Implicitly Assisting Humans to Choose Good Grasps in Robot to Human Handovers

Aaron Bestick, Ruzena Bajcsy, and Anca D. Dragan

Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

Abstract. We focus on selecting handover configurations that result in low human ergonomic cost not only at the time of handover, but also when the human is achieving a goal with the object after that handover. People take objects using whatever grasping configuration is most comfortable to them. When the human has a goal pose they'd like to place the object at, however, the most comfortable grasping configuration at the handover might be cumbersome overall, requiring regrasping or the use of an uncomfortable configuration to reach the goal. We enable robots to purposefully *influence* the choices available to the person when taking the object, *implicitly helping the person* avoid suboptimal solutions and account for the goal. We introduce a *probabilistic model* of how humans select grasping configurations, and use this model to optimize *expected* cost. We present results in simulation, as well as from a user study, showing that the robot successfully influences people's grasping configurations for the better.

1 Introduction

Handovers happen frequently in collaborative manipulation tasks. Be it when cooking a meal or assembling a device in a factory workcell, we need to pass objects to each other in order to work more effectively. As a result, making robot-to-human handovers seamless has been an area of growing importance in robotics research [1–10].

Imagine unloading the dishwasher with a robot. The robot comes to give you a mug so that you can place it in the cupboard. The way the robot presents you the mug (its position, orientation, and the grasp the robot is already occupying on the object) leads to you having a number of options for how to grasp it, some demanding more effort than others. In the end, the robot's choice of grasp and the object's pose in $SE(3)$ affects how *comfortable* the handover is for you, as well as what you can easily do with the object *after* the handover: how easily you can just lay it down in the desired spot in the cupboard.

Naturally, the robot can take this into account when planning its handover. Prior work has focused on selecting robot grasping configurations [2,4,6–8,10] or object handover locations [1, 3, 5, 9] that maximize the number or range of feasible human grasps [2,7,8] or minimize human ergonomic cost [1,3–6,9,10].

In contrast, our work enables the robot to minimize *expected* cost: our insight is that, although we can't *control* the human's grasp directly, we can model the *probability* that the human will select a particular grasping configuration. This probability distribution can then be used to evaluate the ergonomic

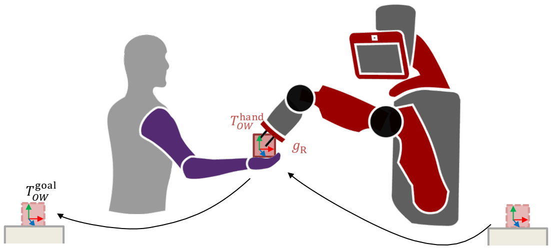

Fig. 1. Setup & Summary. We focus on finding robot handover grasps and object transforms that encourage the human to select good grasps, especially when the human has a next goal for the object. We model how humans select grasping configurations, and leverage that model to minimize expected total ergonomic cost to the human.

cost to the person in expectation, accounting for what they are more or less likely to do.

We propose to model the human as approximately-rational, selecting a grasping configuration with higher probability if its ergonomic cost is lower.

Having such a model of how the human will select a grasp enables the robot to influence the human to select *better* grasps. In particular, we investigate two implications:

Avoiding suboptimal choices for the human, but only when these choices are actually likely: The natural alternatives to having a model of how the user takes the object and minimizing expected cost are either 1) to maximize the total number of grasping configurations available to the user and give them the most flexibility [2], or 2) to minimize average cost to the person [4, 6, 10], without weighting the choices by the probability that the human will actually select them.

Compared to the first, minimizing expected cost enables the robot to produce good configurations as opposed to many configurations. The second, minimizing average cost, also achieves that. However, it also tries to avoid allowing high-cost configurations, because these increase the cost mean. In contrast, in our approach, high-cost configurations do not actually matter, so long as low-cost configurations are available, because the human is very unlikely to select them. Instead, it is suboptimal yet low-cost configurations that are troublesome — these are the configurations that the human might select with high probability, due to the fact that they are not perfectly rational. Our formalism naturally eliminates such choices for the human to the extent possible, helping them select the better options.

