Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

Learning Shared Control by Demonstration for Personalized Wheelchair Assistance

Ayse Kucukyilmaz and Yiannis Demiris

Abstract—An emerging research problem in assistive robotics is the design of methodologies that allow robots to provide personalized assistance to users. For this purpose, we present a method to learn shared control policies from demonstrations offered by a human assistant. We train a Gaussian process (GP) regression model to continuously regulate the level of assistance between the user and the robot, given the user's previous and current actions and the state of the environment. The assistance policy is learned after only a single human demonstration, i.e. in one-shot. Our technique is evaluated in a one-of-a-kind experimental study, where the machine-learned shared control policy is compared to human assistance. Our analyses show that our technique is successful in emulating human shared control, by matching the location and amount of offered assistance on different trajectories. We observed that the effort requirement of the users were comparable between human-robot and human-human settings. Under the learned policy, the jerkiness of the user's joystick movements dropped significantly, despite a significant increase in the jerkiness of the robot assistant's commands. In terms of performance, even though the robotic assistance increased task completion time, the average distance to obstacles stayed in similar ranges to human assistance.

Index Terms—Assistive robotics; assistive mobility; Gaussian processes; haptic shared control; human factors; human-like assistance; intelligent wheelchairs; learning by demonstration; user modelling

1 Introduction

The use of powered wheelchairs is essential to enhance the independent living of individuals with mobility limitations. However, the increase in the utility of powered wheelchairs comes along with a raising concern about the safe use of the technology. Reports indicate a growing number of accidents, especially among elderly users, proportional with the number of devices [1]. One way to promote safe and effective use of powered mobility is to implement intelligent shared control mechanisms onto powered wheelchairs, which integrate the capabilities of the user and the assistive technology in order to provide proactive and effective assistance to the user.

Ideally, the assistance policy applied by an intelligent wheelchair system should provide assistance targeted at individual user needs and benefits. We suggest that an intelligent wheelchair would be most effective if it can emulate the operation of a human assistant, e.g. an occupational therapist. Different studies have investigated human-human interaction to learn from the behavioral mechanisms utilized by humans (e.g. [2], [3], [4], [5], [6]). Recently, Ganesh et al. [7] showed that the motor performance of physically interacting individuals improves through motor adaptation. Later, Takagi et al. [8] proposed a model, where interacting dyads use haptic information to estimate each other's goals to improve individual performance. In essence, the idea of learning the intricacies of the interaction between two humans and replicating them for generating robotic assistance policies on a physical robot is new. In line with this goal, in this study, we take a first step toward developing a shared control wheelchair that provides personalized assistance to its user by learning a continuous policy from a human.

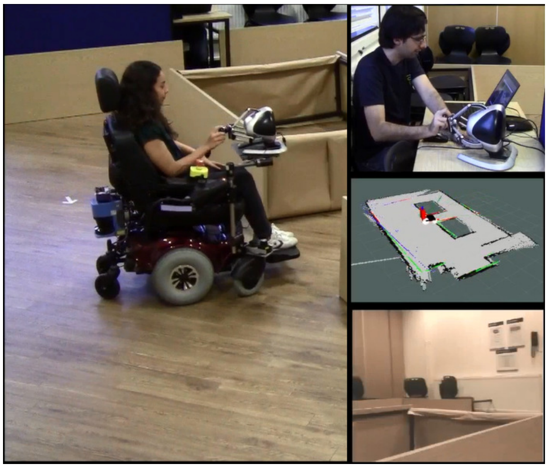

Fig. 1: Learning Shared Control by Demonstration Scenario: A remote assistant (upper right), who is provided with a visualisation (center right) and real-time camera capture (bottom right) of the environment, offers assistance to the user (left) over haptic shared control. The robotic system, in turn, models an assistive policy through observations of the environment and the user commands.

The shared control technique proposed in this paper adopts a user modelling approach to model the assistant's behavior to aid a wheelchair user. We propose a triadic learning shared control by demonstration methodology, in which an intelligent wheelchair learns an assistive policy through modelling the dynamically coupled demonstrations

• *Ayse Kucukyilmaz is with the Lincoln Centre for Autonomous Systems research group in the School of Computer Science, University of Lincoln, Brayford Pool, Lincoln, LN6 7TS, UK.*
Yiannis Demiris is with the Personal Robotics Laboratory, Department of Electrical and Electronic Engineering, Imperial College London, Exhibition Road, London, SW7 2BT, UK.
E-mail: akucukyilmaz@lincoln.ac.uk, y.demiris@imperial.ac.uk

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

given by a remote human assistant.

Learning-by-demonstration (LbD) has been successfully applied to human-robot interaction for over a decade to enable robots to imitate tasks performed by a human. In LbD, the human acts as the *teacher*, and physically demonstrates the task for the robot to generalize from [9], [10], [11], [12]. Up to now, LbD has mostly focused on dyadic human-robot interaction scenarios, where the human is involved in the process only during demonstration, while the robot models the demonstrated actions and reproduces them autonomously.

In this paper, we focus on a complex triadic scenario, where the user continuously interacts with the wheelchair during both demonstration and reproduction. In our setup, the teacher is not the user of the wheelchair, but a remote human, who provides assistance when needed (See Fig. 1). The proposed technique is outlined in the supplementary video supplied as the online resource.

2 Background

Powered wheelchair technologies went through important advances in the last 40 years. Initial efforts implemented simple safeguarding techniques involving collision avoidance and trajectory following to make powered wheelchairs intelligent [13], [14], [15], [16], [17]. Later on, the main trend shifted toward estimating user goals and intentions. Boy et al. proposed a mechanism that modulates the intelligent wheelchair's trajectory corrections based on the user's capabilities by defining control mechanisms tailored to specific disabilities [18]. Carlson and Demiris defined possible actions in a particular environment, and dynamically predicted the most probable actions that shall be taken in near future to correct the orientation of the wheelchair [19], [20]. Demeester et al. used Bayesian decision theory to estimate the certainty of users in a navigation task [21], followed by a comparison of popular approaches, namely Maximum Likelihood (ML), Maximum A Posteriori (MAP), and a greedy Partially Observable Markov Decision Process (POMDP), in order to evaluate navigation plans [22]. Later, Hüntemann used the Bayesian approach in conjunction with Gaussian processes to model user driving behavior to reason about the user's local navigation plan [23].

Related to the estimation of user intentions, is the question of setting the shared control parameters in assisted driving. In essence, the shared control paradigm allows the robot to change its autonomy level when interacting with the human [24]. The autonomy changes can be triggered explicitly by the user or hard coded through the interaction. However, such implementations are inherently problematic, as the former increases the mental load of the user and the latter limits the system's ability to accommodate different user profiles. These shortcomings can be addressed by dynamically controlling the autonomy level of the robot depending on user inputs [25], [26], [27], [28]. Fernandez-Carmona et al. [29] and Li et al. [30] used task performance metrics to adapt the user's control authority during robot-assisted wheelchair operation. Goil et al. trained a Gaussian mixture model to extract task variability in terms of angular velocity as a measure of task complexity [31]. This variance was used to set a blending coefficient to control the allowable angular velocity that the user can apply.

Even though policy blending offers a structured solution to manage the amount of assistance, it does not define the assistance policy itself. Learning the assistive policy is challenging, since wheelchair users' control commands may not always be sufficient per se to construct a consistent assistive model (such as in the case of toddlers and people with extreme cognitive or physical difficulties). Hence, it is imperative to understand how assistance shall be shaped to individual needs.

Recently, Soh and Demiris proposed that a robotic wheelchair can learn from *human guidance*, where a human can correct the user by means of taking over the control of the wheelchair whenever necessary [32], [33]. Assuming the coherence of human guidance, their approach learned *when* to override user commands to assist the user. In the current study, we take this idea one step further by learning continuous assistance policies during interaction, where the assistant does not take over the control of the task, but shares it with the user at all times over a bidirectional haptic interface. We envision that this will provide the user more autonomy in situations where (s)he comes up with navigation plans different from those of the assistant.

Another question relevant to the purpose of this study is whether it is necessary to integrate a powered wheelchair with haptic feedback capabilities, or not. Previous studies indicated the effectiveness of force feedback in reducing collisions [34], [35], [36], [37] and speeding up the learning process to gain the necessary abilities for driving a wheelchair [38]. Marchal-Crespo et al. provided users with modulated haptic feedback, where they gradually decreased the stiffness of the haptic joystick, to guide the users toward a given trajectory [39]. In a controlled study with children of ages between 4 and 9, they showed that modulating the forces in this manner allowed the users to gain necessary driving skills faster when compared to those in the control group. Later, Morere et al. presented a user study with 5 children with disabilities, to illustrate personal differences in driving styles and the received benefit from haptics [40].

