

Human Preferences for Robot-Human Hand-over Configurations

Maya Cakmak, Siddhartha S. Srinivasa, Min Kyung Lee, Jodi Forlizzi and Sara Kiesler

Abstract—Handing over objects to humans is an essential capability for assistive robots. While there are infinite ways to hand an object, robots should be able to choose the one that is best for the human. In this paper we focus on choosing the robot and object configuration at which the transfer of the object occurs, *i.e.* the hand-over configuration. We advocate the incorporation of user preferences in choosing hand-over configurations. We present a user study in which we collect data on human preferences and a human-robot interaction experiment in which we compare hand-over configurations learned from human examples against configurations planned using a kinematic model of the human. We find that the learned configurations are preferred in terms of several criteria, however planned configurations provide better reachability. Additionally, we find that humans prefer hand-overs with default orientations of objects and we identify several latent variables about the robot's arm that capture significant human preferences. These findings point towards planners that can generate not only optimal but also preferable hand-over configurations for novel objects.

I. INTRODUCTION

Personal robots that will assist humans in different environments such as homes, offices or hospitals will inevitably face tasks that require handing over objects to humans. Robots can fetch desired objects for the elderly living in their homes or hand tools to a worker in a factory. Different aspects of this particular kind of physical human-robot interaction have received a lot of attention in robotics. While there has been substantial progress with approaches that use a kinematic model of the human or take inspiration from human-human interactions, we believe that it is valuable to evaluate how humans would prefer being handed an object by a robot. In this paper we address the problem of collecting information about such preferences and incorporating them in the design of the robot's interactions.

For humans, handing objects or taking objects handed by others is often routine rather than deliberative. Humans carry out successful hand-overs on a daily basis with a variety of objects such as credit-cards, coins or plastic bags. Yet we cannot easily remember these instances or identify how exactly we hand-over particular objects. Our long-term goal is to reach this level of seamless and effortless hand-overs between humans and robots.

This work partially supported by Intel Labs Pittsburgh and the National Science Foundation under Grant No. EEC-0540865. M. Cakmak is with the School of Interactive Computing, Georgia Institute of Technology maya@cc.gatech.edu. S. Srinivasa is with Intel Labs Pittsburgh siddhartha.srinivasa@intel.com. M. K. Lee, J. Forlizzi and S. Kiesler are with the Human-Computer Interaction Institute, Carnegie Mellon University [{mklee, forlizzi, kiesler}@cs.cmu.edu">{mklee, forlizzi, kiesler}@cs.cmu.edu](mailto).




Fig. 1. HERB handing over a drink bottle with a configuration planned using a kinematic model of a human (left) and learned from examples given by other humans (right).

The instrumental goal of handing over is to transfer an object from the robot to the human. This goal on its own poses a highly under-constrained problem. There are infinite ways to achieve transfer of an object between two individuals. While this nature of the problem can be exploited with ad-hoc solutions that work *good-enough*, it also presents the challenge of finding the *best* option among available solutions. We believe that users at the receiving end of the interaction are the ultimate evaluators of the hand-over.

Handing over involves several phases starting from picking up the object in a particular way, to retracting the arm after releasing the object. In this sequence, the moment of transfer is of crucial importance. The way that the robot configures the object, as well as its own body at this moment, determines how the person will take the object from the robot. This paper focuses on choosing these *hand-over configurations*. Note that the hand-over configuration is one of the many factors that influence hand-overs. A complete hand-over behavior will need to consider other factors such as the robot's trajectory or the person's posture and gaze direction. Our study provides insights into one of the crucial factors by isolating the others and is complementary to many of the studies in the literature that inform the design of complete hand-over behaviors.

Using a kinematic model of a human in studying robot-human hand-overs has been a common approach in the literature. Different aspects of hand-over interactions have been studied with this approach, including motion control [1], [8], [13], grasp planning [12], [10] and grip forces to be applied during hand-over [15], [9]. The problem of choosing hand-over configurations was addressed in [17] proposing the criteria of *safety*, *visibility* and *comfort*.

The design of hand-over interactions needs to go beyond efficient kinematic optimization in order to achieve *usability*, *naturalness* and *appropriateness*. Research on hand-overs between two humans has partially addressed this concern. Different aspects of hand-overs between humans have been

studied in the literature. This includes grip forced applied by humans during hand-overs [14], trajectories and velocity profiles adopted by humans both in the role of giver and receiver [16], and the social modification of the instrumental movement of pick-and-place in the context of hand-overs [3]. While these studies have interesting implications on human-robot hand-overs, there is not much evidence suggesting that approaches that will work best for human-robot hand-overs are the ones adopted during human-human hand-overs.

