Affordance-Aware Handovers with Human Arm Mobility Constraints

Paola Ardón*†, Maria E. Cabrera†, Èric Pairet*, Ronald P. A. Petrick*, Subramanian Ramamoorthy*, Katrin S. Lohan* and Maya Cakmak†

***Abstract*— Reasoning about object handover configurations allows an assistive agent to estimate the appropriateness of handover for a receiver with different arm mobility capacities. We propose a method that generalises handover behaviours to previously unseen objects, subject to the constraint of a user's arm mobility levels and the task context. We propose a heuristic-guided hierarchically optimised cost whose optimisation adapts object configurations for receivers with low arm mobility. This also ensures that the robot grasps consider the context of the user's upcoming task, i.e., the usage of the object. To understand preferences over handover configurations, we report on the findings of an online study, wherein we presented different handover methods, including ours, to 259 users with different levels of arm mobility. We find that people's preferences over handover methods are correlated to their arm mobility capacities. We encapsulate these preferences in a statistical relational learner (SRL) that is able to reason about the most suitable handover configuration given a receiver's arm mobility and upcoming task. This proposal is published in [1].**

I. Introduction

Many scenarios in which robots assist humans inevitably involve robot-to-human *object handovers*—the transfer of objects from a robot to a human [2]. Beyond successfully transferring objects, handovers should minimise effort needed from the human. This not only includes effort to *take* the object, but also effort to *use* the object afterwards. For example, imagine a robot handing over a bottle to a person who intends to drink from it. The robot's choice of how to grasp and locate the bottle for the exchange determines how the person will take the object. Hence, in making those choices, the robot should aim to minimise the human's need to extend their arm, offering the bottle in a pose that facilitates drinking without needing to re-grasp the bottle. A method able to adapt robot handovers, with the goal of minimising the person's effort, is particularly convenient for users with arm mobility impairments, where usually the mobility condition changes over time [3].

In [1] we present a method for automatically selecting handover grasps and poses by explicitly taking into account differences in the human receiver's arm mobility while minimising effort. We consider the handover to be composed of a suitable robot grasp that considers the receiver's upcoming task, and an object pose that is safe and reachable depending on the user's arm mobility level. A summary of our approach in [1] is depicted in Fig. 1. Firstly, we pose the problem as hierarchical optimisation with a cost model that adapts to the receiver's mobility constraints, while considering the intended use of the object. Secondly, we evaluate our model through an online survey in which 259 participants with mixed arm mobility limitations rate different handover poses, including the ones generated with our method. An analysis of the responses shows that handover preferences vary significantly across users with different arm mobility capacities, with mobility impaired individuals showing higher preference towards handovers selected with our method. Finally, we extend our method to generate handover configurations for previously unseen objects using a statistical relational learner (SRL). Using our SRL we obtain an average handover pose accuracy of 90.8% across different mobility levels and upcoming tasks with novel objects.

This research was done while the first author was on an academic visit to the University of Washington. It is supported by the Scottish Informatics and Computer Science Alliance (SICSA), EPSRC ORCA Hub (EP/R026173/1) and consortium partners. *Edinburgh Centre for Robotics at the University of Edinburgh and Heriot-Watt University, Edinburgh, Scotland, UK. †Paul G. Allen School of Computer Science & Engineering, University of Washington, Washington, USA. paola.ardon@ed.ac.uk

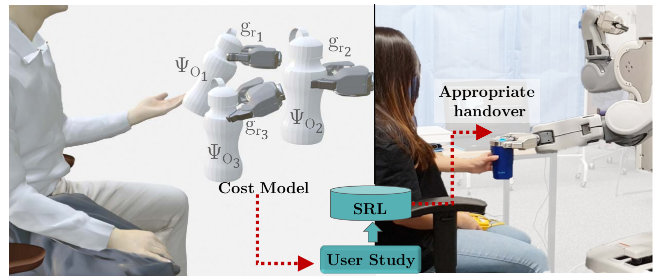

Fig. 1: Summary of [1]. On the left, simulated generation of robot grasps g_r and object poses Ψ_O for handovers using our proposed cost model. On the right, real world deployment of a found suitable handover using our learned SRL model, given the user arm mobility and upcoming task.