Encouraging the human to plan ahead: Usually when we receive an object, it is because we need to do something with it. There is some goal (or set of) goal pose(s) for the object. However, humans are not always very good at planning ahead: they might select a grasping configuration comfortable for taking the object, without thinking of how they will need to manipulate it afterwards. By modeling the human as approximately rational for the handover stage, but myopic to the next stage, we enable robots to minimize expected total cost to the human at both the handover and goal, accounting for this myopia. As a result, the robot avoids handing over an object in ways that allow for low-cost grasps which would have high cost at the goal: if a grasp looks tempting to the human locally, but would make it difficult for the human to satisfy the goal afterwards, then the robot will try to hand over the object in a way that makes such a grasp infeasible.

We tested our approach in simulation and in a user study, suggesting that the robot can successfully influence people to take objects in a way that makes it easier for them to achieve their goal.

2 Related Work

Our main contribution is to explicitly model the probability of the human choosing different available grasps during handover planning, enabling the robot to optimize for *expected* ergonomic cost. A secondary contribution of our work is accounting for the human's goal in the context of minimizing ergonomic cost, enabling the robot to influence the person to select a better grasp. Table 2 categorizes related work along three axes: whether the method accounts for feasibility only or also for ergonomic cost, whether the method accounts for the human's goal, and whether the method accounts for positions of the object only or also grasps.

Table 1. Prior Handover Planning Approaches

	Feasibility Only		Ergonomic Cost	
	H Only	H + G	H Only	H + G
Position Only			[1], [3], [5], [9]	
Grasp Config.	[2]	[7], [8]	[6], [10], [4]	(this paper)

3 Technical Approach

Notation. To choose a handoff configuration, we must select the robot's grasp on the object g_R and the object pose with respect to the world frame T^{hand}_{OW} at which the human will take the object. The object to be handed off allows the human to grasp it at some set of poses $G_H \subset SE(3)$, which we represent as a Task Space Region [11], and discretize to give a finite set of feasible human grasps, so $G_H \triangleq \{g_{H1}, ..., g_{Hk}\}$.

Given a handoff grasp and object pose, $(g_R, T^{\text{hand}}_{OW})$, each possible human grasp g_{Hi} will be reachable with zero or more inverse kinematics (IK) solutions, which we collect into a set $Q^{\text{hand}}_{g_{Hi}}$.

The union of these sets $Q^{\text{hand}}_{(g_R, T^{\text{hand}}_{OW})} \triangleq \bigcup Q^{\text{hand}}_{g_{Hi}}$ gives all the available "taking" configurations available to the human given the robot's choice of $(g_R, T^{\text{hand}}_{OW})$.

A human grasp g_{Hi} also induces IK solutions at the object's goal pose, T^{goal}_{OW}, which we collect in a set $Q^{\text{goal}}_{g_{Hi}}$.

Human Grasp Selection Model. Among possible options, we chose to model the human ergonomic cost as the distance from some ideal nominal resting configuration q^* w.r.t. some metric w:

$$C(q) \triangleq \sqrt{w^{\intercal}(q - q^*)} \quad (1)$$

While we chose this cost function for its simplicity, it would be easy to substitute any other function which maps human limb configurations to ergonomic costs. We would expect superior performance when using cost functions which more accurately capture the human's preferences.

We model the human as approximately-rational, selecting a grasping configuration q at handover time with higher probability when it has lower cost:

$$P(q) \propto e^{-\lambda C(q)} \tag{2}$$

$P(q)$ at the time of handover is normalized over all possible grasping configurations $Q^{\text{hand}}_{(g_R, T^{\text{hand}}_{OW})}$. We can also compute the probability of a grasping configuration given a particular grasp,$P(q^{\text{hand}}|g_H)$, by normalizing over $Q^{\text{hand}}_{g_H}$, and $P(q^{\text{goal}}|g_H)$ at the goal by normalizing over $Q^{\text{goal}}_{g_H}$. Finally, we can compute the probability of a human grasp by summing over all the IK solutions at that grasp: $P(g_H) = \sum_{q \in Q^{\text{hand}}_{g_H}} P(q)$.