These studies use haptic feedback as a guidance mechanism and illustrate its utility when compared to no guidance conditions; but do not address the necessity of having haptic feedback to close the control loop. By turning the haptic feedback on and off while robotic assistance is active, the current study puts an effort to investigate whether humans benefit from haptics when working with an assistant.

3 System Overview

This section describes the haptics-enabled wheelchair ARTA and the shared control framework, through which the remote human assistant presents assistance to the user.

3.1 Robotic Platform

In the experiments, we used ARTA (Assistive Robotic Transport for Adults) powered wheelchair platform, devel-

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

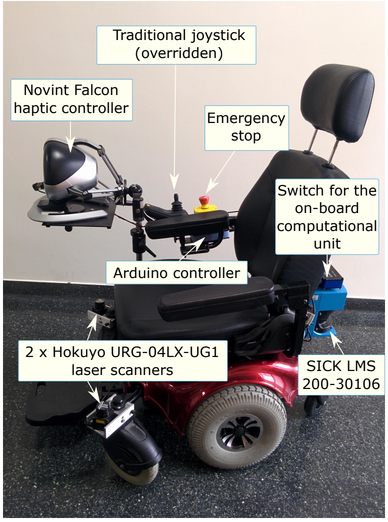

Fig. 2: ARTA (Assistive Robotic Transport for Adults) haptics-enabled powered wheelchair platform

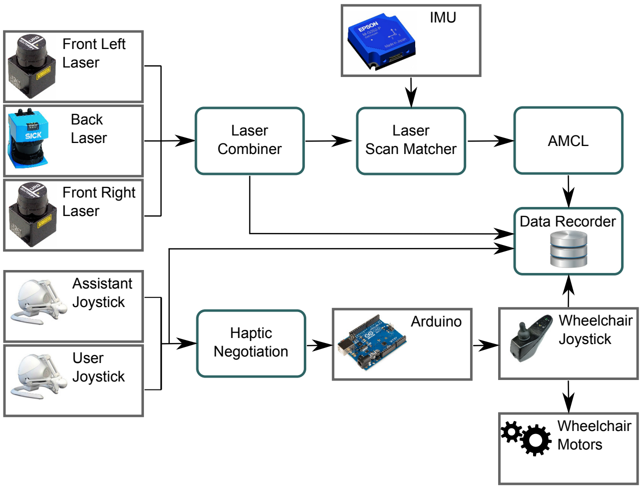

Fig. 3: Software architecture for ARTA

oped at Imperial College London, Personal Robotics Laboratory (Fig. 2). ARTA is based on a mid-wheel powered wheelchair. An Arduino controller is programmed to tap into the PG Drives VSI controller so as to override the joystick commands issued to drive the wheelchair. ARTA is equipped with two Hokuyo URG-04LX-UG1 laser scanners (front left and front right), a SICK Laser Measurement Sensor LMS-200-30106 (back), a Phidgets spatial 3/3/3 inertial measurement unit (IMU), and a wireless router. An on-board computational unit is installed to process and integrate multi-sensory information, manage the haptic shared control module, and run higher-level tasks, such as assistive policy generation.

Fig. 3 summarizes the software components of ARTA, which are built over the system described in [41]. The software is developed using ROS, and each sensor is managed by a separate ROS node, denoted via boxes with rounded corners in Fig. 3. The *Laser Combiner*[1] node combines the laser scan readings from three laser rangers into a single coherent ROS message as explained by [42]. The *Laser Scan Matcher*[2] implements incremental laser scan registration using IMU measurements to estimate the change of the orientation angle of the wheelchair. The *AMCL* node[3] runs the adaptive Monte Carlo localisation algorithm to localise the wheelchair on a pre-existing laser-based map using the combined laser scan and transform messages, in order to output pose estimates. The *Haptic Negotiation*[4] node gets the joystick commands from the user and the assistant, and combines them into a joint command that drives the wheelchair.

3.2 Haptic Shared Control

By default, ARTA is driven by a spring-centered, proportional control joystick that provides full directional control within a 360° circle. In order to enable haptic feedback, the traditional joystick commands are overriden using two haptic controllers, shown in Fig. 2. In our setup, each agent (i.e. the user and the assistant) uses a Novint Technologies Falcon haptic controller, with up to 9 N continuous force feedback capability and 400 dpi position resolution, to collaboratively control the movement of the wheelchair. The user's haptic controller is placed on top of ARTA and directly connected to the computational unit, whereas the assistant's controller is connected to a remote PC that communicates with ARTA over the network.

The wheelchair motion is realized by manipulating the haptic controllers on the horizontal plane: forward/backward movements in z-axis are used for controlling the linear velocity and leftward/rightward movements in x-axis are used for setting the angular velocity. A deadband was used to avoid small joystick commands, the magnitude of which are less than 0.3, from invoking motion on the wheelchair. The position of the haptic interface point directly translates to linear and angular velocity commands, which vary in $[-1\ 1]$ range. These commanding velocities are sent to the controller to generate motor actuation.

The shared control scheme, which combines the operations of the user and the assistant, is based on the *haptic negotiation model* as described in [25]. Fig. 4 depicts the haptic negotiation model that enables shared control on the wheelchair. The user and the assistant drive the wheelchair by moving the haptic controllers, the tooltip positions of which are mapped to haptic interaction points HIP_1 and HIP_2 in the virtual world, where HIP_1 denotes the user's haptic interaction point and HIP_2 denotes that of the assistant. The axes lengths for the positions of HIP_1 and HIP_2 are normalized to range between -1 and 1, regarding the minimum and maximum encoder limits of the haptic device. The operations of the agents are merged by introducing a negotiated interface point (NIP)

1. Laser combiner software, written by Harold Soh, is available from the software section in http://imperial.ac.uk/PersonalRobotics.

2. Software, written by Ivan Dryanovski and William Morris, is available from http://www.ros.org/wiki/laser_scan_matcher.

3. Software, written by Brian Gerkey and Andrew Howards, is available from http://www.ros.org/wiki/amcl.

4. Haptic negotiation software written by Ayse Kucukyilmaz, is available from the software section in http://imperial.ac.uk/PersonalRobotics

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

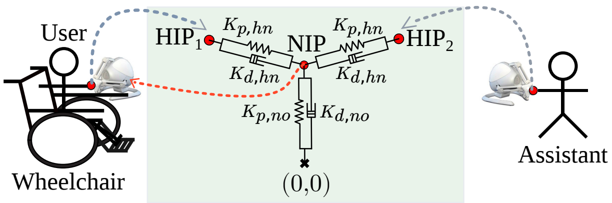

Fig. 4: Haptic negotiation model for controlling ARTA

that is directly connected to the HIPs through virtual spring-damper systems with zero rest lengths. Due to the normalization of HIP ranges, the NIP position is also guaranteed to lie between $[-1\ 1]$, which can be directly used as the commanded velocity sent to the controller in order to generate the required motor actuation. In order to simulate the spring-centered joystick behavior, NIP is also attracted towards the origin by a spring-damper system.

The springs and dampers between interaction points serve as proportional-derivative (PD) controllers driving the interaction points. Hence, at each time step, the position of NIP is influenced by three forces:

$$\begin{aligned}\mathbf{F}_{H1N} =& K_{p,hn}\left(\mathbf{x}_{HIP_1}-\mathbf{x}_{NIP}\right)+K_{d,hn}\left(\dot{\mathbf{x}}_{HIP_1}-\dot{\mathbf{x}}_{NIP}\right)\\ \mathbf{F}_{H2N} =& K_{p,hn}\left(\mathbf{x}_{HIP_2}-\mathbf{x}_{NIP}\right)+K_{d,hn}\left(\dot{\mathbf{x}}_{HIP_2}-\dot{\mathbf{x}}_{NIP}\right)\\ \mathbf{F}_{ON} =& -K_{p,no}\mathbf{x}_{NIP}-K_{d,no}\dot{\mathbf{x}}_{NIP},\end{aligned} \tag{1}$$

where $\mathbf{x}_i$ and $\dot{\mathbf{x}}_i$, $i \in \{HIP_1, HIP_2, NIP\}$, respectively denote the positions and velocities of HIP$_1$, HIP$_2$, and NIP, whereas K_p and K_d values respectively denote proportional and derivative gains. Utilized controller gains are hardware specific and set experimentally as , $K_{d,hn} = 0.25\,\mathrm{N/m}$, $K_{d,no} = 0.2\,\mathrm{N/m}$, and $K_{p,hn} = K_{p,no} = 15\,\mathrm{N/m}$.