II. Handover Optimisation with Mobility Constraints

In [1] we propose a robot-to-human handover method that adapts object configurations to people with different arm mobility. Particularly, we define the handover configurations considering the receiver's (i) upcoming task, to extract an adequate robot grasp, and (ii) arm mobility capacities to adapt the object's pose for the transfer. To achieve such a reasoning model (details in [1]), first we need to design a method that adapts to people with low arm mobility. Then, we need to analyse preferences over handover methods across users with different arm mobility capacities (Section III presents the online user study). This section details the design of our heuristic-guided hierarchically optimised cost model that adapts handovers to users with low arm mobility.

Current robotic handover methods consider preferences over objects and robot grasp configurations that are not designed for receivers with arm mobility impairments. In contrast, with the insight that less effort means more comfort for the receiver [3], [4], we model a handover cost that adapts to users with low arm mobility. The heuristic-guided hierarchically optimised cost model extracts (i) the most suitable robot grasp given the receiver's upcoming task, and (ii) a transfer object configuration located at a reachable yet safe location for the user. We efficiently guide the configuration search through a user-configurable resolution workspace grid map.

The resulting map is composed of $\{\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}\}$ voxels $m_{x,y,z} \in \mathcal{M}_{\{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}\}}$. Each voxel $m_{x,y,z}$ encapsulates: (i) non-controllable human values or constants, in our case the human hand Ψ_{hh}, face pose Ψ_{hf}, and the choice of grasp when receiving the object g_h; and, (ii) cost-constrained variables which are the configurations we want to optimise, in our case the robot grasp g_r, and object pose Ψ_O. As a result, the map is a function of $\mathcal{M}_{\{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}\}}(g_r, \Psi_O)$. We guide the hierarchical optimisation through three costs, as shown in Fig 2. Firstly, we compute an optimal appropriateness cost C_A that gives a suitable robot grasp g_r from a set of grasp affordance configurations $\hat{g}_r \in G_r$. Secondly, using the previously found g_r, we sample for safe object configurations Ψ_O using the safety cost C_S. This cost is constrained to those object poses $\hat{\Psi}_O$ where there is a feasible inverse kinematic solution for the end-effector Ψ_{ree} to proceed with the grasp g_r, denoted as $f(\hat{\Psi}_O, g_r) \neq \emptyset$. Finally, in the reachability cost C_R, we minimise the displacement of the user arm. Given Ψ_O, we inform the search for the closest $m_{x,y,z}$ in $\mathbb{R}^3$ and find the optimal object configuration Ψ_O in SE(3):

$$\begin{gathered}\min_{m \in \mathcal{M}} C_R(\Psi_O) \\ \text{with } \Psi_O = \underset{\hat{\Psi}_O \in m}{\arg\max}\, C_S(\hat{\Psi}_O, g_r) \text{ s.t } \Psi_{ree} \leftarrow f(\hat{\Psi}_O, g_r) \neq \emptyset \\ \text{with } g_r = \underset{\hat{g}_r \in G_r}{\arg\max}\, C_A(\hat{g}_r). \end{gathered} \tag{1}$$

Appropriateness $C_A(\hat{g}_r)$ is calculated in the object affordance space and it extracts the grasp configuration the robot should choose depending on the receiver's future task. Depending on the level of human arm mobility impairment, the hand dexterity may vary considerably and, thus, the human choice of grasps. This is a subject worthy of future study. Although we cannot control the human grasp directly, we can leave the object's part that affords the receiver's chosen action occlusion free. Thus, we implicitly offer the receiver the most suitable grasping region. We reason about g_r, g_h and the object affordances regions a_O using the Markov logic network (MLN) knowledge base (KB) from our earlier work [5]. The KB in [5] is composed of data collected from human users, thus being suitable for the handover task, as well as inferring suitable actions. We consider two sets of grasp configurations: (i) human grasps are configurations inside the object affordance region $g_h \in a_O$, while (ii) robot grasps are outside, $g_r = a_O \setminus g_h$. The final goal in (1) is to choose a g_r that maximises the distance from the closest possible (i.e., most constraining) g_h:

$$C_A(\hat{g}_r) = \min_{g_h \in a_O} d(\hat{g}_r, g_h), \tag{2}$$

thus, guiding appropriate grasps for giver and receiver.