Optimization. When the human does not have a (known) goal, we optimize for expected cost at the handover time:

$$\min_{g_R, T^{\text{hand}}_{OW}} \sum_{q \in Q^{\text{hand}}_{(g_R, T^{\text{hand}}_{OW})}} P(q)C(q) \tag{3}$$

When the human does have a goal, we optimize for expected total cost. The expected cost at the goal is based on the probability of each grasp based on what happened at the handover, $P(g_H)$, and the probability of each configuration given that grasp:

$$\min_{g_R, T^{\text{hand}}_{OW}} \sum_{g_H} \left[\sum_{q \in Q^{\text{hand}}_{g_H}} P(q|g_H)P(g_H)C(q) + \sum_{q \in Q^{\text{goal}}_{g_H}} P(q|g_H)P(g_H)C(q) \right] \tag{4}$$

4 Case Studies

We start with two case studies, highlighting the benefits of our approach: eliminating suboptimal yet tempting grasping configurations.

Expected Cost at Handover Time. Fig. 2 compares optimizing for feasibility, average cost, and expected cost, in a scenario where the PR2 robot is handing over a mug to a human. For each case, we take the robot grasp and object transform that arises from the optimization, and compute: 1) the human grasping configuration of minimum cost; 2) the most "risky" human grasping configuration, that is not high-cost cost enough to be easily discarded by the human; 3) all human grasping configurations available; and 4) the histogram of costs for these configurations.

We find that maximizing the number of feasible options can be dangerous, because it might mean the expected cost is rather high, and the best configuration is not as good. Compared to minimizing average cost, we find that minimizing expected cost will allow more high-cost configurations because there is a very low probability for the human to pick them (marked "unimportant" on the histogram), but will allow fewer configurations that have good cost but not great (marked "problematic" on the histogram). These are config-

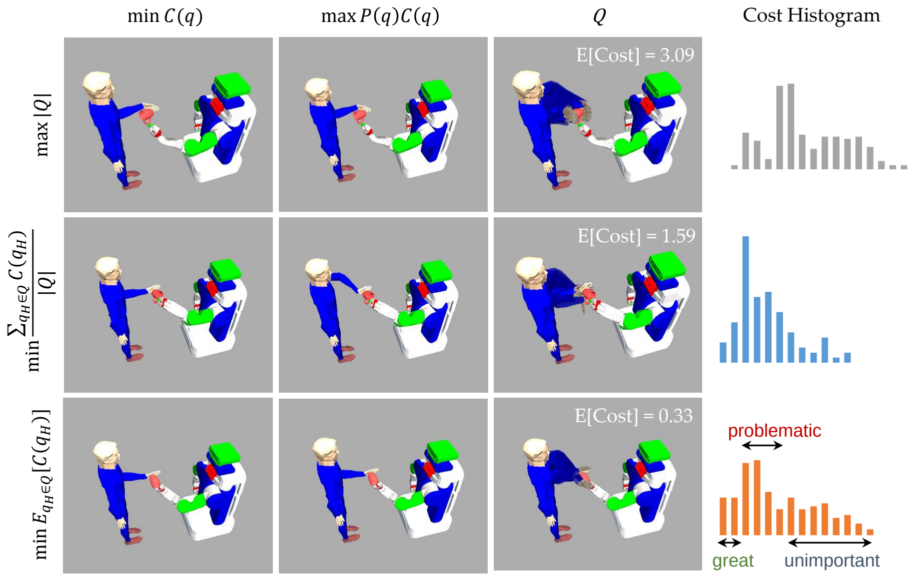

Fig. 2. Case Study w/o Goal - Increasing Great Choices, Reducing OK Choices, Disregarding Bad Choices. A comparison between maximizing the number of feasible human grasping configurations Q (top), minimizing the average ergonomic cost (middle), and minimizing *expected* ergonomic cost (bottom), for the case of a single handover without a known object goal pose. The columns show the most probable human configuration (left), the configuration with the largest contribution to the total cost (middle), and the full space of configurations (right). Our method increases the number of great choices and decreases the number of OK choices which the human might actually pick. It also keeps bad choices if needed, because they have a low probability of being selected anyway.

urations for which the probability is high enough that the human might pick them, but they are not as good as the best configurations.