$\mathbf{F}_{H1N}$ and $\mathbf{F}_{H2N}$ denote the forces that are applied on NIP and reciprocally felt by the user and the assistant, whereas $\mathbf{F}_{ON}$ is the force acting on NIP due to its attraction towards the origin. NIP is represented as a unit mass of $m = 1$ kg, governed by Newton's second law of motion. At the end of each time step of the haptic feedback loop, the position of NIP for the next time step is calculated using semi-implicit Euler integration:

$$\begin{aligned}\mathbf{F}_{NIP} =& \mathbf{F}_{H1N}+\mathbf{F}_{H2N}+\mathbf{F}_{ON},\\ \ddot{\mathbf{x}}^t_{NIP} =& \mathbf{F}_{NIP}/m,\\ \dot{\mathbf{x}}^{t+1}_{NIP} =& \dot{\mathbf{x}}^t_{NIP}+\ddot{\mathbf{x}}^t_{NIP}\Delta t,\\ \mathbf{x}^{t+1}_{NIP} =& \mathbf{x}^t_{NIP}+\dot{\mathbf{x}}^t_{NIP}\Delta t+\ddot{\mathbf{x}}^t_{NIP}\Delta t^2,\end{aligned} \tag{2}$$

where $\ddot{\mathbf{x}}^t$ and $\dot{\mathbf{x}}^t$ respectively denote the acceleration and the velocity of NIP at time t, and $\Delta t = 1$ msec.

This setup allows the agents to get physically coupled, so that each is able to feel the forces exerted by the other party as if holding the same joystick.

4 Methodology

In this study we are interested in making an inference about the relationship between observation inputs (i.e. the environmental information and the user's control commands) and appropriate assistance commands. For this purpose, a Gaussian process (GP) method for regression is chosen [43]. GP is a generalization of a multivariate Gaussian distribution to infinitely many variables. Given training data $\mathcal{D} = (\mathbf{X}, \mathbf{y})$, consisting of N observations of D-dimensional input vectors, $\mathbf{X} = [\mathbf{x}_i^{D\times 1}]_{i=1}^N$, and corresponding real-valued targets, $\mathbf{y} = [y_i]_{i=1}^N$, GP represents a prior distribution over all functions of the form $f : \mathbf{X} \mapsto \mathbb{R}$.

The entire function evaluation is drawn from a multivariate Gaussian distribution, fully specified by a mean vector μ, which is often assumed to be a zero function, and a covariance matrix $\mathbf{K}$:

$$f = \{f_1, \cdots, f_n\}^T \sim \mathcal{N}(\mu, \mathbf{K}). \tag{3}$$

In order to deal with noisy observations, Gaussian noise with a zero mean and a variance of σ_n^2: $y = f(\mathbf{x}+\epsilon)$, where $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$, is incorporated into the kernel matrix:

$$\mathbf{\Sigma}_N = \mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I}, \tag{4}$$

where $\mathbf{I}$ is the $N \times N$ identity matrix and σ_n^2 is the hyper-parameter denoting noise variance.

For a test instance $\mathbf{x}_*$, the predictive distribution of the target value, y_*, is given by

$$\begin{aligned}&p(y_*|\mathbf{x}_*, \mathcal{D}, \hat{\Theta}) =\\ &\mathcal{N}(y_*|\mathbf{k}^\top\mathbf{\Sigma}_N^{-1}\mathbf{y}, k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}^\top\mathbf{\Sigma}_N^{-1}\mathbf{k} + \sigma_n^2),\end{aligned} \tag{5}$$

where $\mathbf{k}^{N\times 1}$ consists of elements $[\mathbf{k}]_i = k(\mathbf{x}, \mathbf{x}_i)$, and $\hat{\Theta}$ denotes the optimal hyper-parameters related to the covariance function and the error variance, which can be estimated by maximizing the marginal-likelihood [44]:

$$p(\mathbf{y}|\mathbf{X}, \Theta) = \mathcal{N}(\mathbf{y}|0, \mathbf{\Sigma}_N). \tag{6}$$

In order to facilitate the generalization of the model to different environments, we used only context-invariant inputs to learn assistive joystick commands. Specifically, the output of the model was the assistant's joystick commands, while the inputs were:

- 16 laser scan values denoting distances to nearby obstacles. Raw scan data were sub-sampled every 22.5° over the full 360° circle around the wheelchair.
- Linear and angular velocity of the wheelchair
- User's joystick commands

The model was trained using a squared exponential covariance kernel:

$$k\left(f(\mathbf{x}_p), f(\mathbf{x}_q)\right) = \exp\left\{-\frac{|\mathbf{x}_p - \mathbf{x}_q|^2}{2l^2}\right\} \quad , \tag{7}$$

where l is a free hyper-parameter that defines the characteristic length-scale of the process, The initial length-scale of the kernel was set to 6 seconds, and optimized using maximum likelihood estimation. In order to reduce the time required to train the model, automatic relevance determination (ARD), which allows the model to identify the directions in the input space that exhibit high relevance, is turned off [45].

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

5 Experiment and Data Collection

An interaction scenario, in which non-disabled participants were asked to manoeuvre with ARTA within the experiment area, as shown in Fig. 1, was designed to a) evaluate the ability of our technique to emulate human assistance, and b) compare usage characteristics for human-human and human-robot interaction scenarios under haptic and non-haptic feedback conditions. Note that as only non-disabled subjects were involved in the study, our purpose is to provide a comparison between human-provided guidance and robotic assistance policies, not to show therapeutic benefits of the proposed technique.

The user study exhibits a mixed design, which includes one between-subjects and two within-subjects variables: 1) interaction type (human-human versus human-robot scenarios), 2) feedback condition (haptic versus non-haptic feedback), and 3) track type (complex versus simple).

5.1 Independent Variables

Interaction Type

The study involved two separate groups of participants: The first group participated in a human-human interaction scenario, called the HHI experiment, where a remote human assistant presented guidance to the user via shared control over the course of all trials. The data collected within the HHI experiment is used to validate the ability of the GP regression model to learn human provided assistance, as will be discussed in Section 6. The second group was involved in a human-robot interaction scenario, dubbed the HRI experiment, where the assistance policy was learned from a single initial human demonstration and was applied by the robot to aid the user in following trials. The inclusion of this factor provides valuable observations that allow a comparison between human assistance and robotic assistance policies. We assume that the human assistance serves as the *ideal* assistance policy, hence it is desirable for the robot to display comparable performance to that of the human. For this purpose, the results of the HRI experiment are contrasted with the results of the HHI experiment to reflect on how the proposed robotic assistance policy compares to actual human assistance, as detailed in Section 7.

Feedback Condition

In order to evaluate the utility of haptics during assisted driving, we included a within-subjects factor that controls the existence of haptic feedback at the user's controller. Hence, apart from comparing human and robotic assistance, this study also explores how haptic communication affects interaction performance and user operation during collaborative driving on a powered wheelchair.

Feedback condition has two levels, which relates to the existence and lack of force feedback at the user's controller:

1) **Haptic Condition (H):** In the HHI experiment, both the assistant and the subject feel the forces exerted by their partner. In other words, the haptic devices are fed back with $\boldsymbol{F}_{d,1} = -\boldsymbol{F}_{H1N}$ and $\boldsymbol{F}_{d,2} = -\boldsymbol{F}_{H2N}$, as computed in (1), so that the assistant's movements are mirrored at the subject's haptic device and vice versa. In the HRI experiment, since the assistant is out of the loop, the haptic controller at the assistant's side is deactivated. However, the assistance commands generated by the model are used to compute $\boldsymbol{F}_{H1N}$ and $\boldsymbol{F}_{H2N}$ in the same way as the HHI experiment, and the user is fed back with $\boldsymbol{F}_{d,1} = -\boldsymbol{F}_{H1N}$.
2) **Non-Haptic Condition (NH):** The subject does not get any force feedback due to the actions of the assistant. This is realized by feeding back the forces to the subject as if his/her haptic device is directly connected to the origin:

$$\boldsymbol{F}_{d,1} = -K_{p,hn}\boldsymbol{x}_{HIP_1} - K_{d,hn}\dot{\boldsymbol{x}}_{HIP_1}. \tag{8}$$

In the background, $\boldsymbol{F}_{H1N}$ is computed using (2) and is used to control the movement of NIP as in the case of the H condition. This enables the wheelchair to act alike under both H and NH. The only difference is that the subject's haptic device is not actuated as a result of the actions of the assistant. In the HHI experiment, the assistant is provided with forces exerted by the partner as in H condition: $\boldsymbol{F}_{d,2} = -\boldsymbol{F}_{H2N}$.

