. On the right: Two examples of early handover attempts by a subject, with the two different trajectories that have low contrast (CN and NN) (Cakmak, Srinivasa, Lee, Kiesler, & Forlizzi, 2011).

In a study described in Cakmak, Srinivasa, Lee, Kiesler, and Forlizzi (2011), we compared four different handover trajectories. These trajectories differed in their carrying and handing configurations. The handing configurations involved either a fully extended arm (as suggested by the studies described in Sections 5.1.1 and 5.1.2) or a partly extended arm. The carrying configurations had either a high contrast, or a low contrast with the handing configuration. These four handover trajectories are illustrated in Figure 5.

The handover trajectories were compared through a within-subject study ($N = 24$), in which participants were brought into a kitchen environment and were instructed to sit on a tall chair in front of a table to receive a drink from the robot. The robot approached the participants from their right. Half of the participants were told to pay attention to the robot approaching them while the rest of the participants were instructed to perform an attention test on a computer, for which they needed to use their right hand. The robot’s overall behavior was similar to that described in Section 5.1.3.

The results demonstrated that high-contrast trajectories resulted in more seamless handovers. We saw that the human waiting time was significantly smaller with high-contrast trajectories ($p < .005$).³ Without the contrast, humans attempted to take the object too early, especially when the carrying configuration was similar to the handing configuration with an extended arm (Figure 5). In addition, the data from participants who were performing the attention test suggested that high-contrast trajectories require less of the person’s attention. Participants missed a smaller number of steps in the tests when they were handed with high-contrast trajectories (2.54, $SD = 1.32$) than when they were handed with low-contrast trajectories (3.50, $SD = 1.41$).

5.3 Taking Objects Handed by Humans

Next, we turn to human-to-robot handovers and explore the questions of how to perceive human readiness to hand over an object to the robot, and how to negotiate the position of the handover. Previous work has addressed this issue in cases where the robot suggests the handover location (Pandey et al., 2011). Here we focus on cases where the human suggests the handover location while the robot complies, such as in the mechanic example where the robot is the assistant. This scenario, of robots in an adaptive receiver role, has received little attention in the literature.

5.3.1 Study 5: Adaptive Reaching in Human–Robot Handovers. We developed a system on our robot HERB to accurately perceive reaching actions by humans and to respond by reaching out to

³The effects of contrast, arm extension, and attention were measured in a three-way mixed factor ANOVA. The reported significance value is for the effect of contrast.

Figure 6. User study conducted in Micelli et al. (2011). Top: The execution of a hand over using the planner. Bottom: The execution of a hand over using the reactive controller.

the handover location negotiated with the human. The main sensor used was the Microsoft Kinect, and the robot exploited a behavior engine that allowed it to operate autonomously. The robot was provided with two reaching behaviors that take the object from the human's hand. The first behavior identifies the object location, plans and executes a trajectory to the perceived handover location, and tries to grasp the object. The second behavior is an adaptive reaching behavior that actively takes the object from the human's hand by continuously tracking and moving toward the position of the object. This reactive behavior is able to adapt the robot's motion to changes in the object's position during the interaction.

In order to study how humans perceive a robot that behaves actively close to their body and to evaluate which are the most critical factors during these kinds of interaction, we performed a user study in which five humans were asked to hand over objects to our robot HERB. We also asked the subjects to focus on a computer task during the handover attempts. Figure 6 shows two handovers from the user study. Each participant performed seven handovers with each robotic reaching behavior, totaling 14 handovers per participant. In both cases, we informed the human about how the robot would behave during the handover attempt, thus allowing the humans to predict the robot's motion. The different characteristics of the two behaviors were useful for eliciting different reactions from the human during the interaction. For analysis, we recorded execution times and gathered survey results from the participants. More details about the robot's implementation and the user study setup can be found in our previous work (Micelli, Strabala, & Srinivasa, 2011).