Experimental Insight 1: ***A robot that models human handover choices can make it more likely that the person will actually select a comfortable handover grasp.***

Expected Total Cost (Handover + Goal). Fig. 3 compares the three approaches from above when accounting for the human goal. Feasibility here accounts for the number of feasible configurations at both the handover and the goal, average cost accounts for cost at the start and goal, and so does expected cost. For each case, we take the resulting robot grasping configuration and compute 1) the human grasping configuration of minimum handover cost, which is what the human will most likely choose if they are being myopic; 2) given this grasp, the configuration of minimum cost at the goal (assuming no regrasp); and 3) the expected cost at the handover and at the goal for each human grasp.

We find that maximizing feasible options can lead to very poor options at the goal. Compared to minimizing average cost, we find that minimizing expected cost is better at eliminating grasps that have low cost at handover time but only allow for high cost at the goal.

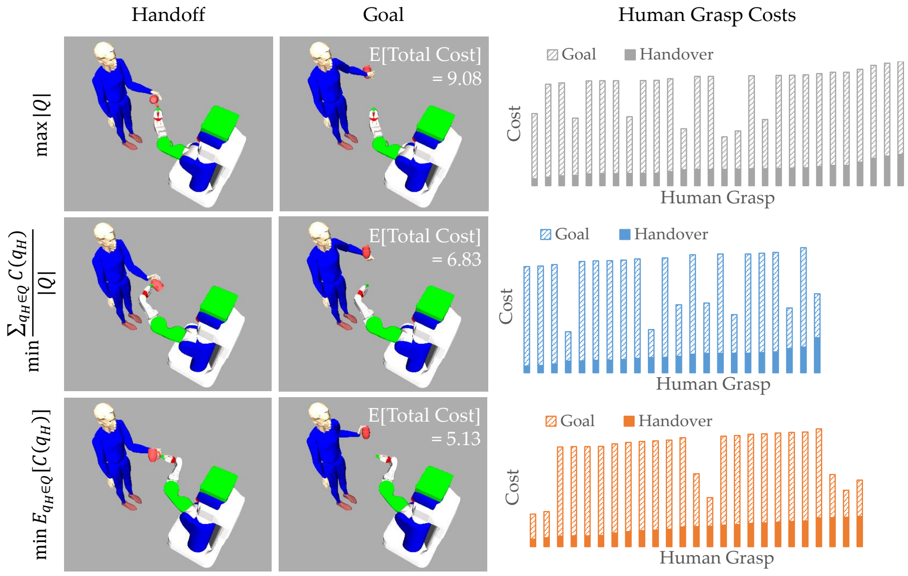

Fig. 3. Case Study w. Goal - Reducing Total Cost. A comparison between maximizing the number of feasible human grasping configurations Q at the handover and goal (top), minimizing the average ergonomic total cost (middle), and minimizing *expected* total cost (bottom). The columns show the most probable human configuration at handover time (left), and at the goal (center), along with a plot of cost for each available grasp to the human. Our method makes it such that the tempting configurations (low cost at handover) also have low cost at the goal.

Experimental Insight 2: ***A robot that models human handover choices can make it more likely that the person will select a handover grasp that also allows for comfortably achieving the goal after the handover.***

5 Simulation Study

Our case studies used a single object and a single goal configuration. Here we expand to an experiment that manipulates both as factors.

5.1 Experimental Design

Manipulated Factors. We manipulate three factors. The first is the *metric* we optimize, as in the case study: maximizing number of feasible options, minimizing average cost, or our metric, minimizing expected cost. The second is the *object* being handed over by the robot: a mug as before, a glass, a pitcher, and a plate, for a total of 4 objects. These objects have vastly different TSR choices. The third is the *goal pose*, for which we use 5 different poses. This leads to a total of 3(metrics) x 4(objects) x 5(goals) = 60 conditions.

Dependent Measures. As in the case studies, we measure expected total cost.

