2015

Distributed and Adaptive Shared Control Systems: Methodology for the Replication of Experiments.

Stoelen, MF

<http://hdl.handle.net/10026.1/4550>

10.1109/MRA.2015.2460911

IEEE Robotics Autom. Mag.

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

GALILEO — IMAGE LICENSED BY INGRAM PUBLISHING.
LAMP — ©ISTOCKPHOTO.COM/JONNY JINI

Distributed and Adaptive Shared Control Systems

Methodology for the Replication of Experiments

By Martin F. Stoelen, Virginia Fernández de Tejada,
Alberto Jardón Huete, Carlos Balaguer,
and Fabio Paolo Bonsignorio

Much work in robotics aimed at real-world applications falls in the large segment between teleoperated and fully autonomous systems. Such systems are characterized by the close coupling between the human operator and the robot, in principle, allowing the agents to share their particular sensing, adaptation, and decision-making capabilities. Replicable experiments can advance the state of the art of such systems but pose practical and epistemological challenges. For example, the trajectory of the system is governed by the adaptation both in the human and the robot agent. What do we need besides (or instead of) data sets for such a system? The degree

ALL-FREE-DOWNLOAD.COM

of similarity between comparable experiments and the exact meaning of replication need to be clarified. Here, we explore replication of a distributed and adaptive shared control for an assistive robot manipulator. We attempt a methodological approach for reporting two virtual human experiments on the system: modeling the complete human-robot binomial, deriving closed-loop performance metrics from the models, and openly publishing the results and experiment implementations.

Replication and Human-Robot Systems

We may think of theoretical/concept papers, proof-of-concept papers, and experimental papers as steps in a research idea life cycle. We believe that more papers of the experimental kind would greatly help the research activities in robotics and the industrial exploitation of the results. This is

consistent with the good experimental methodology (GEM) guidelines [1] promoted by the Euron GEM special interest group (SIG). Such experimental papers should likely contain (at least):

- *Description*: a journal paper with text, figures, and multimedia, according to [1] (or similar).
- *Data sets*: complete set(s) of data for the experiment(s).
- *Code*: complete code identifiers and/or downloadable code (executables may be enough).
- *Hardware*: hardware description or identifier (if it is identifiable).

For work on closely coupled human–robot systems, the experimental papers should likely include the four above-mentioned parts, in addition to the following part, which is specific to the experiments that include human participants:

- *Method*: a complete description of experiments used, e.g., the instructions given to the participants.

This is in line with the recent increase in emphasis on experimental methods in robotics [2]. There are yearly workshops on benchmarking at the major robotics conferences and several open data repositories (see [3]). For assistive technologies in particular, consistent methodologies for benchmarking are essential for comparing different solutions and can help

speed the process of getting them out to the end user. Here, we focus on the replication of a distributed and adaptive shared control approach for an assistive robot manipulator. Two case studies are presented, based on our previous work [4]–[6]. Our approach to replication involves: 1) information-theoretic models of the complete system with the user integrated in the sensorimotor loop, 2) quantitative metrics based on information theory and the speed-accuracy tradeoff of human movements, and 3) the open publication of the complete experiments according to the Euron GEM SIG guidelines. The contributions of this article with respect to previous publications are as follows:

- applies information metrics to the results in [5]
- applies predictive information (PI) for the first time to [4]
- an open repository for sharing experimental papers with human–robot systems has been set up at <http://throughput.sourceforge.net>. The two case studies used here have been posted as examples.

We hope this article can provide the first steps toward replicable experiments on this type of closely coupled human–robot system.

Assistive Manipulators and Shared Control

Assistive robot manipulators aim to increase the independence of their users in performing physical activities of daily living (ADL). Several types of assistive manipulators exist, including static feeders, mobile manipulators, and wheelchair-mounted types. See Figure 1 for examples of the latter. The work described here will focus mainly on ASIBOT, a five degrees of freedom (5DoF) manipulator developed at Universidad Carlos III de Madrid (UC3M) [Figure 1(c)]. Assistive manipulators typically operate in close proximity to the disabled or elderly user. They also aim to perform tasks in real-world, partially structured environments such as a user's home. Most assistive manipulators are, therefore, directly controlled by the user through, for example, a joystick in spite of the high user mental and physical workload required. One potential approach for improving the usability of assistive manipulators is to enable both the user and the robot to use their own sensing, control, and planning capabilities in a cooperative way. This is, here, denoted as shared control.