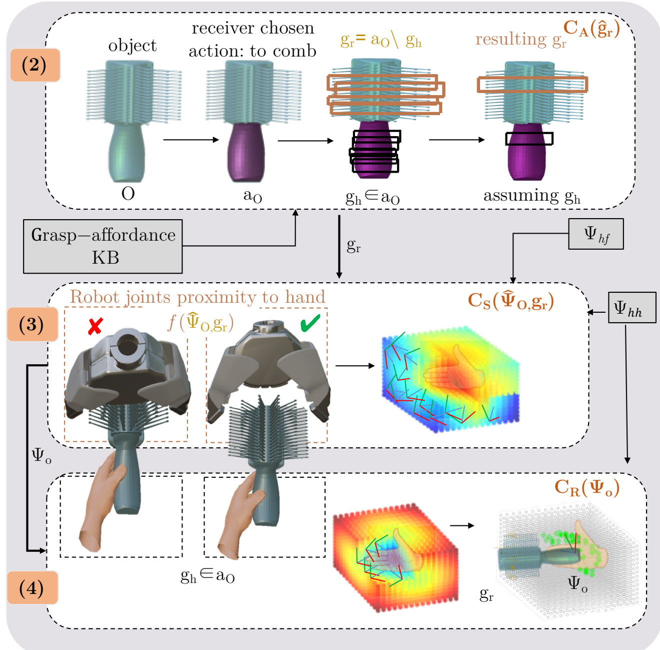

Fig. 2: Summary of our heuristic-guided hierarchically optimised cost model. Each block corresponds to (2)-(4).

Safety $C_S(\hat{\Psi}_O, g_r)$ is considered in terms of distance between the robot to human. The further away from the human user the robot's manipulator is, the safer it is. Thus, we maximise the distance from the object pose $\hat{\Psi}_O$, projected in a_O, to the human hand Ψ_{hh} and face Ψ_{hf}, as well as from the Ψ_{ree} to Ψ_{hh}. We penalise the cost if any of the distances[1] is below a threshold t_h of $5cm$:

$$C_S(\hat{\Psi}_O, g_r) = \begin{cases} d(\hat{\Psi}_O, \Psi_{hh}) + d(\hat{\Psi}_O, \Psi_{hf}) + d(\Psi_{ree}, \Psi_{hh}), & \text{if } d(\cdot) \geq t_h \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

Reachability $C_R(\Psi_O)$ is introduced to minimise the receiver's arm displacement, thus effort [3], [4]. This cost promotes object configurations located as close to the human hand as possible, consequently, adapting to users with low arm capacities. Specifically, $C_R(\Psi_O)$ penalises the human hand movement from the current pose to the implicitly advised grasp g_h. [6] suggested that $75cm$ is a reachable object transfer location, as such, we use it as t_h to penalise greater distances:

$$C_R(\Psi_O) = \begin{cases} d(\Psi_{hh}, g_h), & \text{if } d(\cdot) \leq t_h \\ \infty, & \text{otherwise.} \end{cases} \tag{4}$$

In summary, using (1), the robot obtains the most appropriate robot grasp given the receiver's task and a safe yet reachable object configuration. As a result, adapting handovers to users with low mobility impairments. Fig. 2 summarises the heuristic-guided hierarchically cost model.

III. User Handover Preferences

To implement an inclusive handover method that adapts to people with different arm mobility levels (as presented

[1]The distances between poses are calculated in quaternion using ROS Pose messages.