HERB successfully completed 83% of the 70 handover attempts. This result demonstrates our ability to infer the intention of the human to hand over an object and to comply with the handover location negotiated by the human. After the experiment, subjects were asked about what they liked or disliked about the two control algorithms in terms of naturalness, easiness, safeness, and human-likeness of the interaction. Their comments can be summarized using the following four factors that the subjects found crucial during the interaction:

Forcefulness: The force applied by the robot to the object in the final part of the reaching and during the transfer of the object was very important as it provides useful feedback. Indeed, pre-grasp touching signals to the human the robot's *readiness* to do a handover. However, if the force was too strong, the human partner felt uncomfortable and unsafe. When the robot used the behavior based on the planner, sometimes it reached out too far and pushed the subject's hand back. Subjects

found this behavior uncomfortable. HERB used sensor fusion to apply an appropriate force during the adaptive behavior. The robot used both touch and vision to sense proximity to the object. As the adaptive behavior tracks the position of the subject’s hand during the reaching of the robot, HERB made contact with the object without being intrusive. If the robot did not sense contact with the object, instead of pushing against the object, it relied on the vision system. If the observed position of the object was stationary and close to its hand, the robot tried to grasp the object.

Aggressiveness: Subjects felt the robot was being aggressive when it approached the object quickly. While the velocity of the planner was constant during the execution of the trajectory, the velocity of the reactive controller was directly proportional to the distance from the object. The subjects felt more safe and comfortable when the velocity of the robot’s hand decreased when nearing the human hand. In fact, in the majority of human–human handovers, the receiver quickly approaches the giver’s hand and then slows down to accurately grasp the object (Kajikawa & Ishikawa, 2000).

Predictability: The paths the two behaviors executed when reaching out were sometimes quite different. The reactive controller always took a predictable curved path to the goal, whereas the planner could have different unpredictable trajectories. Predictability makes the humans more comfortable around HERB because they can plan into the future and know that HERB will not do anything strange.

Timing: Subjects pointed out that human–human handovers are faster. The reactive controller was faster than the planner by a factor of nearly 2, with a mean time of 8.46 s ($SD = 2.32$) from detection to grasp. Even if the reactive controller was slow compared with human–human handovers, it can be made faster with more aggressive gains. However, tuning the gains of the controller without making the robot too aggressive is challenging, and a complete redesign of the controller behavior would be necessary to obtain more effective human robot handovers. Another way to make the handover faster was to detect the intent to hand over an object before the human arm begins to move as shown in Strabala et al. (2012).

Although these features are always relevant in human–robot handovers, they are even more prominent when the robot behaves actively to exchange an object. This study identifies these features and evaluates which technical solutions have the best potentials for a handover. Further user studies will be conducted to provide a quantitative evaluation of the parameters described above.

6. Discussion

Like any research, this research has many limitations. The human–human handovers and human–robot handovers were observed and evaluated through laboratory experiments. We implemented human–robot handover behaviors on one robot. The studies were conducted in the USA. Human–robot behaviors were implemented to test specific aspects of the handover structure.

In this paper, we attempted to codify a procedure for seamless human–robot handovers. To do so, we derived a handover structure, designed iteratively through laboratory study, and based on observations of humans handing over items to one another. Our studies with people suggest humans handing over items to other humans coordinate the handover process using context and cues that are physically and verbally communicated, and that robots can adopt some of these signals to coordinate handover activities with people. Throughout all of our studies, we found that people could easily understand human-like cues performed by our robot, and that they preferred these cues over machine-like ones.

Following from our framework, we offer the following design recommendations for seamless human–robot handovers. First, HRI designers should rely on human-like gestures and cues for seamless handovers. Second, HRI researchers should model social norms that are well codified and heavily relied on in certain social settings. Third, HRI designers should implement robots with capabilities to detect how people want to establish the *what*, *when*, and *where* of handovers. It is

also important for robots to respond using human-like gestures and signals (so that people know the robots are responding to their signals). However, in addition to human-like signals, special signals can be used when the human and robot share the meaning of these signals in a common ground.

Two exemplary scenarios for seamless human–robot handovers are presented below.