Hypothesis. Our metric is designed to optimize expected total cost (the dependent measure), so we already know it will perform the best. The question remains whether our metric will be better by a significant margin. Our hypothesis is that it will: *Our metric will result in a significant improvement in expected cost compared to the baselines.*

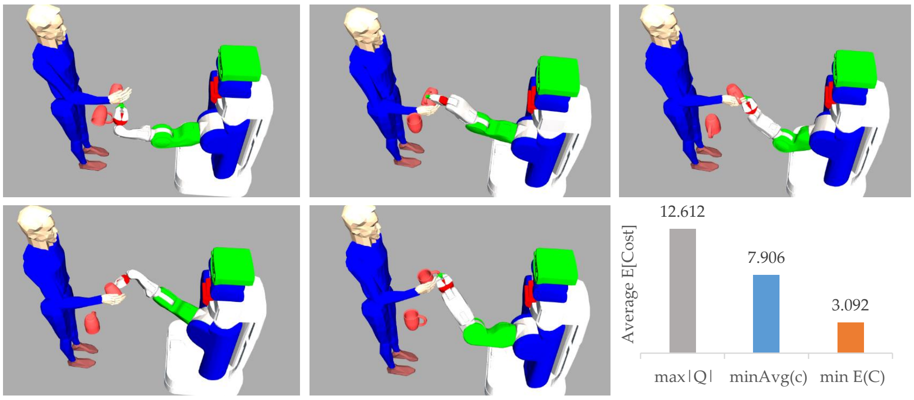

Fig. 4. Optimal Handover for Different Goal Poses. The different goal poses in our experiment lead to different optimal handover configurations for the robot, each selected to minimize expected total cost at the handover time *and* at that particular goal. The chart averages the expected total (handoff + goal) ergonomic costs for each of the three metrics.

5.2 Analysis

We ran an ANOVA with metric as a factor to test differences among the three metrics across objects and goal poses. We found a significant main effect, $F(2, 58) = 1031.07$, $p < .0001$. A post-hoc analysis with Tukey HSD showed that all three metrics were significantly different from each other, with the average cost outperforming maximum feasibility ($p < .0001$) and our metric outperforming average cost ($p < .001$), in line with our hypothesis.

Fig. 4 shows how the robot's grasping configuration changes as the goal pose for the human changes. The robot will present the mug so that the person grabs it by the top when it needs to be placed right side up, by the side when it needs to be placed upside down, etc. In line with our hypothesis, the expected cost was three times lower with our approach compared to the maximum feasibility baseline, and two times lower compared to the minimum cost baseline.

Fig. 5 shows how the robot's grasping configuration changes, for a given goal pose, as the object changes. The robot holds the objects in different ways to ensure that the person can easily grasp them by the side and set them down vertically with ease.

6 User Study

The previous sections tested our method in simulation, assuming users who act according to our model. Real people do not. We conducted a user study to test whether the simulation results generalize, and to explore whether users perceive the improvement brought about by our method.

6.1 Experimental Design

Manipulated Factors. We manipulated three factors. We manipulated the *metric* the robot used to compute its handover configuration, using our metric based on the user model we proposed, $\min E[C]$, and the maximum feasibility baseline $\min |Q|$.

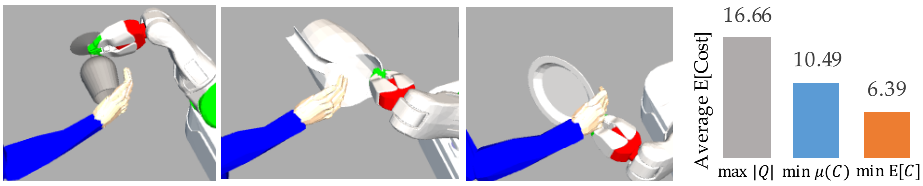

Fig. 5. Optimal Handover for Different Objects. The different objects in our experiment lead to different optimal handover configurations for the robot for a given goal. The chart averages the expected total (handoff + goal) ergonomic costs for each of the three metrics across objects.

We used the mug as the handover object for this experiment, and manipulated the *goal pose* using 10 different poses. In these poses, the mug was placed upside down, upright, and to the side to ensure variance.

Finally, we manipulated *whether the user knows the goal* (Fig. 6). We did this because we wanted to separate the two assumptions our method is making: that users select grasping configurations based on ergonomic cost, and that users are myopic or greedy in this selection, only accounting for ergonomic cost at handover time but not at the goal. Therefore, manipulating the user's knowledge of the goal enables us to test not only how our method performs overall (in realistic situations in which users have a goal and are aware of it), but also whether our method is influencing the users' grasp choice in the way we expected, assuming users are actually myopic (which in reality might or might not be the case). Altogether, this led to 2(metrics) x 10(goals) x 2(knowledge) = 40 conditions.