User studies involving actual human-robot hand-overs are particularly valuable in guiding the design of hand-over interactions. [11] analyzes human preferences about the robot's hand-over behaviors in terms of the approach direction as well as height and distance of the object. User preferences between two velocity profiles for handing over is analyzed in [8] in terms of several behavioral measures as well as participant's rating of human-likeness and feeling of safety. [7] presents a study which demonstrates the effectiveness of a simple hand-off mechanism that automatically drops the object without any sensing. In [5] user preferences between direct delivery versus delivery by setting on a plain surface was analyzed. Robot poses that convey the intent of handing over are determined with an online survey in [4].

In this paper we present a user study to collect data about human preferences for hand-over configurations and we perform a human-robot interaction (HRI) experiment to evaluate how hand-over configurations learned from such data compare to configurations planned using a kinematic model of the human. We find that the learned configurations are preferred and found more natural and appropriate, while the planned configurations provide better reachability of the object. We also analyze the data from our user study about human preferences and identify several latent variables along which participants show significant preference.

II. APPROACH

A. Hand-over configurations for robots

This paper focuses on the problem of choosing the configuration of the robot and the object at which the hand-over occurs, which we refer to as a *hand-over* configuration.

A hand-over configuration can be fully specified by three variables $C_{handover}^r = (P_{grasp}^r, C_{arm}^r, P_{base}^r)$ where P_{grasp}^r denotes the grasp pose of the robot's hand relative to the object¹, C_{arm}^r denotes the robot's arm configuration and P_{base}^r denotes the robot's position relative to receiver.² These variables will have different degrees of freedom depending on the robot platform. Note that when $C_{handover}^r$ is fixed, the 6D configuration of the object, C_{obj} , is also fixed.

Any collision-free point in the space defined by these variables is a viable hand-over configuration. However the effort required by the receiver can vary a great deal. For instance the robot facing away from the receiver will require going around the robot to take the object. Without a value

function defined over the space of hand-over configurations this is as viable as any other configuration.

This paper advocates incorporating human preferences while specifying such value functions. We present a simple approach for achieving this, and compare it against a planning approach similar to [19] and [12]. Both approaches are explained in more detail in the rest of this section.

B. Planning hand-over configurations

Using a kinematic model of the human, a robot can simulate how an object will be taken by the receiver in each of its hand-over configurations. It can then choose the configuration that provides the easiest or most comfortable taking configuration for the receiver. This approach goes beyond pre-defined, ad-hoc hand-over configurations by providing a way for the robot to plan object specific hand-over configurations.

While some hand-over configurations may provide more than one taking configuration for the human, others might not allow any. We represent a taking configuration with $C_{take}^h = (P_{grasp}^h, C_{arm}^h)$ asserting that the human should not need to move in order to take the object. Note that C_{take}^h is dependent on $C_{handover}^r$ in that C_{take}^h is limited to a subset of all possible take configurations given $C_{handover}^r$. The set of all possible take configurations given a hand-over configuration is denoted by $\mathbb{S}_{take}^h | C_{handover}^r$.

The value of the overall hand-over can then be considered to have two components for giving and taking. We simplify the dependence of take configurations on hand-over configuration to a dependence on the 6D configuration of the object, C_{obj} . The value function can then be expressed as:

$$f = f_{handover}(C_{handover}^r) + f_{take}(\mathbb{S}_{take}^h | C_{obj}) \quad (1)$$

Using this value function the robot can evaluate each hand-over configuration and choose the best one. Thus the problem of choosing a hand-over configuration is converted to finding the optima of this function. For this we use *hierarchical optimization*, which constrains the search on variables by defining a series of optimization problems that are solved in a predetermined order [2].

The hand-over configuration is chosen in a hierarchical manner as follows. First we find the object configuration, C_{obj} , that maximizes f_{take} . Note that, f_{take} should account for all possible take configurations provided by an object configuration. Using a value function that evaluates individual take configurations, this can be achieved by obtaining an average or a maximum over \mathbb{S}_{take}^h . Instead we use the size of this set, without accounting for how good each configuration is. Thus we pick the object configuration that provides the maximum number of take configurations to the receiver.