When the robot is a giver: The robot receives a request from a person about what to retrieve or decides on its own that it should give an object to a human. The robot should carry the object in a distinctive carrying posture while approaching the person (Sections 3.1 and 5.2). When the robot nears the person, the robot should observe the person’s eye gaze (i.e., whether the person is looking at the robot) and interruptibility (i.e., not holding objects) (Section 3.2). Upon finding a good moment to interrupt, the robot should reach out with the object toward the torso of the person (Section 3.1). If the person reaches out with an arm toward the robot before the robot reaches out, the robot should reach out in response to hand over the object to the person (Sections 3.1 and 5.3). If the robot cannot find a good moment to interrupt while traveling, it stands near the person and reaches out its arm toward the person (Sections 3.1 and 3.2).

When the robot is a receiver: The robot requests an object from the human, or the human decides to hand over an object to the robot. The robot should approach the person with its arms close to its body (in order to communicate that the robot is not ready to receive) (Section 3.2). When the robot nears the person, it should observe the person to see when he/she reaches out an arm with the object (Sections 3.1 and 3.2). Once the person starts to reach out an arm, the robot should respond by reaching out its arm toward the object (Sections 3.1, 3.2, and 5.3). If the person does not reach out an arm, the robot reaches out with its arm and opens its hand to signal its readiness to receive an object (Sections 3.1 and 3.2).

Extending this work by applying our framework and guidelines to handovers will uncover further design and research questions that are of interest. For example, in human–robot handovers, how can seamlessness be maintained when robots have a primitive arm, are not anthropomorphic, or have very simple sensing and actuation capabilities? An interesting question is how robots should behave based on their limited knowledge of the social and cognitive context of a situation. This work suggests new opportunities for research in human–robot handovers and for exploring the role of social information, cognitive information, and context, in improving interactions between people and the robots that work closely with them.

Acknowledgments

The work presented in this article was carried out as part of the research project Multidisciplinary University Research Initiative (MURI), funded by the Office of Naval Research (ONR MURI N00014-09-1-1031). We gratefully acknowledge funding from the NSF Quality of Life Technologies ERC (NSF-EEC-0540865) and the Intel Embedded Computing ISTAR. Maya Cakmak was partially supported by the Intel Summer Fellowship. Vincenzo Micelli was partially supported by research funds of the University of Parma, Italy.

References

- Agah, A., & Tanie, K. (1997). Human interaction with a service robot: Mobile-manipulator handing over an object to a human. In *Robotics and automation* (pp. 575–580). <http://dx.doi.org/10.1109/ROBOT.1997.620098>.
- Aleotti, J., Micelli, V., & Caselli, S. (2012). Comfortable robot to human object hand-over. In *Robot and human interactive communication* (pp. 771–776). <http://dx.doi.org/10.1109/ROMAN.2012.6343845>.
- Basili, P., Huber, M., Brandt, T., Hirche, S., & Glasauer, S. (2009). Investigating human–human approach and hand-over. In *Human centered robot systems: Cognition, interaction, technology* (Vol. 6, pp. 151–160). Berlin, Heidelberg: Springer. http://dx.doi.org/10.1007/978-3-642-10403-9_16.