Subject Allocation. We recruited 9 users (6 male, 3 female, ages 22-29). All of the factors were *within-subjects*, meaning each user experienced all conditions. We counterbalanced the order of the metrics to avoid order effects, and randomized the order of the goals. We split the experiment in two parts, the first in which the user did not know the goal, and the second in which they did:

In Part 1, the robot handed the object to the person at each of the 20 optimal handover configurations (one for each metric and goal pose), but the user was not told the goal used by the planner. We instructed the user to take the object from the robot and immediately drop it in a box. This ensured that no notion of a goal pose would impact the subject's choice of object grasp. This portion of the experiment evaluated the two algorithms' ability to influence the subjects to select a particular grasp when the subject was not aware of a goal, i.e. when the myopic/greediness assumption holds.

In Part 2, a pictoral marker was placed on a table next to the subject indicating the object's goal pose during each handoff. The subject was told that two different algorithms, "Program 1" and "Program 2," would be used during this part of the experiment. We conducted handovers at the same 20 configurations as before, but this time the subject was instructed to place the object at the indicated goal pose. We told the users before each handover which of Programs 1 and 2 was in use. This portion of the experiment evaluated the algorithm's ability to influence people to select ergonomically optimal grasps even when they know the goal, i.e. they are not necessarily myopic. Furthermore, it enabled us to ask users to compare the two methods, seeing if their

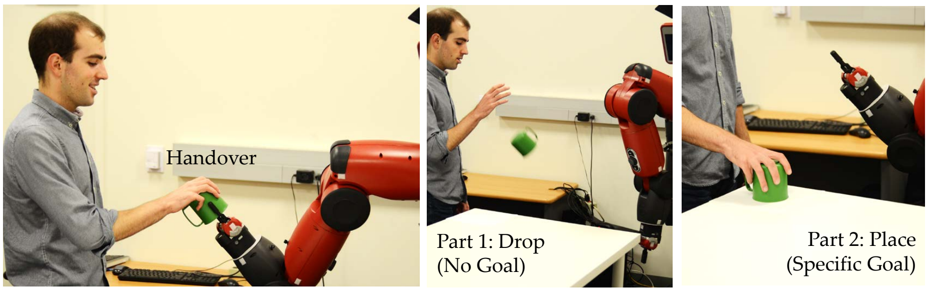

Fig. 6. User Study Setup.

notion of comfort matches ours. If people are actually myopic about the goal when selecting a grasping configuration, then we expect results for this second part to match those from the first part.

Dependent Measures. We used both objective and subjective measures.

Objective: We annotated for each condition which of the 6 TSRs for the mug the person selected. From this, we computed expected cost over all IK solutions at the goal, for all grasps that were feasible at handover time (i.e. had feasible IK solutions), making 2 assumptions: 1) the person follows our ergonomic model, and 2) we know the human kinematics:

$$\text{OM1: } E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in \text{IK(g)}, \forall g \in \text{TSR feas. at handover} \quad (5)$$

To alleviate bias in our results induced by the two assumptions, we introduce 3 additional metrics that break each assumption separately as well as both assumptions together: we break the first assumption by computing average cost (which is the expected cost using a uniform distribution, i.e. $\lambda = 0$) instead of expected ergonomic cost , and we break the second assumption by allowing infeasible grasps that a person with different kinematics might have chosen:

$$\text{OM2: } E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in \text{IK(g)}, \forall g \in \text{TSR feas. at handover} \quad (6)$$

$$\text{OM3: } E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in \text{IK(g)}, \forall g \in \text{TSR} \quad (7)$$

$$\text{OM4: } E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in \text{IK(g)}, \forall g \in \text{TSR} \quad (8)$$

We did not estimate cost at handover time, because we were specifically interested in whether the robot successfully influences users to select grasps that are good at the goal. Indeed, we might see lower handover time costs for the baseline condition because it restricts the users less.

Subjective. After each complete experiment, the subject answered a series of 1-7 Likert-scale survey questions about which program they preferred, which program made their goal easier to accomplish, and which program inspired the most trust in the robot. These capture each subject's subjective opinion about which metric was more effective at making interaction with the robot comfortable and effective.

Hypotheses.