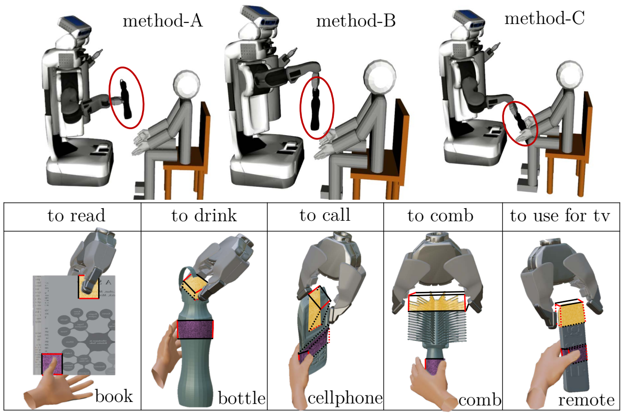

Fig. 3: First row: video frame samples of the methods presented to the users. Second row: objects in the online study with our detected g_h (purple) and g_r (ochre).

in [1]), we need to explore users preferences on handover methods. In this section, we summarise the data collection and evaluation of the users' perception presented in [1].

A. User Study Setup

For the user study we consider three different handover methods. As *method-A*, we implement a handover technique following the guidelines in [6], [7]. These works extract the optimal object transfer point. As in [6], [7], for method-A we set the object transfer point at a distance of $75cm$ from the human body and an arbitrary robot grasp. As *method-B*, we use [8]'s suggested transfer location at $50cm$ and a robot grasp that considers the receiver's upcoming task. As *method-C*, we use our proposal in [1]. The first row of Fig. 3 shows an example of an object's final pose for each handover method. To the participants, neither the methods' details nor name were disclosed. For the reminder of this manuscript *method-C* will be referred to as *ours*.

B. Data Collection

We collected our data through an online survey to guarantee social distancing rules to our participants. Contrary to previous works, our goal is to achieve an inclusive robot-human handover technique. This was done in collaboration with Chest Heart and Stroke Scotland (CHSS)[2] to recruit participants that suffer from arm mobility impairments. Through CHSS, we recruited a total of 9 volunteers. Additionally, we used Amazon's Mechanical Turk platform, to obtain opinions from 250 people with varied arm mobility capacities. We presented to each participant 3 different short clips. Each clip with a different handover method, explained in Section III-A. The 3 clips were randomised among 5 objects using counterbalance design. The 5 objects are of common use on activities of daily living (ADL) by people with amyotrophic lateral sclerosis (ALS) [9] (Fig. 3, second row).

C. Systematic Analysis of User Input

We examine the participants' responses to detect handover preferences. Guided by our hypothesis, we analyse the data to show the influence of arm mobility level and handover technique interaction. Detailed results can be found in [1]. Fig. 4 shows a summary of the findings as extracted from the 5-point scale metric. The higher on the scale the safest, more comfortable or appropriate the handover method is. As mentioned in Section III-B, we create animations for the users to identify their arm mobility level. The participants identified themselves in either of the 4 shown animations: high (H), high-medium (H-M), low-medium (L-M), and low (L) arm mobility. For each of the levels, we illustrate the mean and standard deviation of the three handover methods included in the study. Our analysis involves a normally distributed two-way repeated-measures analysis of variance (ANOVA), using the handover methods and users arm mobility as factors, and user ID as repetitions.

For *perceived safety*, there is no significant difference across methods as rated by users in groups H and H-M. For *perceived appropriateness* of the robot grasp given the user upcoming task, there is a significant difference across method-A and ours for all arm levels. The gap between the method-A and the other two methods is noticeable. On average, our method is perceived positively by the users. Nonetheless, the preference for our method over the other two setups is clearer in participants that reported lower levels of arm mobility (i.e., L-M and L).

We also asked the participants to choose their overall preferred handover method. Table I shows a summary of the preference distribution as related to arm mobility levels.