- Becchio, C., Sartori, L., & Castiello, U. (2010). Toward you: The social side of actions. *Current Directions in Psychological Science*, 19(3), 183–188, <http://dx.doi.org/10.1177/0963721410370131>.
- Cakmak, M., Srinivasa, S., Lee, M. K., Forlizzi, J., & Kiesler, S. (2011). Human preferences for robot–human hand-over configurations. In *Intelligent robots and system* (pp. 1986–1993). <http://dx.doi.org/10.1109/IROS.2011.6094735>.
- Cakmak, M., Srinivasa, S., Lee, M. K., Kiesler, S., & Forlizzi, J. (2011). Using spatial and temporal contrast for fluent robot–human hand-overs. In *Human–robot interaction* (pp. 489–496). <http://dx.doi.org/10.1145/1957656.1957823>.
- Castiello, U. (2003). Understanding other people’s actions: Intention and attention. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2), 416–430.
- Chakraborty, B. (2007). Feature selection and classification techniques for multivariate time series. *Innovative Computing, Information and Control*, 42, <http://dx.doi.org/10.1109/ICICIC.2007.309>.
- Chan, W. P., Parker, C. A., Loos, H. M. Van der, & Croft, E. A. (2012). Grip forces and load forces in handovers: implications for designing human–robot handover controllers. In *Human–robot interaction* (pp. 9–16). <http://dx.doi.org/10.1145/2157689.2157692>.
- Clark, H. H. (1996). *Using language*. Cambridge University Press.
- Edsinger, A., & Kemp, C. (2007). Human–robot interaction for cooperative manipulation: Handling objects to one another. In *Robot and human interactive communication* (pp. 1167–1172). <http://dx.doi.org/10.1109/ROMAN.2007.4415256>.
- Flash, T., & Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7), 1688–1703.
- Glasauer, S., Huber, M., Basili, P., Knoll, A., & Brandt, T. (2010). Interacting in time and space: Investigating human–human and human–robot joint action. In *Robot and human interactive communication* (pp. 252–257). <http://dx.doi.org/10.1109/ROMAN.2010.5598638>.
- Goffman, E. (1959). *The presentation of self in everyday life*. Garden City, N.Y.: Doubleday.
- Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3, 1157–1182.
- Hoffman, G., & Breazeal, C. (2007). Cost-based anticipatory action selection for human–robot fluency. In *Robotics* (Vol. 23, pp. 952–961). <http://dx.doi.org/10.1109/TRO.2007.907483>.
- Huber, M., Radrich, H., Wendt, C., Rickert, M., Knoll, A., Brandt, T., et al. (2009). Evaluation of a novel biologically inspired trajectory generator in human–robot interaction. In *Robot and human interactive communication* (pp. 639–644). <http://dx.doi.org/10.1109/ROMAN.2009.5326233>.
- Huber, M., Rickert, M., Knoll, A., Brandt, T., & Glasauer, S. (2008). Human–robot interaction in handing-over tasks. In *Robot and human interactive communication* (pp. 107–112). <http://dx.doi.org/10.1109/ROMAN.2008.4600651>.
- Kajikawa, S., & Ishikawa, E. (2000). Trajectory planning for hand-over between human and robot. In *Robot and human interactive communication* (pp. 281–287). <http://dx.doi.org/10.1109/ROMAN.2000.892509>.
- Kajikawa, S., Okino, T., Ohba, K., & Inooka, H. (1995). Motion planning for hand-over between human and robot. In *Intelligent robots and systems* (pp. 193–199). <http://dx.doi.org/10.1109/IROS.1995.525796>.
- Kim, J., Park, J., Hwang, Y. K., & Lee, M. (2004). Advanced Grasp Planning for Handover Operation Between Human and Robot: Three Handover Methods in Esteem Etiquettes Using Dual Arms and Hands of Home-Service Robot. In *Autonomous robots and agents* (pp. 34–39).
- Koay, K., Sisbot, E., Syrdal, D., Walters, M., Dautenhahn, K., & Alami, R. (2007). Exploratory study of a robot approaching a person in the context of handing over an object. In *Multidisciplinary collaboration for socially assistive robotics* (pp. 18–24).
- Lee, M. K., Forlizzi, J., Kiesler, S., Cakmak, M., & Srinivasa, S. (2011). Predictability or adaptivity? Designing robot handoffs modeled from trained dogs and people. In *Human–robot interaction* (pp. 179–180). <http://dx.doi.org/10.1145/1957656.1957720>.
- Lopez-Damian, E., Sidobre, D., DeLaTour, S., & Alami, R. (2006). Grasp planning for interactive object manipulation. In *Robotics and automation*. <http://dx.doi.org/10.1.1.106.9217>.
- Mainprice, J., Gharbi, M., Simeon, T., & Alami, R. (2012). Sharing effort in planning