Next, the robot needs to choose a hand-over configuration, $C_{handover}^r$, given the chosen object configuration, as to maximize $f_{handover}$. This is further divided into two steps. First the robot chooses a P_{grasp}^r and P_{base}^r that provides the highest number of C_{arm}^r , then it chooses the C_{arm}^r that is furthest away from the joint limits. Note that a number of evaluation functions, other than distance to joint limits can

¹In this study we pre-compute a database of grasps for different objects and use them for handing over.

²Superscript r denotes association with *robot* and superscript h denotes association with *human*.

be used; such as effort required to carry the object in the given configuration. Similarly a different order or grouping of optimization steps is possible. In this paper we are not concerned with the optimality of the solutions, but we are interested in finding a solution that reflects consideration for human kinematics.

C. Learning hand-over configurations from users

The way in which human preferences can be incorporated in the robot's behavior is dependent on how such preferences are represented. This representation has strong implications on how the information about preferences will be obtained from the human. In this paper we explore two ways of getting human input.

Good and bad examples. The user is given control of the variables of the hand-over configuration and asked to configure the robot in what they think is a good or a bad configuration.

Systematic comparisons. The user is asked to pick one of two configurations that differ in one or more variables.

While both types of input can be used in estimating a value function, in this paper we restrict our comparative evaluation to a value function obtained from good and bad examples. In addition we analyze both types of input in detail.

Instead of letting the users configure each joint of the robot's arm, we choose a more intuitive variable set for the hand-over configuration: $C_{handover}^r = (C_{obj}, P_{grasp}^r, P_{base}^r)$. Note that these variables do not fully specify the configuration of the robot, since more than one arm configuration might be possible. While we let users choose among the possible arm configurations provided by their choice of other variables, by default the arm configuration is chosen in the same way as in the planning approach.

In order to avoid making assumptions about the underlying distributions, we use an instance based model to estimate the value function from given examples. Each configuration is evaluated based on how similar it is to good examples and how different it is from bad examples given by users. The value function is written as:

$$f = \frac{\frac{1}{|\mathbb{S}_{good}|} \sum_{C_j \in \mathbb{S}_{good}} d(C_{handover}^r, C_j)}{\frac{1}{|\mathbb{S}_{bad}|} \sum_{C_i \in \mathbb{S}_{bad}} d(C_{handover}^r, C_i)} \quad (2)$$

Here \mathbb{S}_{bad} and \mathbb{S}_{good} are the set of collected good and bad examples, and $d(C_1, C_2)$ is a similarity function defined between two configurations. It takes a maximum value of 1.0 when the two configurations are exactly the same and goes to zero as the configurations become dissimilar. For the similarity between object configurations and base poses we use the 1.0 minus the Euclidian distance normalized to the maximum distance between any two examples. Examples further than the maximum distance are not considered. The similarity between two grasps is 1.0 when the grasps are the same and zero otherwise. The overall distance is the average of the distances of three variables. Among available hand-over configurations the robot picks the one that maximizes this functions.



Fig. 2. Objects used in experiments and their simulated models: (1) Plate, (2) Notebook, (3) Bottle, (4) Shaker, (5) Mug.

III. EXPERIMENTS

A. Platform

Our research platform is HERB (Home Exploring Robot Platform) (Fig.1) developed for personal assistance tasks in home environments [18]. HERB has two 7-DoF WAM arms, each with a 4-DoF Barrett hand with three fingers. The WAM arms provide position and torque sensing on all joints. Additionally their stiffness can be set to an arbitrary value between 0 (corresponding to gravity compensation mode) and 1 (corresponding to maximally stiff). The sensing for objects being pulled from HERB's gripper is based on end effector displacements detected while the arm has low stiffness. OpenRAVE [6] is used for simulating kinematics of the robot and the human and for grasp planning.

A human model with 8-DoF arms and 17-DoF hands is used for the planning approach described in Sec. II-B. The model is 162cm tall. The joint limits of the human model are adjusted such that any configuration of the joints is a physically possible one. Both approaches are evaluated with 5 different objects shown in Fig.2.

Two additional simplifications are made for the experiments in this paper. First, the position of the robot relative to the human is restricted to a single line facing the human. Thus choosing P_{base}^r is reduced to choosing the distance of the robot. This choice is in line with the conclusion in [11] that humans prefer being approached from the front sector of their personal space in hand-over interactions. Secondly, the space of variables is discretized and limited within a feasible region in front of the humans right hand. This is done mainly to provide real-time interactivity in the graphical user interface used for getting user input by pre-computing all inverse kinematic solutions prior to the study.