The participants ranked, from most to least important, the following aspects: (a) safety, (b) comfort, (c) naturalness of the handover, (d) appropriateness of the robot grasp given the receiver's upcoming task, and (e) that the robot moves more than the human to reach the object transfer location (i.e., shared effort). Fig. 5 illustrates the median rank consensus per arm mobility level. The resulting ranking demonstrates the difference in priorities, especially on the extremes of the arm mobility level spectrum. For example, for H and H-M feeling safe and comfortable is the top priority. In contrast, for L-M and L the preference fluctuates between the robot moving more than the human to transfer the object and obtaining the object in a configuration that they can easily use afterwards. As in Table I, Fig. 5 reiterates that users with lower arm mobility prefer a technique that brings the object closer to the hand. Detailed statistics are shown in [1].

Finally, we asked in an *open-ended question* how the users feel about robot-human handover collaboration tasks. Table II shows a sample set of responses. Per arm level,

	participants	Handover method preference overall		
		method-A	method-B	ours
H	179	23.5%	**73.7%**	2.8%
H-M	27	23.1%	**65.4%**	11.5%
L-M	18	7.1%	28.6%	**64.3%**
L	35	3.2%	16.1%	**80.7%**

TABLE I: Distribution of participants per arm mobility level as they chose their preferred handover method.

[2]Health charity for rehabilitation `https://www.chss.org.uk/`

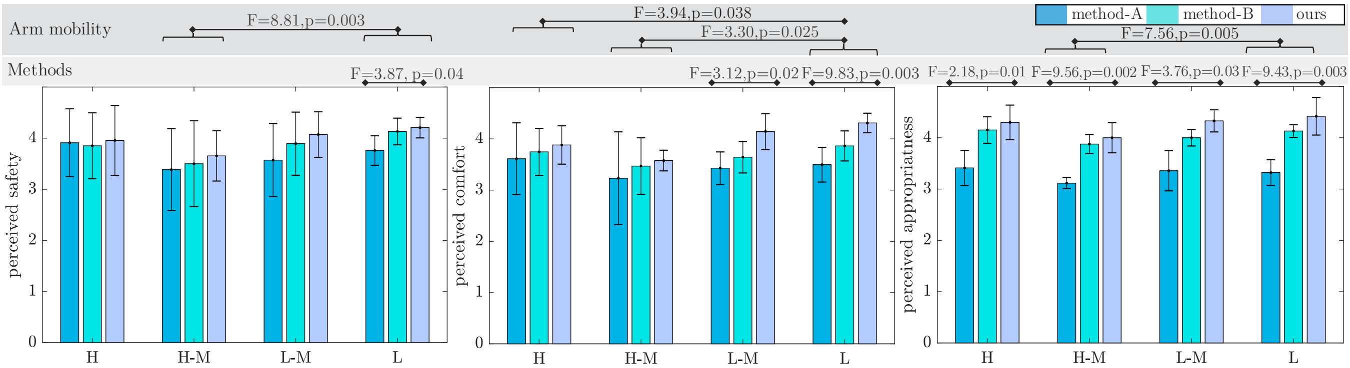

Fig. 4: Mean and variance of participants perception on safety, comfort and appropriateness in a 5-point scale metric. Evaluation includes two-way repeated ANOVA. Only significant statistical results are shown, i.e., with a $p < 0.05$.

	%	Sample of positive responses	%	Sample of negative responses
H	23.5	"... AS LONG AS it is safe I WOULD BE COMFORTABLE"	24	"It DEPENDS ... because of the risk of something happening that the robot cannot ADAPT"
H-M	57.7	" I'D FEEL COMFORTABLE IF the robot's help is convenient ..."	7.7	"It DEPENDS on the object and the setting"
L-M	50	"... I FEEL COMFORTABLE, it behaves USEFULLY"	7.14	"... make sure that I am able to override its functions with voice commands if it malfunctions or behaves UNEXPECTEDLY"
L	64.5	"It looks COMFORTABLE. I am the primary care taker of my sister. This could be really USEFUL for her ..."	3.32	"I WOULD ONLY be worried about dangerous objects"

TABLE II: High recurrent words in users' responses. The % indicates the appearance events of the KEYWORDS sets.

we created a word count of the responses and extracted sets of words appearing with higher frequency. Some of the extracted sets imply a positive opinion about the task, while others suggest a negative or doubtful perception of the robot's performance. Some examples are shown in Table II. By putting these words in context, it is clear that, depending on the mobility level, some participants accept the collaboration with reservations while others perceive the robot as a helper, thus supporting the ranking on Fig. 5.