Shared control has been explored extensively in the context of intelligent wheelchairs, typically by assisting the user through distance sensing and cameras mounted on the wheelchair. The user maintains the locus of control but with assistance on avoiding collisions or performing specific tasks like moving through a doorway. Such systems have been shown to reduce both collisions and cognitive workload [7], and it seems feasible to adapt the assistance automatically to each wheelchair user [8]. Some approaches attempt a higher-level interaction through speech or two-dimensional interfaces combined with some robot autonomy on specific tasks. This is also denoted as shared autonomy [9]. Here, we explore a complementary low-level collision-limitation behavior, dubbed *shared distributed adaptive control (SDAC)*. We consider a tight

Figure 1. Examples of assistive robot manipulators mounted to wheelchairs, also known as wheelchair-mounted robot arms (WMRA). The robots are intended for disabled users in wheelchairs with a low amount of mobility in the upper body. An assistive manipulator can give these users more independence with daily tasks such as preparing a meal. (a) The exact Dynamics AMOR (photo courtesy of Exact Dynamics), (b) the Kinova JACO are current commercial examples (photo courtesy of Kinova), and (c) the UC3M ASIBOT is a research prototype.

Figure 2. A conceptual representation of SDAC, with the current implementation on the AMOR assistive robot manipulator at UC3M in Spain. This is an adaptive approach to shared control, i.e., it attempts to adapt the use of the robot sensing and decision making to each user and task set. The main idea is to help the user limit collisions with the environment through a set of m proximity sensors distributed over the arm. The adaptation occurs in real time by associating proximity sensors with specific collisions experienced (a set of n collision sensors). The approach thus uses distributed sensing but also allows the computations needed for learning and control to be distributed over hardware on the arm.

integration of user and robot with interaction through direct control interfaces as well as through the environment. The behavior is intended both for aiding the disabled user [4], [5] and for aiding the remote teleoperation by a family member or caregiver [6].

Shared Distributed Adaptive Control

The original idea behind SDAC was to let the robot help enforce the accuracy requirements of a task, indirectly allowing the operator to move with greater speed. This collision-limitation idea was merged with a distributed approach to sensing, learning, and control inspired by DAC [10]. This is based on the gradual association and activation of reflexive actions with stimuli that predict future dangerous situations, such as collisions. DAC is, as such, related to classical conditioning in psychology and constitutes: 1) a predefined value system, expressed in combinations of unconditioned stimuli and unconditioned reflexes, and 2) a mechanism for associating conditioned stimuli (CS) representations to unconditioned stimuli representations. The SDAC applies DAC to an assistive robot manipulator, where the stimuli come from distributed tactile and proximity sensing and where the reflexes are mechanisms for slowing the robot down before a collision (Figure 2).

The operator of the robot provides Cartesian velocities $\vec{v}_{user}$, which are resolved into Cartesian velocities for each link $\vec{v}_{link}$. The actual Cartesian velocity followed by the robot, $\vec{v}_{robot}$, is the commanded velocity reduced by a factor r_{max} , the maximum proximity ratio, (A) in Figure 2. The local control component, the equivalent of a reflex, is represented in each collision sensor i by the local proximity ratio r_i , (B) in Figure 2. The proximity ratio is high when: 1) there is a high local forward velocity normal to the collision sensor (high u_i) and 2) there is a high activity in the collision neuron c_i . This proximity ratio is initially only driven by the

collision sensor (with signal $coll_i$), the unconditioned stimuli in classical conditioning, that is hardwired to the local collision neuron (c_i), (C) in Figure 2. Therefore, before beginning, learning the proximity ratio is high only when there is a collision and when the robot locally moves in a direction that maintains or worsens this collision.