B. Collecting information on user preferences

We conducted a two-part study to get input from users on how the robot should hand-over different objects. In the first part participants are asked to give *good* and *bad* examples of hand-over configurations through a graphical user interface (Fig.III-B). The interface provides sliders to change each degree of freedom of the hand-over configuration variables described in Sec. II-C. This gives 8 sliders: 6 for the position and rotation of the object (C_{obj}), one for the distance of the robot from the human (P_{base}^r), and one for the grasp type (P_{grasp}^r). An additional dynamic slider is provided to let the user choose alternative arm configurations, if any, different




Fig. 3. User interface for collecting good and bad examples of hand-over configurations.

from the default one. When there are no arm configurations that support the combination of variables to which the sliders are set, a sign that says “N/A” appears above the robot. Position variables are discretized with 10cm resolution and orientations with 45 degree resolution. A total number of 10-20 grasps are available for each object.

Users can view the configurations from multiple angles by navigating the scene in 3D using the mouse. Two buttons let the user submit a configuration as *good* or *bad*. Participants are asked to give 4 good and 4 bad examples for five different objects. The order of objects and the initial configuration of sliders for each object are randomized for each participant.

In the second part of the study, participants are presented with a set of image pairs of hand-over configurations. Each pair is presented side-by-side and users pick the configuration they prefer by clicking on the corresponding image.

The images of configurations are obtained with the same program used in the first part of the user study using an isometric perspective. Each pair is obtained by varying one or more of the variables that participants could manipulate. For each of the 5 objects, 9 pairs are obtained by varying variables between the two extremes of the sliders provided to users while keeping all other variables constant at a default value. If the extreme configuration is not possible, the configuration that is closest to the extreme is used in the comparison. For the grasp type variable we choose two grasps that have contact points on the object as far from each other as possible. For the arm configuration variable we choose two configurations in which the elbow joint is as far from each other as possible. An additional 16 pairs are obtained by varying multiple variables simultaneously, or by varying a variable between one extreme and one default value. This results in a total of 61 comparisons presented to each participant.

Before starting, participants are briefed about the goal of the study. In order to give a sense of HERB’s real size, they are given an introduction about its capabilities while physically standing in front of HERB. They are also shown the actual five objects to get a sense of their weights. Finally, participants are given a short demonstration of the user interfaces for both parts of the study. The experimenter explains each variable by moving the sliders to demonstrate

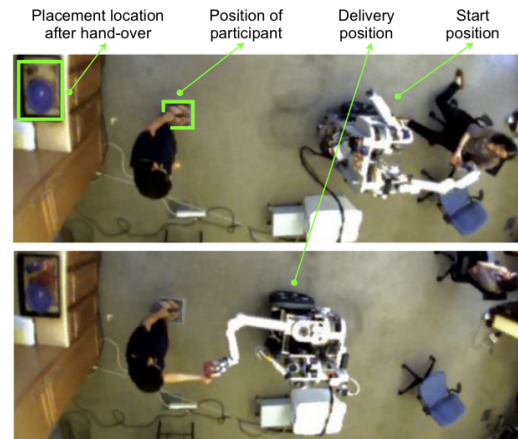


Fig. 4. Setup for the human-robot interaction study for comparing two approaches.

the effect. Participants are told that the receiving human is right-handed.

C. Comparing two approaches

A human-robot interaction experiment was conducted to evaluate hand-over configurations obtained with the two approaches described in Sec. II. For each of the five objects a single hand-over configuration is obtained with both approaches (Fig.5). These configurations are used for delivering the objects to the participant.

In this experiment participants are asked to stand on a square marked on the ground (Fig.4). The robot starts at about 1m away from the participant facing away from the participant. The experimenter hands the object to the robot in the respective grasp type and the robot configures its arm according to the respective hand-over configuration. It turns 180 degrees to face the participants and moves toward them until it is at a distance specified by the configuration. Then it says “Please take the object.” and waits to sense a pull on the arm. When a pull is detected it opens the gripper and moves the end-effector 10cm toward itself. Next it gets back to the starting point. Participants are told to place the object on a tray behind them after they take it.

The robot delivers two of each object one after the other using the configurations generated with the two approaches. After the two deliveries, participants are asked to compare the two hand-over configurations by answering four questions:

- 1) Liking: *Which one did you prefer?*
- 2) Naturalness: *Which one looked more natural?*
- 3) Practicality: *Which one was easier to take?*
- 4) Appropriateness: *Which one was more appropriate?*

The survey also includes comment boxes for specific remarks that participants might want to express. The order of five objects and order of the two hand-overs for each object are randomized for each participant.