IV. Conclusions and Future Work

In summary, although users in general prefer a method that considers their upcoming task, there are different preferences related to user arm mobility capacities. Receivers with low levels of arm mobility prefer the robot to perform most of the handover task, while users with high mobility choose to have some freedom and share the task effort. These preferences are then encapsulated in a SRL, the generalisation capabilities of such learned model are detailed in [1].

Our proposal motivates future research in different directions. First, in-person human-robot interaction settings to study the receiver's acceptance levels. This includes further research on the trade-off between reachability and safety optimisation criteria by considering a weighted or a SRL

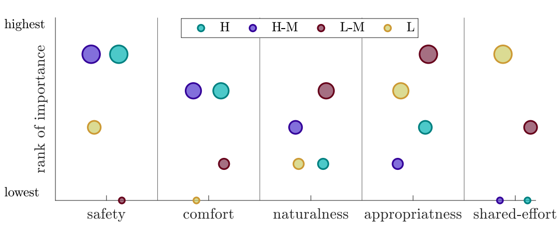

Fig. 5: Median rank consensus per arm mobility of the considered aspects that influence a handover task.

influenced cost function that allows to adapt online. Second, the refining of the appropriateness cost with human grasps extracted from human-human handover studies. Finally, the study of failure and recovery alternatives for cases when the robot grasp is not socially acceptable for the handover task, and ways to enrich our SRL model to prevent such scenarios.

References

[1] P. Ardón, M. E. Cabrera, È. Pairet, R. Petrick, S. Ramamoorthy, K. S. Lohan, and M. Cakmak, "Affordance-aware handovers with human arm mobility constraints," *IEEE Robotics and Automation Letters to be presented in ICRA*, 2021.

[2] V. Ortenzi, A. Cosgun, T. Pardi, W. Chan, E. Croft, and D. Kulic, "Object handovers: a review for robotics," *arXiv:2007.12952*, 2020.

[3] P. Sampson, C. Freeman, S. Coote, S. Demain, P. Feys, K. Meadmore, and A.-M. Hughes, "Using functional electrical stimulation mediated by iterative learning control and robotics to improve arm movement for people with multiple sclerosis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, no. 2, pp. 235–248, 2015.

[4] N. Schweighofer, Y. Xiao, S. Kim, T. Yoshioka, J. Gordon, and R. Osu, "Effort, success, and nonuse determine arm choice," *Journal of neurophysiology*, vol. 114, no. 1, pp. 551–559, 2015.

[5] P. Ardón, È. Pairet, R. P. A. Petrick, S. Ramamoorthy, and K. S. Lohan, "Learning grasp affordance reasoning through semantic relations," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4571–4578, 2019.

[6] H. Nemlekar, D. Dutia, and Z. Li, "Object transfer point estimation for fluent human-robot handovers," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2627–2633, IEEE, 2019.

[7] J. Mainprice, M. Gharbi, T. Siméon, and R. Alami, "Sharing effort in planning human-robot handover tasks," in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pp. 764–770, IEEE, 2012.

[8] S. Parastegari, B. Abbasi, E. Noohi, and M. Zefran, "Modeling human reaching phase in human-human object handover with application in robot-human handover," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3597–3602, IEEE, 2017.

[9] Y. S. Choi, T. Deyle, T. Chen, J. D. Glass, and C. C. Kemp, "A list of household objects for robotic retrieval prioritized by people with als," in *2009 IEEE International Conference on Rehabilitation Robotics*, pp. 510–517, IEEE, 2009.