H1. *IF humans are actually myopic when selecting grasping configurations (e.g. when they are not even aware of the goal), our method successfully influences them to select configurations with lower cost at the goal compared to the baseline.*

Table 2. Estimated Human Ergonomic Costs at Goal (Part 1: Users not aware of goal)

Objective Measure	$\min \lvert Q \rvert$	$\min E[C]$
$E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR feas. at handover	12.43	6.02
$E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR feas. at handover	12.41	6.30
$E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR	12.18	11.26
$E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR	12.28	11.45

Table 3. Estimated Human Ergonomic Costs at Goal (Part 2: Users aware of the goal)

Objective Measure	$\min \lvert Q \rvert$	$\min E[C]$
$E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR feas. at handover	11.42	5.37
$E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR feas. at handover	11.52	5.61
$E[C(q^{goal})], \lambda = 10, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR	11.72	11.02
$E[C(q^{goal})], \lambda = 0, \forall q_{goal} \in$ IK(g), $\forall g \in$ TSR	11.84	11.23

H2. *Our method influences people to select configurations with lower cost at the goal compared to the baseline, even when they are aware of the goal.*

H3. *People prefer to work and are more comfortable with a robot using our method compared to the baseline.*

6.2 Analysis

H1. We used results for part 1 of the study, when users are not aware of the goal, to test H1. We first computed Cronbach's α for the four objective measures, which was high at .9036. We thus computed an aggregate goal cost using all four measures.

We then ran a repeated-measures factorial ANOVA on this aggregate, with goal and metric as factors. We found a significant main effect for metric, as expected ($F(1, 179) = 377.83$, $p < .0001$), and a significant main effect for goal ($F(9, 171) = 26.79$, $p < .0001$). However, there was also a significant interaction effect, and so we conducted a Tukey HSD post-hoc, comparing all pairs but compensating for multiple comparisons. The analysis revealed that the expected cost (our) metric resulted in significantly lower cost at the goal than the baseline for 7 out of the 10 goals, all with $p < .03$.

This supports our hypothesis H1, but suggests that the benefit of our method does depend on the choice of the goal pose, with the maximum feasibility baseline being sufficient for some goals.

Table 2 shows the goal ergonomic costs estimated by each of the four measures, averaged across all nine study participants for this part of the study. It shows that pose optimization with $\min E[C]$ gives consistently lower ergonomic cost at the goal than optimization with $\min \lvert Q \rvert$. This difference is particularly marked for the first two measures, which consider only grasps feasible at the handover. These results suggest that expected ergonomic cost can be used to influence humans to choose grasps with good ergonomic properties even when they are completely unaware of the goal.

H2. For part 2, when users were given specific goals, our objective measures again had high item reliability, Cronbach's $\alpha = .8830$. We again computed an

Table 4. Post-Study Survey Results

Statement	min $\lvert Q\rvert$	min $E[C]$	$t(8)$	p
"I prefer Program __"	2.0	6.2	9.73	<.0001
"The robot was helpful when running Program __"	3.4	6.4	6.80	<.0001
"I trust the robot running Program __"	3.7	6.1	4.4	<.01
"The robot understood my goal when running Program __"	2.8	6.4	5.33	<.001
"It was physically easy to do the task when the robot was running Program __"	2.8	6.2	6.50	<.001
"The robot running Program __ handed me objects in a way that made the task easier"	2.0	6.3	9.19	<.0001
"If you had to choose a program you prefer, which would it be?"	0%	100%	-	-

aggregate cost. We again ran a factorial repeated-measures ANOVA, and the results, as expected, were analogous to the results from part 1. We again saw significant main effects, but also a significant interaction between the factors. As before, a post-hoc with Tukey HSD corrections showed that 7 out of the 10 goals saw significantly lower costs at the goal with our method than with the baseline. The set of these 7 goals was almost identical to the one in part 1, with the exception of one goal no longer showing a significant difference, and one goal starting to show a significant difference.

This supports out hypothesis H2: ***our method does not only help users improve performance when we force them to be myopic by not making them aware of the goal – it helps in realistic situations, when users have a goal that they are aware of***. This suggests that people are indeed myopic in their selections of a grasp configuration.

Table 3 shows the goal ergonomic costs estimated by each of the four objective metrics, averaged across all nine study participants for Part 2 of the study, where subjects were instructed to place the object on a pictoral marker at the goal pose after each handover. We see a similar improvement in ergonomic costs when minimizing $E[C]$ versus maximizing $|Q|$.