Track Type

The user was asked to perform the task on two different trajectories. Fig. 5 depicts a bird's eye view sketch of the experiment area, where the desired trajectories are marked in dotted red lines. The first trajectory, as shown in Fig. 5(a), involved maneuvering in between two obstacles, and looked as if the subject is drawing a figure 8 on the track, and will be called the *figure 8 track* in the rest of the paper. The second trajectory, as depicted in Fig. 5(b), required the subjects to drive around the two obstacles on a rectangle shaped track, and will be called the *rectangular track* in the rest of the paper. The rectangular track is designed to serve

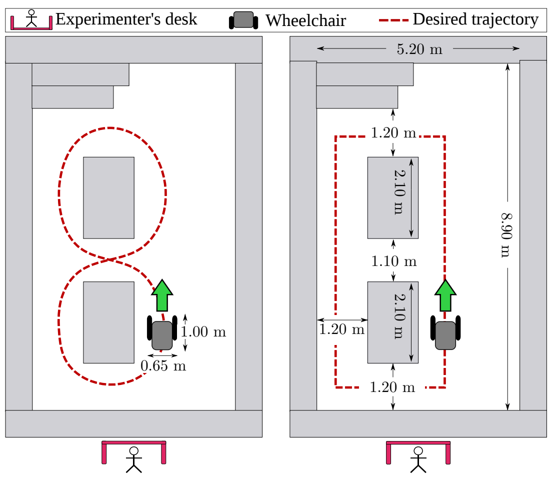

(a) Trajectory 1: figure 8 track (b) Trajectory 2: rectangular track

Fig. 5: Bird's eye view of the lab area and the tracks used in the experiments. The red dotted lines denote the desired trajectories. The wheelchair, experimenter's desk, and obstacles are shown.

as a simpler task than the figure 8. Also, as the user model is learned over the figure 8 trajectory, the rectangular track performance is used as indicative of the generalizability of the assistance model to different trajectories.

5.2 Procedure

At the beginning of the experiment, the subjects gave informed consent for their participation in the study. They were verbally instructed about the wheelchair and its operation with the haptic joystick, and were informed about the trajectories. During the experiment, the experimenter sat behind a curtain, which hid the haptic device from the subject's sight. The subjects were asked to pay attention to any assistance they might feel during the sessions, noting that it is possible that there might not be any assistance at all. However, they were not informed about the conditions nor about the existence of external human assistance.

The experimenter acted as a remote assistant, who monitors the operation of the wheelchair by observing its motion through rviz 3D visualization tool for ROS. She was provided with a frontal camera view of the environment, as well as laser scan readings matching the map of the environment as shown in Fig. 1. The assistant was experienced in teleoperating the wheelchair using the haptic joystick, and was instructed to observe the task and assist the user only if his/her operation may lead to dangerous situations, e.g. too fast or prone to collisions.

The HHI and HRI experiments are performed on independent groups. Within each group, the subjects experimented under both feedback conditions (H and NH) to drive on both tracks consecutively in a single day. The conditions H and NH were presented to the subjects in balanced random order to avoid ordering and learning effects. In particular, the subjects were randomly assigned to two equal-sized sub-groups, where the first sub-group was presented with H, followed by NH, and the second experimented with NH, followed by H. The experimenter was also blinded from the condition, but was aware of the details of the study. Since the experimenter's interface was always programmed to feedback the forces from the task and the user, it was not possible for her to know whether the user is working under H or NH condition. In this sense, the experimenter bias was controlled.

Each experiment consisted of four sessions in a single day: a practice session, followed by a demonstration session and two experimental sessions, which were identical in terms of the task that the subjects were asked to perform. In the practice session, the subjects were asked to complete a single lap on the rectangular track, followed by another lap on the figure 8 track. After the practice session, the subjects were asked to complete a single lap on the figure 8 track, where the assistant provided remote guidance, which acted as the one-shot shared control demonstration. At the end of the demonstration session, the users were asked to fill in a questionnaire to provide demographics information as well as particulars on past experience in driving, powered wheelchairs, and haptic devices. This brief period was sufficient to train the personalized user model in the HRI experiment, and was used to keep the subjects busy until the training is over without informing them about the background operations.

In the experimental sessions, the subjects were presented with either H or NH, followed respectively by NH and H, depending on their group. In each session, the subjects were asked to drive through the same trajectories to complete eight laps (five on the figure 8 track, and three on the rectangular track) as fast as they can and avoiding collisions with walls and obstacles. The initial position, where the subjects started the experiment was marked on the floor and the subjects were instructed to pass through this spot to complete the laps. A bell sound was played by ARTA each time the subjects complete a lap; and at the end of the first 5 laps they were notified to switch to the second trajectory. At the end of each session, the subjects were given a questionnaire and were asked to comment on their experience. Table 1 summarizes the experimental protocol.

TABLE 1: The experimental protocol

SESSION	TASK
Practice	1 lap x trajectory 2 (rectangular track) 1 lap x trajectory 1 (figure 8 track)
Demonstration	1 lap x trajectory 1 (figure 8 track)
Demographics and personal experience questionnaire (Model is trained in HRI experiment)	
Experimental Session -A-	5 laps x trajectory 1 (figure 8 track) 3 laps x trajectory 2 (rectangular track) User experience questionnaire
Experimental Session -B-	5 laps x trajectory 1 (figure 8 track) 3 laps x trajectory 2 (rectangular track) User experience questionnaire

5.3 Participants

24 able-bodied subjects (9 females and 15 males), aged between 19 and 36 (mean = 26.8, std = 4.2), participated in our study. All subjects were right-handed and used their dominant hands when interfacing with the wheelchair. The subjects interacted with the same remote assistant, who was experienced in driving the wheelchair remotely using the haptic interface. Prior to the experimental sessions, the subjects rated their past driving experience and acquaintance with the wheelchair and haptic technologies on a 4-pt Likert scale (1 = None, 2 = Little, 3 = Some, 4 = A Lot). According to questionnaire responses, most subjects had previous driving experience (median = 3, IQR = 2), but few had previous acquaintance with haptic devices (median = 1, IQR = 1) or powered wheelchairs (median = 1, IQR = 0).

6 Model Evaluation

In order to evaluate the ability of our technique to emulate human assistance, the data collected over the HHI condition is analyzed. For this purpose, a GP model is trained using the observations collected during the demonstration trial; and then used to estimate the assistance signals offered by the human in response to varying user commands and environmental context in the upcoming trials.

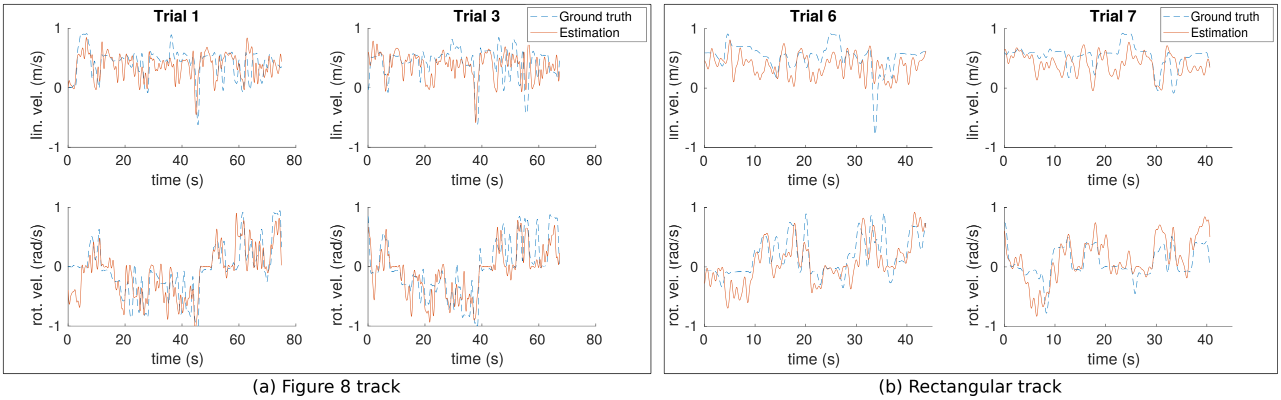

Fig. 6: Assistance command estimations for a representative user over four trials (a) on the figure 8 track (b) on the rectangular track

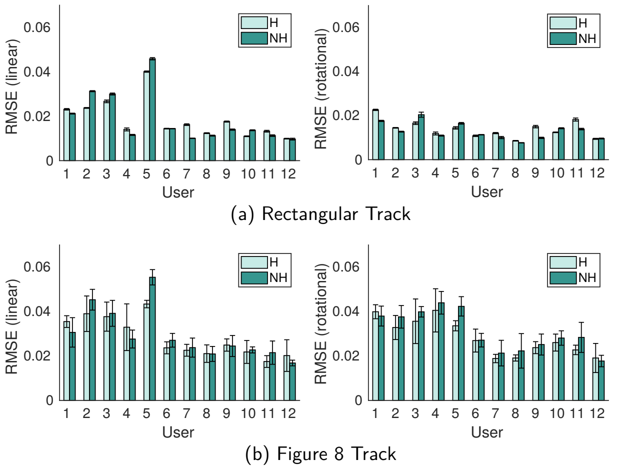

Fig. 7: RMSE scores showing the normalized estimation error over the trials for all users

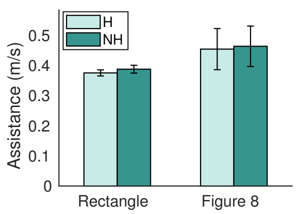

Fig. 8: Mean magnitude of the assistance commands over the trials for all users. Error bars indicate standard deviations.