Prior to the interactions, participants are briefed about the goal of the study and the interaction sequence is explained to them pointing to all objects and the relevant locations.




Fig. 6. Position and orientation of all objects averaged over good examples (Top) and averaged over bad examples (Middle) collected in the user study. (Bottom) Positions and orientations that are reachable to the human model mapped onto the discrete values used in examples given by participants. Transparency decreases with higher occurrence.

They are told to pay attention to the position and orientation in which the object is presented to them as well as the arm configuration of the robot. The four questions based on which they will compare the hand-over configuration pairs are read to them by the experimenter. They are also told that other aspects of the hand-over, such as the force with which they need to pull, or the speed with which the robot approaches will be the same in all interactions. This is done to avoid small variations in these aspects of the hand-over effect the participant's preference. Finally they are asked to take an object (different from the ones used in the experiment) from the robot so that they get a sense of how much they need to pull the object. Participants are told to use their right hand to take all objects delivered by the robot.

IV. RESULTS

A. Analysis of user input

We present some observations about the hand-over configuration examples given by 10 participants (8 male, 2 female between the ages of 19-36) and their preferences in the comparison of systematically chosen configuration pairs.

1) *Good and Bad Examples:* Analyzing the good and bad examples configured by participants we make the following observations.

a) *A common understanding of good:* Fig.6 shows the the positions and rotations of the objects across all good and bad examples given by the participants. We observe that the good examples given across participants are concentrated around few values of each variable. The distribution of examples is unimodal and has a small variance. This indicates that preferences for the object configuration in the hand-over is similar across different people.

b) *A sense of reachability:* The positions of the object in the good examples given by participants are often well chosen in terms of their reachability for the human model.

TABLE I

REACHABILITY AND PREFERENCE FOR DEFAULT ORIENTATIONS OF OBJECTS IN THE EXAMPLES GIVEN BY PARTICIPANTS.

Object	Reachable examples		Default examples	
	Good	Bad	Good	Bad
Bottle	83%	63%	65%	40%
Mug	63%	30%	36%	10%
Notebook	53%	18%	50%	40%
Plate	93%	50%	85%	20%
Shaker	40%	28%	68%	28%

We refer to object configurations that provide at least one take configuration for the human as his *reachable space*. A high overlap in the positions can be observed for some objects in Fig.6, comparing the distribution of good examples given by participants and the reachable space of the human model mapped onto available discrete space.

Orientations, on the other hand, do not overlap as much. Table I reveals the percentage of good and bad examples that are reachable. We find that for some objects, such as the shaker and the notebook reachability of good examples is rather low. This points towards the necessity of a planning approach which makes sure an object is reachable to the human. In addition, bad examples given by participants are much less reachable than good examples. This shows that it is important for the robot to present an object in a reachable configuration for it to be considered a good hand-over.

c) *Preference for default orientations:* We refer to the orientation in which an object is viewed most frequently in everyday environments as its *default orientation*. These orientations are often the most stable orientation for the object. Rotations around the vertical axis often do not effect the stability of the object, however for non-symmetric objects it can result in different functional properties. In this study the default configurations are chosen as any upright orientation for the bottle, plate and shaker; upright positions with the handle on the right side for the mug; and the lying in a readable orientation for the notebook.

Table I gives the percentage of configurations in which the object is in its default orientation in good and bad examples. We observe that this number is rather high considering that there are many alternatives. The number is lower for objects that are not symmetric around the vertical axis (mug and notebook) since we limited the default orientation to a particular rotation around the vertical axis. We also see that the percentage of default-orientation examples is much higher in good examples than bad examples. This again shows that participants generated bad examples by altering this property which they thought was important.

This finding is important since it can be applied to novel objects if their default orientation is known. Rotating an object around different axes can have different effects on different objects and may cause the object to appear in an unfamiliar orientation. Thus, it is a safe approach for the robot to hand-over objects in their default orientation.




Fig. 5. Handing configurations chosen for comparative evaluation on the physical robot. (Top) Planned using a kinematic model of a human. (Bottom) Learned from examples given by the users.

TABLE II

PREFERENCE ON DIFFERENT MODIFIABLE VARIABLES (TOP) AND LATENT VARIABLES (BOTTOM) IN COMPARISONS OF CONFIGURATIONS PAIRS. REFER TO TEXT FOR A DESCRIPTION OF THE VARIABLES.