- human–robot handover tasks. In *Robot and human interactive communication* (pp. 764–770). <http://dx.doi.org/10.1109/ROMAN.2012.6343844>.
- Micelli, V., Strabala, K. W., & Srinivasa, S. (2011). Perception and control challenges for effective human–robot handoffs. In *Robotics: Science and systems workshop on rgb-d cameras*.
- Mitchell, T. (1997). *Machine learning*. McGraw-Hill.
- Mumm, J., & Mutlu, B. (2011). Human–robot proxemics: physical and psychological distancing in human–robot interaction. In *Human–robot interaction* (pp. 331–338). <http://dx.doi.org/10.1145/1957656.1957786>.
- Mutlu, B., Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2009). Nonverbal leakage in robots: communication of intentions through seemingly unintentional behavior. In *Human–robot interaction* (pp. 69–76). <http://dx.doi.org/10.1145/1514095.1514110>.
- Nagata, K., Oosaki, Y., Kakikura, M., & Tsukune, H. (1998). Delivery by hand between human and robot based on fingertip force-torque information. In *Intelligent robots and systems* (Vol. 2, pp. 750–757). <http://dx.doi.org/10.1109/IROS.1998.727283>.
- Pandey, A. K., Ali, M., Warnier, M., & Alami, R. (2011). Towards multi-state visuo-spatial reasoning based proactive human–robot interaction. In *Advanced robotics* (pp. 143–149). <http://dx.doi.org/10.1109/ICAR.2011.6088642>.
- Sadigh, M., & Ahmadi, H. (2009). Safe grasping with multi-link fingers based on force sensing. In *Robotics and biomimetics* (pp. 1796–1802). <http://dx.doi.org/10.1109/ROBIO.2009.5420425>.
- Satake, S., Kanda, T., Glas, D., Imai, M., Ishiguro, H., & Hagita, N. (2009). How to approach humans? Strategies for social robots to initiate interaction. In *Human–robot interaction* (pp. 109–116). <http://dx.doi.org/10.1145/1514095.1514117>.
- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2), 70–76. <http://dx.doi.org/10.1016/j.tics.2005.12.009>.
- Shibata, S., Tanaka, K., & Shimizu, A. (1995). Experimental analysis of handing over. In *Robot and human communication* (pp. 53–58). <http://dx.doi.org/10.1109/ROMAN.1995.531934>.
- Sisbot, E., & Alami, R. (2012). A human-aware manipulation planner. In *Robotics* (pp. 1–13). <http://dx.doi.org/10.1109/TRO.2012.2196303>.
- Sisbot, E., Alami, R., Simeon, T., Dautenhahn, K., Walters, M., & Woods, S. (2005). Navigation in the presence of humans. In *Humanoid robots* (pp. 181–188). <http://dx.doi.org/10.1109/ICHR.2005.1573565>.
- Srinivasa, S. S., Berenson, D., Cakmak, M., Collet, A., Dogar, M. R., Dragan, A. D., et al. (2012). HERB 2.0: Lessons Learned From Developing a Mobile Manipulator for the Home. In *Proceedings of the IEEE* (Vol. 100, pp. 2410–2428). <http://dx.doi.org/10.1109/JPROC.2012.2200561>.
- Strabala, K., Lee, M. K., Dragan, A., Forlizzi, J., & Srinivasa, S. (2012). Learning the communication of intent prior to physical collaboration. In *Robot and human interactive communication* (pp. 968–973). <http://dx.doi.org/10.1109/ROMAN.2012.6343875>.
- Takayama, L., & Pantofaru, C. (2009). Influences on proxemic behaviors in human–robot interaction. In *Intelligent robots and systems* (pp. 5495–5502). <http://dx.doi.org/10.1109/IROS.2009.5354145>.

Authors’ names and contact information: Kyle Strabala, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. Email: strabala@cmu.edu. Min Kyung Lee, Human–Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, USA. Email: mk-lee@cs.cmu.edu. Anca Dragan, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. Email: adragan@cs.cmu.edu. Jodi Forlizzi, Human–Computer Interaction Institute and School of Design, Carnegie Mellon University, Pittsburgh, PA, USA. Email: forlizzi@cs.cmu.edu. Siddhartha S. Srinivasa, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. Email: siddh@cs.cmu.edu. Maya Cakmak, Willow Garage, Menlo Park, CA, USA. Email: mcakmak@willowgarage.com. Vincenzo Micelli, Università Degli Studi di Parma, Parma, Italy. Email: micelli@ce.unipr.it.