The CS are the inputs from a set of proximity sensors on the same link. When there is a collision, the weight of a synapse from a linear proximity sensor neuron p_j that is receiving a high input at that time will increase in strength. The change $\Delta w_{i,j}$ is driven by a Hebbian associative learning rule running in real time, (D) in Figure 2. That is, the synapse weight $w_{i,j}$ is strengthened when there is coactivation in collision sensor i and proximity sensor j . The increase in synapse weights means the proximity sensors will gradually start driving the collision neuron when close to obstacles, which activates the collision-limitation behavior. The learning rate is γ . Thus, the robot begins slowing down before repeating previously experienced collisions. This is then the CR of the system. A simple forgetting term is added, which reduces each weight continually in proportion to the current strength $w_{i,j}$ (with discount rate ϵ). The operator can, therefore, reduce the assistance received by moving slower and with more care. Note also that by providing more assistance to the operator, there will typically be fewer collisions. This means that the weights have a tendency to stabilize for a given operator and situation.

In this article, two experiments with simulated robots are used as case studies. However, distributed proximity sensing is well within the scope of current technology, as shown in Figure 2. In fact, distributed robot arm proximity sensing with a tenfold density was developed and tested over 25 years ago [11].

Modeling the System

Understanding the potential and the limits of a closely coupled human-robot system like SDAC is not trivial. By sharing

Figure 3. The directed acyclic graph representing the human-robot system for assistive manipulators with SDAC. Circles: components of the system as random variables. Arrows: probabilistic dependence. Subscripts: time-steps. The meaning of the random variables used is as follows: U —user agent; H —human intent; D —input device; Z —noise, typically representing a simplified physical disability; X —overall system state, including environment, physical robot, and physical user; R —robot agent; A —actuation by robot; C —collision sensor; and P —proximity sensor.

control, both human and robot agents are also able to affect the outcome of the experimental evaluation. Modeling the complete system can give us a better understanding of how to compare such systems and what to measure. A model of a complete shared control system will typically have to include the user, an adaptive and self-motivated agent, inside the control loop. For such systems, we believe it is fruitful to investigate the information that is transmitted in the system during usage and its usefulness in accomplishing the user's goals in an effective and safe way. Interestingly, Touchette and Lloyd [12] proposed an information-theoretic approach to the study of control systems. We have applied this approach to represent and evaluate human-robot systems [4], which is extended here. An entropy approach to the problems of the control of stochastic nonlinear systems was previously proposed (with different mathematical formalisms) by, among others, Saridis [13]. The advantage of the methods proposed in [12] with respect to those proposed by Saridis is that, while sharing the conceptual framework, they lead to a more direct quantification.

Now armed with information theory [14], what can we say about a complex human-robot system like SDAC? Figure 3 shows the complete human-robot system assumed here as a directed acyclic graph. For each time step t , the system has a random variable representing the current state X_t , with values $x_t \in \mathcal{X}_t$ and a future state X_{t+1} . The state of the system is assumed to include the physical robot, the physical human, and the physical objects and obstacles relevant to the task. That is the physical configuration of the system. There is also a random variable for the end user operating the robot, U_t . The user here includes the random variable representing the human intent, H_t , which the user intends to actuate over an input device, D_t . As can be seen from Figure 3, D_t depends on both the user's intent H_t and a noise Z_t . The noise repre-

sents the part of the interaction that does not depend statistically on the intended message of the human acting as the source. Here, the noise Z_t is assumed to be a simplified physical disability (see the "Modeling the User?" section). The human intent, H_t , depends probabilistically on the current state. Thus, the human user is able to observe the robot in the environment, an important part of the human-robot interaction. The robot has access to n collision sensors C and m proximity sensors P that depend directly on the state. The robot actuation, A_t , is assumed to represent the robot's interpretation of the action to be executed, based on the sensed state and the noisy commands received from the user. The robot's actuation generates the future state X_{t+1} , shifting the system downward as shown in Figure 3.