Preferred value of variable	# of pairs	% preferred	Significance
<i>High (over low)</i>	5	86%	$\chi^2(1, N=50)=25.92, p<.001$
<i>Left (over right)</i>	5	90%	$\chi^2(1, N=50)=32.00, p<.001$
<i>Near (over far)</i>	5	74%	$\chi^2(1, N=50)=11.52, p<.001$
<i>Default orient.</i>	19	84%	$\chi^2(1, N=190)=730.00, p<.001$
<i>Extended arm</i>	4	92%	$\chi^2(1, N=40)=25.12, p<.001$
<i>Arm consistent-x</i>	4	100%	$\chi^2(1, N=40)=34.00, p<.001$
<i>Arm consistent-y</i>	6	80%	$\chi^2(1, N=60)=27.92, p<.001$
<i>Arm consistent-z</i>	12	72%	$\chi^2(1, N=120)=152.08, p<.001$
<i>Natural</i>	17	76%	$\chi^2(1, N=170)=442.88, p<.001$

2) Systematic comparisons:

a) *Preference on directly modifiable variables:* Preferences observed in the good and bad examples given by participants are supported by their responses in the comparisons. On Table II we present participants' choices on different variables for which a significant preference toward one value was observed.

In terms of the position of the object relative to the human, we find clear preferences for *high over low*, *left over right* and *near over far*. While this is consistent with the reachability of the objects, it may also indicate a concern for the *visibility* of the object to the human in that an object is more visible when it is closer, higher and centered in the field of view.

We find no consistent preference for individual rotation variables, however considering comparisons in which the configurations differ in whether the object is in its default orientation or not, we find a strong preference for the default orientation, supporting our earlier observation.

We can not directly talk about a preference on the grasp type variable as it is not shared across different objects. The grasp type variable is an index on the grasp database which has no meaningful order. It is also difficult to comment about




Fig. 7. Two configurations that differ only by the grasp type variable. The configuration in (a) is consistent in the x-axis whereas the one in (b) is not. Both configurations are consistent in y and z axes.

grasp preferences on a particular object since only one or two comparisons are made for differing grasp types. The same problem applies to arm configurations variable that is limited to a small subset. In the following we define some latent variables that capture the effect of these parameters better.

b) *Preference on latent variables:* We believe that the arm configuration has an important role in making a hand-over configuration look natural and communicate the intent of handing the object. In our user study the participants did not have direct control on the arm configuration since it was limited by the choice of other variables. However combinations of different variables can result in variations in arm configurations in terms of different properties. We refer to these variations as *latent variables*. We define three latent variables as follows.

Arm extension is a measure of the distance of the end-effector to the robot's base. It is defined as the ratio of the current distance to the maximum distance. Arm extension is 100% when the arm is fully extended.

Consistency is defined as the monotonic arrangement of the elbow joint, the wrist joint and the point on the object that is furthest away from the finger contacts. Consistency can be considered along the three axes. For instance, in a configuration that is consistent along the vertical axis, either the wrist is lower than the elbow and the end point of the object is lower than the wrist, or this order is reversed. For instance the arm configuration given in Fig.7(a) is consistent

in the x-axis while the one in Fig.7(b) is inconsistent. Both configurations are consistent in the y and z axes.

Naturalness is loosely defined as mappability to a human configuration given the degrees of freedom and joint limits of the human. We assume a correspondence between the shoulder and elbow joints of the robot and the human. In addition we assume a correspondence between the single finger on the gripper and the human thumb. We look at whether the arrangement of the positions of shoulder, elbow and wrist joints (given the finger mappings) is achievable with a human model. By arrangement we refer to their orders in some direction on all three axes.

On the 61 configuration pairs compared by the participants we identify the pairs that differ in terms of these latent variables. For arm extension we choose pairs in which one of the configurations is extended at least 20% more than the other. The preferences on these three latent variables are given in Table II. We find a strong preference on all three variables. This is not surprising as our goal in defining these latent variables was to capture such preferences. This provides a good set of properties to constrain the choice of arm configurations in a hand-over configuration.

B. Human-robot interaction experiment

Our HRI experiment for comparing the planning and learning approaches described in Sec. II was completed by 10 right-handed participants (6 male, 4 female between the ages of 20-32). The results from the survey comparing the two approaches are summarized in Table III and Table IV. We find that the hand-over configurations learned from user examples is preferred more than the configurations produced with planning in all dimensions. The difference in preferences is most significant for *naturalness*, which shows that humans' notion of a *good* hand-over configuration includes *naturalness* and the planning approach does not spontaneously produce natural looking configurations. While we did not find a preference for configurations produced with planning in terms of *practicality*, this was the dimension that was least in favor of learned configurations.