Fig. 6 plots the ground truth assistance and the estimation signals acquired using the GP model for a representative user over two trials for each of figure 8 and the rectangular tracks. Evidently, model-based estimations follow trends similar to ground truth human assistance in both linear and rotational velocity commands; hence is able to mimic the operation of the human assistant. As seen in Fig. 6(b), the model trained on the figure 8 track is able to generalize well to the rectangular track. In line with this, Fig. 7 presents the root mean squared error (RMSE) values between the assistance estimations and the actual assistance given by the human for linear and rotational velocity commands. The model attains relatively low ($<$ 0.06) errors for command velocities for both trajectories, illustrating good and consistent estimation performance over the course of the experiment to predict human assistance. We note that RMSE values are typically lower for the rectangular track, since the level of assistance for this easier track is lower. As shown in Fig. 8, the mean magnitude of assistance is slightly lower than 0.5 for both tracks. This indicates that the assistant was active for roughly half of the trials.

Fig. 9 plots some trajectories followed by two users. Overlaying density plots mark the locations on the paths, where assistance was offered to the user either by the human assistant or through the estimations achieved by the GP model. The darker red areas indicate regions where higher amounts of assistance was applied. Fig. 9(a) plots the assistance density for the initial demonstration session. Figs. 9(b) and 9(c) plot the paths overlayed with ground truth assistance density as applied by the human and the density of the assistance estimations, respectively under H and NH conditions. Clearly, the model is able to capture the intentions of the human assistant to reproduce his/her operation under trajectories that are different than the demonstration; hence has a generalization capability. These results indicate that the model does not memorize where the assistance shall be applied, but functions as a context and user-aware technique for assistive mobility.

Upon closer inspection, Fig. 9 exemplifies two different scenarios. In the first one, as shown in the upper set of figures, the human assistant applies a not-so-strong assistance policy during the demonstration. On the other hand, (s)he diplays comparably stronger assistance during the experimental trials. In response to the weak assistance demonstration, the robot learns to act in a way that is less active than the actual human assistance. In contrast, the bottom set of figures show a demonstration where the human assistant is dominant. As the figure indicates, in this case, the model predictions tend to be stronger.

7 Assistive Mobility Evaluation

This section provides an evaluation of the model under an actual driving scenario. Learned assistance policy is compared to a baseline condition of human assistance.

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

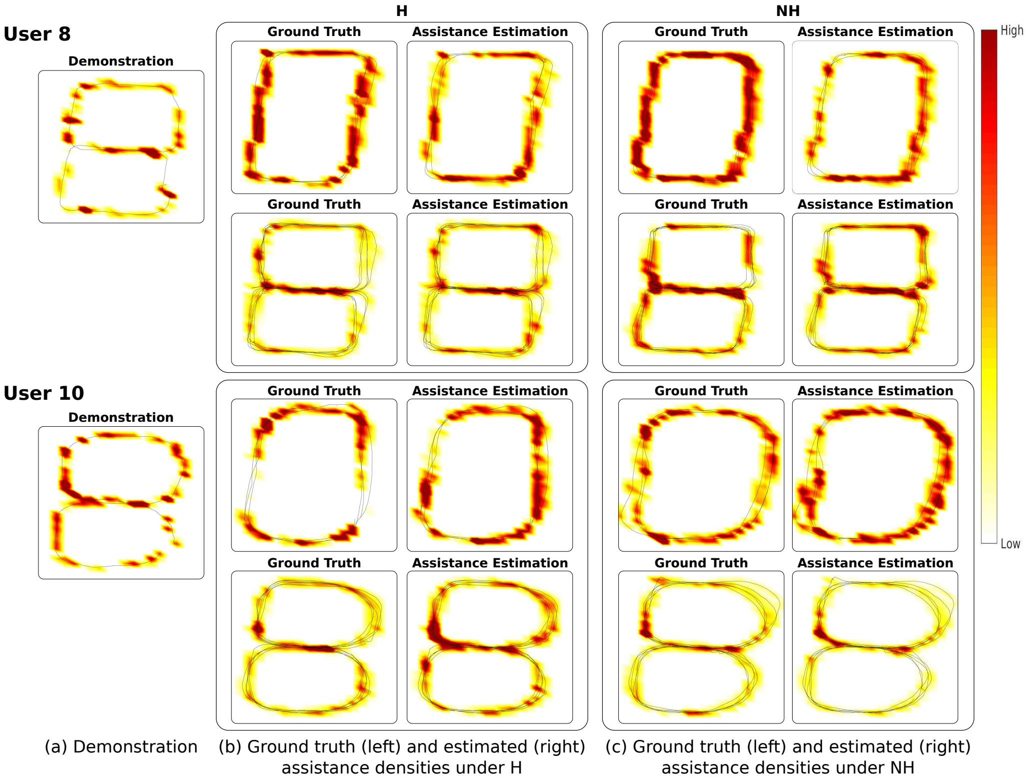

Fig. 9: Example trajectories for two users, with overlaying density plots indicating the location and amount of the offered assistance. Darker areas indicate higher amounts of assistance. a) Assistance density in the demonstration trial. These data were used to train the GP regression model. b) Ground truth (left) and estimated (right) assistance densities under H. c) Ground truth (left) and estimated (right) assistance densities under NH.

The analyses that will be presented in the rest of this section investigate the statistically significant effects of the interaction type (HRI vs. HHI), feedback condition (H vs. NH), and the track (rectangular vs. figure 8) using three-way mixed ANOVA. If no interaction effects are observed for a dependent variable, only the main effects are reported. In case of statistically significant interactions, follow up tests exploring simple effects are carried out.

During operation, timestamped sensor data were collected at 20 Hz to record the AMCL estimation for ARTA's pose in the inertial frame, the forces applied by the agents ($\boldsymbol{F}_{HIP_1}$ and $\boldsymbol{F}_{HIP_2}$), and the joystick commands issued on the wheelchair by the user and the assistant. In the evaluations, only the data collected within the experimental sessions are used.

7.1 Task performance

Fig. 10 shows mean lap completion times and the mean proximity to obstacles for each condition and standard errors of the means. ANOVA results indicate a statistically significant effect of the interaction type ($F(1,376) = 11.631$, $p < .005$) on lap completion time. Particularly, in the HRI experiment, subjects spend more time to complete the task than they do in the HHI experiment. For both interaction types, HHI and HRI, the lap completion times are consistent across feedback conditions H and NH ($F(1,376) = 2.760$, $p = .111$) and they are significantly lower for the rectangular track ($F(1,376) = 80.031$, $p < .001$).

Regarding the mean driving proximity to obstacles, ANOVA results indicate no main effect of interaction type ($F(1,376) = .677$, $p = .419$) or the feedback condition ($F(1,376) = .394$, $p = .537$), but a significant effect of the track ($F(1,376) = 33.928$, $p < .001$), with the proximity to obstacles getting lower (i.e. less safer) when driving on the Figure 8 track.

These results indicate that the addition of haptics does not significantly affect the performance of the users. However, the integration of robotic shared control slows down task execution.