Then, the conditional probabilities representing the correlations of the different random variables can be defined. For example, the probabilistic dependence of the actuation of the robot can be written as $\text{prob}(a_t | d_t, c_t^1, \dots, c_t^n, p_t^1, \dots, p_t^m)$ for the model in Figure 3. These correlations can be thought of as noisy communication channels between the random variables. The information in a correlation can be represented formally by the mutual information [14]. See (1) for the case of two random variables, X and Y . Probability distributions can be estimated by a sampling process that counts how often a given value of one or more random variables occur, and metrics based on mutual information can, therefore, be calculated from the experimental data. See, for example, Lungarella et al. [15].

$$I(X; Y) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \text{prob}(x, y) \log_2 \frac{\text{prob}(x, y)}{\text{prob}(x) \text{prob}(y)}. \quad (1)$$

The overall goal of the system design can then be phrased and potentially quantified experimentally to maximize the mutual information between the user's intent and the robot actuation, i.e., maximize $I(H; A)$. However, for many tasks, the user's intent is not easily predicted. A more readily available information may be the noise Z , which can, for example, be a simulated disability in an experimental setting (see the "Modeling the User?" section).

Modeling the User?

A mathematical model of each end user could help design for personalized assistance but is challenging to achieve in practice. Human movement models are typically only feasible for specific actions or tasks, while a large set of diverse ADLs is required. In addition, physical disabilities come in many forms and can be hard to model exactly. So what can be done?

The tradeoff between speed and accuracy is a generic property of most human targeted movements. Models include Fitts' law [16] and the related Steering law, both of which have been used extensively in input device comparisons. For the higher-DoF movements involved in operating assistive manipulators, there are fewer models available; however, the necessity of trading off velocity with accuracy remains. In this type of experimental setting, the tradeoff

can be quantified by the mean time (MT) to complete a task and the number of collisions between the manipulator and the environment. It has been postulated that the speed–accuracy tradeoff stems from a signal-dependent neuromotor noise. That is, faster movements require greater forces in the muscles, which again introduces more neuromotor noise and movement variability [17]. An increase in the signal-dependent neuromotor noise has been related to stroke-related motor deficiencies [18].

It is clear that real end-user participants are needed to validate the clinical credibility of any assistive technology. However, the involvement of people with disabilities in a lengthy and tiring systematic activity of validation in intermediate stages of development must be reasonably limited. It can also be challenging to find a homogeneous set of disabled participants. Experimental comparisons can be made more reliable and easier to replicate by simulating a consistent disability for a set of able-bodied participants. That is, by introducing controlled perturbations in the perception-action loop. This shortens the prototype development and facilitates the experimentation on more robotics-specific research issues. However, the simulation of disabilities in able-bodied people is not obvious, as the neurophysiological modeling of disabilities is still a research issue. When a specific model cannot be obtained, or when it is not appropriate to focus on only one disability, a complementary approach is to assume a simplified physical disability that reduces the ability of the operator to control the system. Previous examples of this approach include assessments of computer mouse movements by disabled people [19] and of shared control with joystick input for assistive wheelchairs [8].

For the experiments reported here, an increase in the signal-dependent noise was used as a first-order approximation of a physical disability. A Gaussian noise (Z in Figure 3) was added to the raw able-bodied user input, in analogy with [19]. The noise was filtered to below 2 Hz to be comparable with typical human movements in the frequency domain. It was also made to increase in strength with the magnitude of the velocities commanded by the user, a signal-dependent noise making the speed–accuracy tradeoff more difficult [16]–[18]. Note that a disability can also require the use of less dexterous parts of the body for robot control, which are also typically subject to the speed–accuracy tradeoff. While the noise added does not necessarily correspond exactly to a specific real disability, it is an attempt at emulating the negative effect a disability could have on the ability of the user to accurately control the manipulator.