There are some differences on individual objects. For instance we observe that the planned configuration for the notebook was preferred by more participants and was thought to be more practical and appropriate. One of the participants who preferred the planned configuration mentioned that it was "better suited for how [she] naturally orients [her] hand while reaching out to grasp the notebook". Another subject commented that "the robot made the hand-over more convenient, [and that she] did not have to stretch [her] hand out as much". Preference on the shaker configurations was almost equally distributed. Referring to these configurations some participants noted that they were "almost the same" or that they "could not see a difference".

Other comments by the participants supported their preferences. Referring to naturalness of planned configurations subjects mentioned that the arm "was awkward looking" or "had awkward direction of joints". Referring to appropriateness of the planned configuration for the plate, one

TABLE III
OVERALL COMPARISON OF TWO APPROACHES ON SURVEY QUESTIONS.

Criteria	Preference		Significance
	Planned	Learned	
Liking	38%	62%	$\chi^2(1,N=50)=2.88, p=.09$
Naturalness	36%	64%	$\chi^2(1,N=50)=3.92, p=.05$
Practicality	46%	54%	$\chi^2(1,N=50)=0.32, p=.57$
Appropriateness	38%	62%	$\chi^2(1,N=50)=2.88, p=.09$

TABLE IV
COMPARISON OF TWO APPROACHES ON SURVEY QUESTIONS FOR INDIVIDUAL OBJECTS. NUMBER OF PARTICIPANTS OUT OF 10 WHO PREFERRED LEARNED CONFIGURATIONS ARE GIVEN.

Criteria	Bottle	Mug	N.book	Plate	Shaker
Liking	7	8	4	6	6
Naturalness	6	8	6	7	5
Practicality	5	7	4	5	6
Appropriateness	8	7	4	6	6

TABLE V
COMPARISON OF TWO APPROACHES BASED ON ANALYSIS OF VIDEOS.

		Planned	Learned
Number of event occurrences for each object	Bottle	2	6
	Mug	5	1
	Notebook	2	7
	Plate	6	3
	Shaker	13	19
Number of occurrences of each event across objects	Bending	15	17
	Stepping forward	2	6
	Extending arm	11	13

subject mentioned that the "the slanted orientation of the plate did not seem appropriate, [because] you expect to have the plane horizontally". About the appropriateness of the learned configuration for the mug, 6 subjects referred to the handle being at the right place.

Recordings of the interactions are coded for the occurrence of one or more of three particular events that are believed to be indicative of problems in terms of *practicality* of the hand-over configuration. These are *bending* forward, *stepping forward* and taking the object with a *fully extended arm*. A snapshot from each event is shown in Fig.8(Top). The counts of each event for individual object for both approaches are given in Table V.

We observe that overall the planned configurations have less occurrences of these problems. However this can differ for particular objects. The mug is an interesting example. In the planned configuration It is presented in an orientation facing down, with the handle being towards the human. The choice of this configuration in the planning approach is based on the human grasping the object with a comfortable grasp, *i.e.* such that they would carry it in this orientation. However, some people grasp the mug such as to rotate it its default orientation. Three examples of participants taking the object in this way are shown in Fig.8(Bottom).

Note that there is a large difference between the height



Fig. 8. Handing configurations chosen for comparative evaluation on the physical robot. (Top) Learned hand-over configuration for the bottle. (Bottom) Planned hand-over configuration for the mug.

of the human model and the height of participants, which is 175.50cm ($SD=11.46\text{cm}$). Therefore, while the differences in the occurrence of events are informative, the absolute number of the counts can be misleading. For instance one participant of height 192cm had to bend for all of the hand-overs. This points towards importance of customization.

Overall, our evaluation shows that learned configurations were preferred by users in all aspects, including practicality. However analysis of the videos showed better reachability for objects presented with planned configurations. Individual objects presented exceptions to these results.

V. CONCLUSION

We present two user studies that address the question of human preferences for hand-over configurations. In the first study, we collect data about human preferences for hand-over configurations. We analyze this data to identify significant preferences and we define several latent variables that reflect these preferences.

In the second study we evaluate hand-over configurations that incorporate human preferences in a human-robot interaction experiment by comparing them to configurations that are planned using a kinematic model of a human. We find that the learned configurations are preferred in terms of different criteria, however the planned configurations provide better reachability of the object. While a planning approach has the potential to produce configurations that are *practical*, it is insufficient in addressing *usability*, *naturalness* and *appropriateness*. Configurations that were learned from examples given by users implicitly encoded these properties, and therefore were preferred over planned configurations. In this paper we tried to explicitly define some latent variables that might capture these properties.