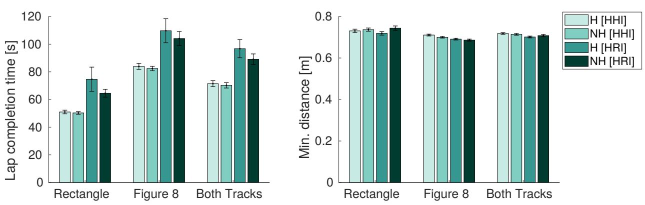

Fig. 10: Mean lap completion times (left) and the proximity to obstacles (right), and standard errors of the means

7.2 Smoothness of motion

The smoothness of motion is quantified as a function of jerk, i.e. the time derivative of acceleration. The jerk cost equals to the time normalized integral of the square of jerk in wheelchair position [46], and is computed separately to quantify the smoothness of the joystick commands applied by the user (JU) and the assistant (JA):

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

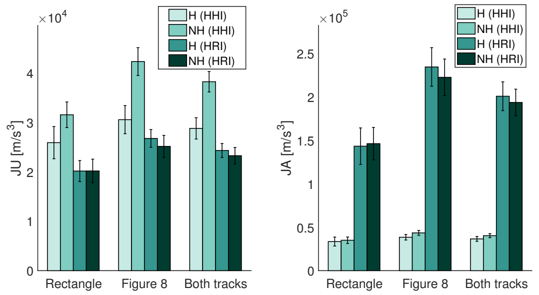

Fig. 11: Mean jerk costs in user's (left) and assistant's (right) joystick movements and standard errors of the means

$$JU = \frac{1}{T}\int_0^T \|\dddot{\boldsymbol{x}}_{HIP1}\|^2 dt \tag{9}$$

$$JA = \frac{1}{T}\int_0^T \|\dddot{\boldsymbol{x}}_{HIP2}\|^2 dt \quad , \tag{10}$$

where $\boldsymbol{x}_{HIP1}$ and $\boldsymbol{x}_{HIP2}$ respectively denote the HIP position of the user and the assistant. In order to reduce differentiation errors due to noise, in all jerk computations, position data are filtered with a 4$^{\text{th}}$ order lowpass Butterworth filter with 5 Hz cut-off frequency. A two-way filter is used to avoid phase shifts.

The jerk costs for the agents' joystick movements, and the standard errors of the means are shown in Fig. 11. A statistically significant main effect of the interaction type ($F(1, 376) = 4.317$, $p < .05$) for JU indicates that the user's hand motion is smoother under HRI than it is under HHI. Also, a statistically significant main effect of track type ($F(1, 376) = 19.832$, $p < .001$) illustrates that the joystick motion characteristics of the subjects are jerkier under the harder figure 8 track for both interaction types.

This result is accompanied by a significant increase in the mean jerk cost in the assistant's hand movement under HRI ($F(1, 376) = 13.690$, $p < .005$). The jerkiness of the assistant's hand motion is consistent accross H and NH ($F(1, 376) = .012$, $p = .913$), and is significantly higher in the figure 8 track only under HRI ($F(1, 376) = 26.977$, $p < .001$).

These indicate that haptic feedback does not affect the smoothness of the subjects' joystick motions. However, the users tend to perform more smoothly when interacting with robotic assistance, the signals of which are significantly jerkier than the joystick trajectory of a human assistant.

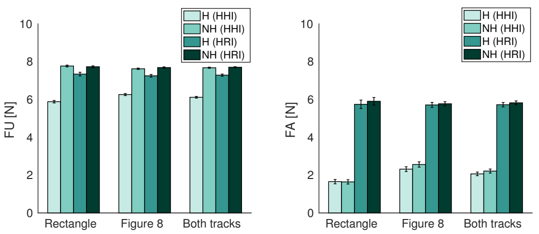

Fig. 12: Mean force contribution of the user (left) and the assistant (right) and standard errors of the means

7.3 Individual forces applied by the agents

Mean magnitudes of individual interaction forces $\boldsymbol{F}_{HIP_1}$ and $\boldsymbol{F}_{HIP_2}$, are reported to quantify the individual force contributions of the user (FU) and the assistant (FA) :

$$FU = \frac{1}{T}\int_0^T \|\boldsymbol{F}_{HIP_1}\| dt \tag{11}$$

$$FA = \frac{1}{T}\int_0^T \|\boldsymbol{F}_{HIP_2}\| dt \quad . \tag{12}$$

Fig. 12 illustrates the mean magnitude of the forces applied by the user and the assistant. The interaction type has a significant simple effect on the amount of individual forces applied by the user only under H ($F(1, 376) = 34.368, p < .001$), where the user applies lower forces when working with a human assistant in the presence of haptic feedback from the partner. The forces applied by the user in HRI and HHI settings stay in similar higher ranges under NH. Additionally, there is a significant main effect of the feedback condition on the net force applied by the user ($F(1, 376) = 132.539, p < .001$), where higher forces are applied in the lack of haptic feedback from the assistant, i.e. under NH. The forces are increased for the figure 8 track when compared to the rectangular track under H, only for HHI ($F(1, 376) = 19.975, p < .001$).

On the other hand, the forces applied by the assistant are significantly increased in HRI ($F(1, 376) = 104.697, p < .001$). The feedback condition does not affect the forces applied by the assistant ($F(1, 376) = 3.728, p = .066$), which is indicative of consistent assistance behavior for both HRI and HHI settings. There is a significant effect of track for the assistant's force contribution only under HHI ($F(1, 376) = 28.683, p < .001$), whereas assistance forces stay consistent across tracks under HRI ($F(1, 376) = .303, p = .582$).

7.4 Subjective Evaluation

At the end of each session, the subjects were asked to indicate their level of agreement or disagreement on a 7-point Likert scale (1 = Totally Disagree, 7 = Totally Agree) for a series of statements. The questionnaire consisted of questions, six of which are taken from NASA-TLX task load index [47], as well as those developed specifically for the purpose of this experiment:

- *Mental Demand:* The task required a large amount of mental and perceptual activity (e.g. thinking, deciding, calculating, remembering, looking, searching, etc.)
- *Physical Demand:* The task required a large amount of physical activity (e.g. pulling, pushing, turning, controlling, activating, etc.)
- *Temporal Demand:* I needed to be quick to perform the task.
- *Performance:* I was successful in accomplishing the goals of the task set by the experimenter (or myself).
- *Effort:* I had to work hard (mentally and physically) to accomplish the task.
- *Frustration Level:* I felt irritated / stressed / annoyed during the task.

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

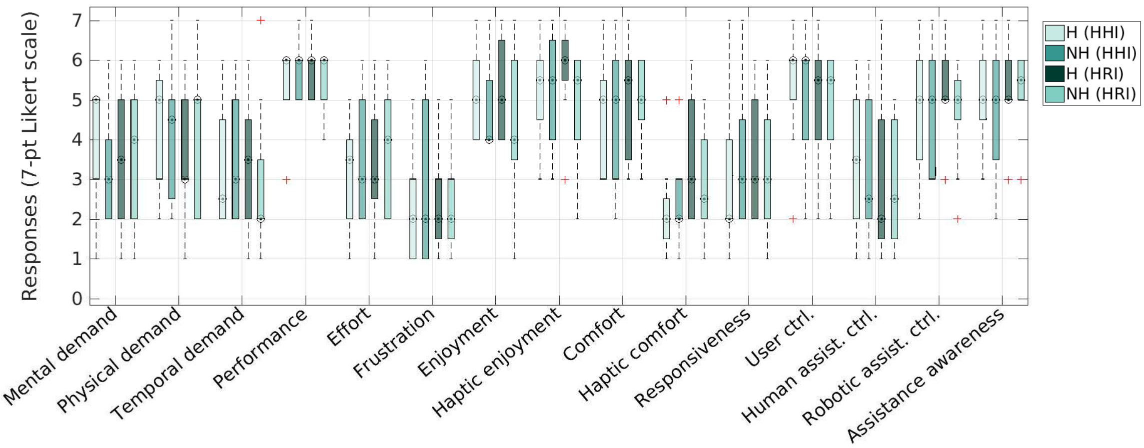

Fig. 13: Boxplots showing the median and interquartile ranges of the subjective scores for each variable (1 = Totally disagree, 7 = Totally agree)

- *Enjoyment:* I enjoyed driving the wheelchair.
- *Haptic enjoyment:* I enjoyed driving the wheelchair particularly because I was using the haptic joystick.
- *Comfort:* I felt comfortable driving the wheelchair.
- *Haptic comfort:* Using the haptic joystick improved my comfort level during the experiment.
- *Responsiveness:* Wheelchair responded to my actions.
- *Assistance awareness:* I got some external assistance during operation.
- *Control source:* Three questions investigated the users' perception of who is controlling the operation:
  1) *User control:* The robotic wheelchair moved under my command.
  2) *Human assistant control:* The robotic wheelchair was operated remotely by the experimenter.
  3) *Robotic assistant control:* The robotic wheelchair (i.e. not the experimenter) was interfering with my operation.