Performance Metrics

Given the previously mentioned models of system and user, what metrics are suitable for measuring performance? The MT to completion is a natural metric when measuring performance on tasks involving movements. But, given the speed–accuracy tradeoff, the high speeds associated with low completion times will typically lower the accuracy and, thus, increase errors. In some experiments on simple movements

(e.g., the Fitts’ law paradigm), the participants are instructed to maintain the error rate below a certain level (typically 2–5%). This helps disambiguate results with low times but a high number of errors, and vice versa, but is difficult to control for more complex tasks. Here, we include the consequences of an error in the MT by making each collision cost time. Put in the context of the signal-dependent noise in the human sensorimotor system [17], the user has to adjust the speed of execution to achieve a variance of the trajectories performed (variance of X in Figure 3) that statistically minimizes the average MT over attempts. This is the main performance metric used here, but it says little about the inner workings of the human–robot system.

Controllability is a useful metric for controllers acting in a closed loop in general. Loosely speaking, a system is controllable if it can be commanded to any final state from any initial state. For complex stochastic systems, a different formulation with broader conditions is desirable, such as the information-theoretic approach presented in [12]. Following this approach, we showed that controllability is also relevant for the specific case where the controller is a disabled user, H , trying to overcome the noise produced by his/her disabilities, Z , to successfully perform a manipulation task with a physically assistive robot [4]. For the experimental evaluation performed here, the controllability is approximated using the unconditional mutual information over the robot actuation and the noise added (A_t and Z_t in Figure 3). That is, we used $I(A_t; Z_t)$, measured over a finite set of tasks and a fixed number of repetitions for each task. The controllability metric, thus, measures the correlation of the noise added with the actuation of the robot. A lower correlation indicates that the system is better at rejecting the noise and has a higher controllability from the user’s perspective. The simplification assumes that: 1) the user attempts movements that are repeatable (i.e., passing through similar trajectories of X_t and H_t) and 2) the movements are close to optimal.

Enforcing such conditions is not always possible, however. A random user input could, for example, artificially inflate the controllability approximation used. Bialek et al. [20] proposed PI, the mutual information over the past and the future, as a general measure of the complexity of a time series. The measure can be said to quantify the total information of past experience that can be used for predicting future events, and has, among others, been used to drive the behavior of mobile robots in unknown environments. Applied to the robot actuation in Figure 3, it may be used to compare how predictable the actuation is with different shared-control approaches. The one-step mutual information was here used, with the PI of actuation thus being defined as $I(A_t; A_{t+1})$. The measure can also be used to identify cases where the controllability metric is no longer valid.

Case Study 1

Two case studies are used here as examples for the approach to replication outlined in the previous section. Both case

Figure 4. The robot implementations and sensors used for the case studies. (a) Case study 1: simplified robot implementation. The ranges of simplified proximity sensors are shown. (b) and (c) For case study 2, collision sensors (black squares) and simulated field of view for each proximity sensor (square pyramids) are shown: (b) the detailed view of the end-effector and (c) the manipulator.

Figure 5. The virtual experiment settings used for the case studies. (a) The simulated environment for case study 1. (b) The example translational trajectories for case study 1: participant 1, shared control. (c) The simulated environment for case study 2, with the robot in the initial resting position. (d) The example translational trajectories of the end-effector for case study 2: participant 6, shared control.

studies apply SDAC on a virtual assistive manipulator, and the system is tested on reaching tasks in simulated environments. The end users are simulated with able-bodied participants and simulated stochastic disabilities. In the first case study, a simplified model of an assistive robot manipulator was used, represented by a free-floating end-effector, as shown in Figure 4(a). The simulated environment as seen by the participants is shown in Figure 5(a). The five tasks performed consisted in moving a can from an initial position in front of the

user to a target position in a kitchen cabinet or on a kitchen desktop. A transparent and enlarged target can was used to indicate the position and tolerance (by the scale) required to finish each task. The user could control the translational velocities of the end-effector in the end-effector frame, $\vec{v}_{\text{robot}} = [v_x, v_y, v_z]^T$. For this case study, 20 proximity sensors were used, simulated as point distance detectors with a nominal range of 50–300 mm. The 20 discrete collision sensors were simulated with the same sensors, with a range of 10 mm.