Based on both user studies we outline some modifications to a planner that might produce more usable, natural and appropriate hand-over configurations. *Visibility* of the object should be taken into account in the position of the object while satisfying reachability. The robot should try to present objects in the *default orientation* when possible. *Affordances*

of the object (such as handles) should also be taken into account in the configuration of the object. The arm of the objects should preferably be as *extended* as possible, while complying with as many *latent variables* for *consistency* and *naturalness* as possible. Hand-over planners with these modifications might be able to generate not only optimal but also preferable hand-over configurations for novel objects. Implementation and evaluation of such planners will be explored further in future work.

REFERENCES

- [1] A. Agah and K. Tanie. Human interaction with a service robot: Mobile-manipulator handing over an object to a human. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, 575–580, 1997.
- [2] G. Anandalingam and T. Friesz. Hierarchical optimization: An introduction. *Annals of Operations Research*, 34(1):1–11, 1992.
- [3] C. Becchio, L. Sartori, and U. Castiello. Toward you: The social side of actions. *Current Directions in Psychological Science*, 19(3):183–188, 2010.
- [4] M. Cakmak, S.S. Srinivasa, M.K. Lee, S. Kiesler, and J. Forlizzi. Using Spatial and Temporal Contrast for Fluent Robot-Human Hand-overs In *Proc. of Human-Robot Interaction*, 2011.
- [5] Y.S. Choi, T.L. Chen, A. Jain, C. Anderson, J. Glass, and C.C. Kemp. Hand it over or set it down: A user study of object delivery with an assistive mobile manipulator. In *Proc. of IEEE RO-MAN*, 2009.
- [6] R. Diankov and J. Kuffner. OpenRAVE: A Planning Architecture for Autonomous Robotics. Technical Report CMU-RI-TR-08-34, Robotics Institute, July 2008.
- [7] A. Edsinger and C. Kemp. Human-robot interaction for cooperative manipulation: Handing objects to one another. In *Proc. of IEEE RO-MAN*, 2007.
- [8] M. Huber, M. Rickert, A. Knoll, T. Brandt, and S. Glasauer. Human-robot interaction in handing-over tasks. In *Proc. of IEEE RO-MAN*, 107–112, 2008.
- [9] I. Kim and H. Inooka. Hand-over of an object between human and robot. In *Proc. of IEEE RO-MAN*, 1992.
- [10] J. Kim, J. Park, Y. Hwang, and M. Lee. 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 *Proceedings of the Intl. Conf. on Autonomous Robots and Agents*, 34–39, 2004.
- [11] K. Koay, E. Sisbot, D. Syrdal, M. Walters, K. Dautenhahn, and R. Alami. Exploratory study of a robot approaching a person in the context of handing over an object. In *Proceedings of the AAAI Spring Symposium on Multidisciplinary Collaboration for Socially Assistive Robotics*, 18–24, 2007.
- [12] E. Lopez-Damian, D. Sidobre, S. DeLaTour, and R. Alami. Grasp planning for interactive object manipulation. In *Proc. of the Intl. Symp. on Robotics and Automation*, 2006.
- [13] J. Mainprice, E. Sisbot, T. Simeon, and R. Alami. Planning safe and legible hand-over motions for human-robot interaction. In *Proceedings of the IARP Workshop on Tech. Challenges for Dependable Robots in Human Environments*, 2010.
- [14] A. Mason and C. MacKenzie. Grip forces when passing an object to a partner. *Experimental Brain Research*, 163:173–187, 2005.
- [15] K. Nagata, Y. Oosaki, M. Kakikura, and H. Tsukune. Delivery by hand between human and robot based on fingertip force-torque information. In *Proceedings of the IEEE/RSJ Intl. Conference on Intelligent Robots and System (IROS)*, 750–757, 1998.
- [16] S. Shibata, K. Tanaka, and A. Shimizu. Experimental analysis of handing over. In *Proceedings of the IEEE International Workshop on Robot and Human Communication*, 53–58, 1995.
- [17] E. Sisbot, L. Marin, and R. Alami. Spatial reasoning for human robot interaction. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and System (IROS)*, 2007.
- [18] S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. Weghe. Herb: A home exploring robotic butler. *Autonomous Robots*, 2009.
- [19] N. Vahrenkamp, D. Berenson, T. Asfour, J. Kuffner, and R. Dillmann. Humanoid motion planning for dual-arm manipulation and re-grasping tasks. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and System (IROS)*, 2009.