Fig. 13 plots the mean values of the subjects' responses to the statements in the questionnaire. The differences between the means for H and NH as well as HHI and HRI groups are investigated through 2-way mixed ANOVA. No statistically significant differences are observed for any of the subjective measures between the feedback conditions H and NH or between HHI and HRI. This indicates that the robotic assistance was not perceived very differently from actual human assistance, and haptic cues from partners did not cause the subjects to feel an increased level of work load, nor did they affect the subjects' comfort or enjoyment level. The subjects did not report any particular positive tendency toward the haptic sensations that made the task more comfortable or more enjoyable. In general, the subjects thought that the wheelchair was responding to their actions, but were aware that the wheelchair did not move completely under their command. Interestingly, the subjects were unable to distinguish whether the external commands were exerted by a remote human being, or were due to some programmed assistance function of the robotic wheelchair. However, in all conditions, the subjects tended to think that they interacted with a robot, not a human.

8 Discussion and Conclusions

In this study, we proposed a technique to learn shared control policies from human-human interactions. The capability of this machine-learned model to generate and generalize assistive policies for different trajectories is illustrated by comparison with human data. The utility of the proposed technique is further investigated through a user study, where robotic shared control policies (HRI) are compared with actual human assistance (HHI) for guiding a wheelchair user. As a side goal, we investigated the utility of haptic feedback during shared control with an assistant. In our setup, the driver of the wheelchair was not informed about whether (s)he was physically coupled with another human or a robot during the task. This was done on purpose to bypass the cognitive adaptation process and focus clearly on motor adaptation [7].

As a result of the user study, we observed that the trajectories under HRI and HHI scenarios were similar in terms of the wheelchair's average distance to obstacles. Lap completion time was significantly increased when the user operated with robotic assistance; yet the addition of haptic feedback on the shared control setup did not affect task performance. The forces applied by the human under HRI stayed in consistent levels with those observed under HHI, particularly when no haptic feedback is provided. This indicates that the effort requirement of the users were comparable between HRI and HHI settings. However, with the inclusion of haptic feedback, the effort requirements of the users dropped in HHI. Even though there is no direct evidence that humans deliberately use forces to realize interpersonal communication during physical joint action, this observation may imply that the motor responses of the users are affected by the existence of haptic sensations from their partners, which reduces the effort requirements of the human. However, why this effect is only visible for human-human interaction and not apparent in a human-robot setting arises an interesting research question that requires further exploration.

On the other hand, the behaviors of the robotic and human assistants, as defined by the force contribution and the jerkiness of the assistance commands, differed significantly. We observed that these variables were not affected by the presence or absence of force feedback for HRI, yet we noticed some changes under HHI across feedback conditions. Interestingly, these differences did not manifest themselves in subjective task load scores. This is particularly appealing, as it may mean that from the perspective of the user, robotic assistance can replace a human assistant without introducing extra task load, or negatively affecting task comfort. However, as mentioned earlier, an investigation of mutual adaptation mechanism under both human-robot and human-human settings shall be conducted before reaching any conclusions.

In the experiment, the same human, who was experienced with the wheelchair in both collocated and teleoperated settings, acted as the assistant. She attempted to observe the task and assist the user if his/her operation felt dangerous (e.g. too fast or prone to collisions). Despite these qualifications, the assistant's consistency during the trials cannot be guaranteed. We designed this experiment with the assumption that discrepancies in human behavior would not affect the tested hypotheses. In planning future studies, interaction scenarios with different assistants displaying different assistance behaviors will be devised.

This paper has presented observations on general tendencies toward a robotic assistant in physical shared control. However, personal differences can also be highly effective in determining the human responses to the use of assistive technologies. Morere et al. have presented a user study with disabled children on the use of haptics for wheelchair control [40]. Their findings indicate personal differences in receiving the benefits offered by the haptic technology. In line with this, in [48], we had discussed that any assistive system needs to be trained to accommodate personal characteristics and should exploit individual preferences of humans it interacts with. Our findings indicated that haptic feedback did not play a very important role in this shared control setup, especially when working with a robot. As a future direction, we will investigate personal differences in how individuals respond to machine learned assistive policies, and in particular, process haptic information when used as a communication medium.

Acknowledgements

This work was supported by EU grant H2020-PHC-643783 "PAL": Personal Assistant for Healthy Lifestyle, and TUBITAK 2219 International Postdoctoral Grant 1649B030800061. The experiment was conducted at the Personal Robotics Laboratory of Imperial College London. The authors thank Joshua Elsdon, Dr Miguel Sarabia, and Phil Jones for their contributions in building ARTA; and Dimitri Ognibene for his helpful comments.

References

[1] W.-Y. Chen, Y. Jang, J.-D. Wang, W.-N. Huang, C.-C. Chang, H.-F. Mao, and Y.-H. Wang, "Wheelchair-related accidents: relationship with wheelchair-using behavior in active community wheelchair users," *Archives of physical medicine and rehabilitation*, vol. 92, no. 6, pp. 892–898, 2011.

[2] C. Madan, A. Kucukyilmaz, M. Sezgin, and C. Basdogan, "Recognition of haptic interaction patterns in dyadic joint object manipulation," *Haptics, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2014. [Online]. Available: http://dx.doi.org/10.1109/TOH.2014.2384049

[3] A. Melendez-Calderon, V. Komisar, G. Ganesh, and E. Burdet, "Classification of strategies for disturbance attenuation in human-human collaborative tasks," in *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*. IEEE, 2011, pp. 2364–2367.

[4] N. Stefanov, A. Peer, and M. Buss, "Role determination in human-human interaction," in *WHC'09: IEEE World Haptics Conference*, Salt Lake City, USA, 2009, pp. 51–56.

[5] R. Groten, D. Feth, R. Klatzky, and A. Peer, "The role of haptic feedback for the integration of intentions in shared task execution," *IEEE Transactions on Haptics*, vol. 6, no. 1, pp. 94–105, 2013.

[6] N. Jarrassé, V. Sanguineti, and E. Burdet, "Slaves no longer: review on role assignment for human–robot joint motor action," *Adaptive Behavior*, vol. 22, no. 1, pp. 70–82, 2014.

[7] G. Ganesh, A. Takagi, R. Osu, T. Yoshioka, M. Kawato, and E. Burdet, "Two is better than one: Physical interactions improve motor performance in humans," *Scientific reports*, vol. 4, 2014.

[8] A. Takagi, G. Ganesh, T. Yoshioka, M. Kawato, and E. Burdet, "Physically interacting individuals estimate the partner's goal to enhance their movements," *Nature Human Behaviour*, vol. 1, 2017.

[9] L. Rozo, P. Jiménez, and C. Torras, "A robot learning from demonstration framework to perform force-based manipulation tasks," *Intelligent service robotics*, vol. 6, no. 1, pp. 33–51, 2013.

[10] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer, 2008, pp. 1371–1394.

[11] D. Korkinof and Y. Demiris, "Online quantum mixture regression for trajectory learning by demonstration," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 3222–3229.

[12] Y. Wu, Y. Su, and Y. Demiris, "A morphable template framework for robot learning by demonstration: Integrating one-shot and incremental learning approaches," *Robotics and Autonomous Systems*, vol. 62, no. 10, pp. 1517–1530, 2014.

[13] R. Gelin, J. Detriche, J. P. Lambert, and P. Malblanc, "The sprint of coach," in *Systems, Man and Cybernetics, 1993. 'Systems Engineering in the Service of Humans', Conference Proceedings., International Conference on*, vol. 3, 1993, pp. 547–552.

[14] T. Gomi and A. Griffith, "Developing intelligent wheelchairs for the handicapped." in *Assistive Technology and Artificial Intelligence*, ser. LNCS, V. O. Mittal, H. A. Yanco, J. M. Aronis, and R. C. Simpson, Eds., vol. 1458. Springer, 1998, pp. 150–178.

[15] R. Simpson, E. LoPresti, S. Hayashi, S. Guo, D. Ding, W. Ammer, V. Sharma, and R. Cooper, "A prototype power assist wheelchair that provides for obstacle detection and avoidance for those with visual impairments," *Journal of Neuroengineering and Rehabilitation*, vol. 2, no. 30, october 2005. [Online]. Available: http://dx.doi.org/10.1186/1743-0003-2-30

[16] S. P. Levine, D. A. Bell, L. A. Jaros, R. C. Simpson, Y. Koren, and J. Borenstein, "The navchair assistive wheelchair navigation system," *IEEE Trans. on Rehabilitation Engineering*, vol. 7, pp. 443–451, 1999.

[17] H. A. Yanco, "Wheelesley: A robotic wheelchair system: Indoor navigation and user interface," in *Assistive Technology and Artificial Intelligence, Applications in Robotics, User Interfaces and Natural Language Processing*. London, UK: Springer-Verlag, 1998, pp. 256–268.

[18] E. S. Boy, C. L. Teo, and E. Burdet, "Collaborative wheelchair assistant," in *IEEE/RSJ International Conference on Intelligent Robots and System*, 2002, pp. 1511–1516.