Here, we found it interesting that the costs dropped slightly across the board. ***This suggests that perhaps when people are aware of the goal they perform slightly better, but that still our method can significantly help them to further improve their performance.***

H3. Table 4 summarizes users' subjective ratings. t-tests showed that ***our method outperformed the baseline in user overall preference, how helpful they thought the robot was, how much they trusted the robot, and how easy it was to do the task. Users thought that the robot understood their goal and that it handed them objects in a way that made their task easier.***

The users' comments were particularly enlightening (here Program 1 refers to the baseline and Program 2 refers to our method):

*"With Program 2, I could move straight from grip to the target with a natural motion. With Program 1, I would sometimes have to contort my arm unnaturally to place the mug correctly."; "Program 1 made it easier to pick up objects but harder to achieve the goal. Program 2 sometimes made it more difficult to pick up objects but achieving the goal was easier."; "Program 1 is an a**hole."*

7 Discussion

Summary. We introduced a model of how people take an object from the robot, and used it to select robot actions that lead to better outcomes for the person. Especially when the person has a goal for the object after the handover, but they are myopic or greedy in their selection of their grasp and do not account for the goal, we have shown that the robot can influence the person's grasp to help them achieve better comfort across the task – at the handover time, but also at the goal time.

Limitations and Future Work. Our work is limited in many ways. We optimize for total ergonomic cost to the person, but it is not clear what this ergonomic cost should be, and it will likely differ from human to human. Furthermore, our study did not measure exactly the ergonomic cost at the goal. Future work might address this by instrumenting the person and the object. Nonetheless, we are encouraged by the subjective results, which align well with our objective estimates. Thus far we only looked at cost at the handover and at the goal, but not at the trajectory the human would plan from one to the other.

Conclusion. Despite these limitations, we are encouraged to see robots being able to influence human actions in a helpful way, making it more likely for them to find good solutions for the task. We are excited to explore further applications of this idea beyond handovers, to human plans more broadly.

References

1. A. Bestick, S. Burden, G. Willits, N. Naikal, S. S. Sastry, and R. Bajcsy, "Personalized kinematics for human-robot collaborative manipulation," in *IEEE International Conference on Intelligent Robots and Systems*, 2015.
2. M. Cakmak, S. S. Srinivasa, M. K. Lee, J. Forlizzi, and S. Kiesler, "Human preferences for robot-human hand-over configurations," in *IEEE International Conference on Intelligent Robots and Systems*, pp. 1986–1993, IEEE, 2011.
3. C.-m. Huang, M. Cakmak, and B. Mutlu, "Adaptive coordination strategies for human-robot handovers," in *Robotics: Science and Systems (RSS)*, 2011.
4. 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," *2nd International Conference on Autonomous Robots and Agents*, pp. 34–39, 2004.
5. J. Mainprice, E. A. Sisbot, T. Siméon, and R. Alami, "Planning safe and legible hand-over motions for human-robot interaction," vol. 2, p. 7, 2010.
6. V. Micelli, K. Strabala, and S. Srinivasa, "Perception and control challenges for effective human-robot handoffs," *RSS 2011 RGB-D Workshop*, 2011.
7. A. H. Quispe, H. B. Amor, and M. Stilman, "Handover planning for every occasion," *IEEE-RSJ International Conference on Humanoid Robots*, 2014.
8. A. H. Quispe, H. Ben Amor, H. Christensen, and M. Stilman, "It takes two hands to clap: towards handovers in bimanual manipulation planning," in *Robotics: Science and Systems (RSS)*, 2014.
9. E. Sisbot and R. Alami, "A human-aware manipulation planner," *IEEE Transactions on Robotics*, vol. 28, pp. 1045–1057, oct 2012.
10. K. Strabala, M. K. Lee, A. Dragan, J. Forlizzi, S. S. Srinavasa, M. Cakmak, and V. Micelli, "Towards seamless human-robot handovers," *Journal of Human-Robot Interaction*, vol. 1, no. 1, pp. 112–132, 2012.
11. D. Berenson, S. Srinivasa, and J. Kuffner, "Task Space Regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.