[19] T. Carlson and Y. Demiris, "Human-wheelchair collaboration through prediction of intention and adaptive assistance." in *ICRA*. IEEE, 2008, pp. 3926–3931.

[20] ——, "Collaborative control for a robotic wheelchair: Evaluation of performance, attention, and workload." *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 42, no. 3, pp. 876–888, 2012.

[21] E. Demeester, A. Hüntemann, D. Vanhooydonck, G. Vanacker, A. Degeest, H. V. Brussel, and M. Nuttin, "Bayesian estimation of wheelchair driver intents: Modeling intents as geometric paths tracked by the driver," in *IEEE/RSJ International Conference on*

Preprint version • Citation information: DOI: 10.1109/TOH.2018.2804911, IEEE Transactions on Haptics

Intelligent Robots and Systems, IROS 2006, October 2006, pp. 5775–5780.

[22] E. Demeester, A. Hüntemann, E. Vander Poorten, and J. De Schutter, "ML, MAP and greedy POMDP shared control: comparison of wheelchair navigation assistance for switch interfaces," in *International Symposium on Robotics*, 2012, pp. 1106–1111.

[23] A. Hüntemann, E. Demeester, E. Vander Poorten, H. Van Brussel, and J. De Schutter, "Probabilistic approach to recognize local navigation plans by fusing past driving information with a personalized user model," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 4376–4383.

[24] J. d. R. Millán, R. Rupp, G. R. Müller-Putz, R. Murray-Smith, C. Giugliemma, M. Tangermann, C. Vidaurre, F. Cincotti, A. Kübler, R. Leeb, *et al.*, "Combining brain–computer interfaces and assistive technologies: state-of-the-art and challenges," *Frontiers in Neuroscience*, vol. 4, 2010.

[25] A. Kucukyilmaz, T. M. Sezgin, and C. Basdogan, "Intention recognition for dynamic role exchange in haptic collaboration," *IEEE Transactions on Haptics*, vol. 6, no. 1, pp. 58–68, First 2013. [Online]. Available: http://dx.doi.org/10.1109/TOH.2012.21

[26] A. Moertl, M. Lawitzky, A. Kucukyilmaz, T. M. Sezgin, C. Basdogan, and S. Hirche, "The role of roles: Physical cooperation between humans and robots," *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1656–1674, 2012. [Online]. Available: http://ijr.sagepub.com/content/31/13/1656.abstract

[27] D. A. Abbink, M. Mulder, and E. R. Boer, "Haptic shared control: smoothly shifting control authority?" *Cognition, Technology & Work*, vol. 14, no. 1, pp. 19–28, 2012.

[28] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.

[29] M. Fernandez-Carmona, B. Fernandez-Espejo, J. Peula, C. Urdiales, and F. Sandoval, "Efficiency based collaborative control modulated by biometrics for wheelchair assisted navigation," in *2009 IEEE International Conference on Rehabilitation Robotics*. IEEE, 2009, pp. 737–742.

[30] Q. Li, W. Chen, and J. Wang, "Dynamic shared control for human-wheelchair cooperation," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 4278–4283.

[31] A. Goil, M. Derry, and B. D. Argall, "Using machine learning to blend human and robot controls for assisted wheelchair navigation," in *Rehabilitation Robotics (ICORR), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1–6.

[32] H. Soh and Y. Demiris, "When and how to help: An iterative probabilistic model for learning assistance by demonstration," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2013*, Nov 2013, pp. 3230–3236.

[33] ——, "Spatio-temporal learning with the online finite and infinite echo-state gaussian processes," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 26, no. 3, pp. 522–536, March 2015.

[34] A. Fattouh, M. Sahnoun, and G. Bourhis, "Force feedback joystick control of a powered wheelchair: preliminary study," in *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, vol. 3, 2004, pp. 2640–2645. [Online]. Available: http://dx.doi.org/10.1109/ICSMC.2004.1400729

[35] M. H. Abdelkader, G. Bourhis, and B. Cherki, "Haptic feedback and motor disability: a preliminary study," in *International Conference on Human-Machine Interaction, HuMaN07*, march 2007.

[36] M. A. H. Abdelkader, G. Bourhis, and B. Cherki, "Haptic feedback control of a smart wheelchair," *Applied Bionics and Biomechanics*, vol. 9, no. 2, pp. 181–182, 2012. [Online]. Available: http://dx.doi.org/10.3233/ABB-2012-0067

[37] G. Bourhis and M. Sahnoun, "Assisted control mode for a smart wheelchair," in *Rehabilitation Robotics, 2007*, 2007, pp. 158–163. [Online]. Available: http://dx.doi.org/10.1109/ICORR.2007.4428422

[38] X. Chen, C. Ragonesi, J. Galloway, and S. Agrawal, "Training toddlers seated on mobile robots to drive indoors amidst obstacles," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 19, no. 3, pp. 271–279, 2011. [Online]. Available: http://dx.doi.org/10.1109/TNSRE.2011.2114370

[39] L. Marchal-Crespo, J. Furumasu, and D. Reinkensmeyer, "A robotic wheelchair trainer: design overview and a feasibility study," *Journal of NeuroEngineering and Rehabilitation*, vol. 7, no. 1, pp. 1–12, 2010. [Online]. Available: http://dx.doi.org/10.1186/1743-0003-7-40

[40] Y. Morère, M. H. Abdelkader, K. Cosnuau, G. Guilmois, and G. Bourhis, "Haptic control for powered wheelchair driving assistance," *IRBM*, vol. 36, no. 5, pp. 293–304, 2015.

[41] M. Sarabia and Y. Demiris, "A humanoid robot companion for wheelchair users," in *Social Robotics*, ser. Lecture Notes in Computer Science, G. Herrmann, M. Pearson, A. Lenz, P. Bremner, A. Spiers, and U. Leonards, Eds., vol. 8239. Springer International Publishing, 2013, pp. 432–441.

[42] H. Soh and Y. Demiris, "Towards early mobility independence: An intelligent paediatric wheelchair with case studies," in *IEEE/RSJ IROS2012 Workshop on Navigation and Manipulation Assistance for Robotic Wheelchairs*, October 2012.

[43] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, ser. Adaptative computation and machine learning series. The MIT Press, 2006.

[44] S. Park and S. Choi, "Hierarchical gaussian process regression." in *ACML*, 2010, pp. 95–110.

[45] C. M. Bishop, *Pattern recognition and machine learning*, 4th ed. springer New York, 2006.

[46] T. Flash and N. Hogan, "The coordination of arm movements: an experimentally confirmed mathematical model," *The journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.

[47] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," *Advances in psychology*, vol. 52, pp. 139–183, 1988.

[48] A. Kucukyilmaz and Y. Demiris, "One-shot assistance estimation from expert demonstrations for a shared control wheelchair system," in *Robot and Human Interactive Communication (RO-MAN), 2015 24th IEEE International Symposium on*. IEEE, 2015, pp. 438–443. [Online]. Available: https://doi.org/10.1109/ROMAN.2015.7333600

Ayse Kucukyilmaz is a Lecturer in Computer Science at University of Lincoln, U.K and a member of the Lincoln Centre for Autonomous Systems (L-CAS). Her research interests include haptics, physical HRI, assistive robotics, and machine learning. Between 2013 and 2015, she was a research associate at the Personal Robotics Laboratory of Imperial College London, working in the EU FP7 project, ALIZ-E. Later, she worked as an assistant professor at the Department of Computer Engineering at Yeditepe University between 2015 and 2017. She received her Ph.D. in Computational Sciences and Engineering from Koc University, in 2013, and obtained her B.Sc. and M.Sc. degrees in Computer Engineering from Bilkent University, Turkey, respectively in 2004 and 2007. Dr. Kucukyilmaz was the recipient of the Academic Excellence Award at her graduation from Koc University in 2013, and has received fellowships from TUBITAK, Turkey and Marie Sklodowska-Curie Actions.

Yiannis Demiris (SM03) received the B.Sc. (Hons.) and Ph.D. degrees from the Department of Artificial Intelligence, University of Edinburgh, Edinburgh, U.K. He is a Professor with the Department of Electrical and Electronic Engineering at Imperial College London, where he heads the Personal Robotics Laboratory. His research interests include human robot interaction, machine learning, user modeling, and assistive robotics; he has published more than 150 journal and peer reviewed conference papers on these topics. Professor Demiris was the Chair of the IEEE International Conference on Development and Learning in 2007 and the Program Chair of the ACM/IEEE International Conference on Human Robot Interaction in 2008. He was a recipient of the Rectors Award for Teaching Excellence in 2012 and the FoE Award for Excellence in Engineering Education in 2012. He is a senior member of the IEEE, and a Fellow of the IET, BCS and the Royal Statistical Society.