Simulated End Users

The simulation of the robot in the environment shown in Figure 5(a) was displayed on a 20-in computer monitor at a distance of about 1 m. The input device used was a SpaceNavigator six-DoFs joystick. Noise was added to the user input, according to (2). This was Gaussian noise, low-pass filtered at 2 Hz and generated independently for each Cartesian component of the noise vector ($\vec{z} = [z_x, z_y, z_z]^T$). The magnitude of the velocity caused by the noise increased proportionally to the magnitude of the translational velocities commanded by the user, with some noise existing also when the user did not indicate movement (nonzero α_{noise}). See Figure 5(b) for example trajectories.

$$\vec{v}_{\text{user}} = \vec{v}_{\text{input}} + \vec{v}_{\text{noise}},$$

where

$$\vec{v}_{\text{noise}} = \vec{z}(\alpha_{\text{noise}} + \beta_{\text{noise}} \|\vec{v}_{\text{input}}\|). \quad (2)$$

Five able-bodied participants were used, all university graduate students at UC3M. There were two females and three males, all right-handed. All had experience with three-dimensional (3-D) input devices. The mean age was 26.8, ranging from 26 to 27.

Experimental Protocol

The testing lasted 1 h for each participant, and informed consent was obtained. First, each participant was given two training sessions to get familiar with the setup, followed by six main sessions. Each session had 25 trials per participant, with five repetitions of five different tasks. The shared control was activated for the last four sessions. A timer was shown to the participants, which began running when the participant pressed an indicated button. If the hand collided with the environment, it was reset to the initial position, while the timer kept running. Trials that had collisions were rerun. The participants were instructed to attempt to achieve the lowest times possible while keeping in mind that collisions were costly in terms of time. The experiment data were recorded at 50 Hz. The data used for calculating the mutual information were normalized and discretized to ten states. The mutual information was then estimated by the histogram method. The InfoMeth MATLAB toolbox [15] was used. The time series was the x , y , and z Cartesian components of the robot velocity ($\vec{v}_{\text{robot}}$) and the noise velocity ($\vec{v}_{\text{noise}}$), averaged in time over every four points recorded. All five successful attempts for all five tasks were used.

Figure 6. The completion time for both case studies. The upper whisker represents the most extreme data point below the limit: 1.5 times the interquartile range beyond the third quartile; similarly for the lower whisker and the first quartile. The notch is based on a 95% confidence interval for the median. The outliers are not shown for clarity. (a) Case study 1: Each box plot is based on data from five participants with ten repetitions of five tasks (250 observations). (b) Case study 2: Each box plot is based on data from eight participants with six repetitions of five tasks (240 observations).

Results

The completion time distributions for case study 1 can be seen in Figure 6(a). All participants had a reduction in the MT, with the shared control ranging from 13.2 to 29.2% improvement. There was a 23.1% reduction in the average metric over participants. A Shapiro–Wilk test indicated that the results in Figure 6(a) were not normally distributed, with $W = 0.764$, $p < 0.001$ for the benchmark condition and $W = 0.869$, $p < 0.001$ for the shared control condition. A Wilcoxon signed-rank test was, therefore, used to test for statistical difference. The MT over all tasks and attempts was used, giving one sample for each of the five participants. A weekly significant result was found, with $Z = 2.023$ and $p = 0.063$. A complementary paired t-test was also run, as the Wilcoxon signed-rank test has a limited resolution when used with only five samples. The difference was significant, with $t(4) = 4.82$, $p = 0.009$.

The results for the controllability for case study 1 can be seen in Figure 7(a). There was a decrease in the mutual information, and, thus, an increase in the controllability, for all velocity components with the shared control. The distributions were found to be likely nonnormal, and the Wilcoxon signed-rank test was, therefore, used. There were no strict significant differences found, although the z component was weakly significant, with $Z = 2.023$ and $p = 0.063$. Figure 7(b) shows the results for the predictability of execution for case study 1. There was an increase in the PI for all velocity components with the shared control. The distributions were found to possibly be nonnormal, and the Wilcoxon signed-rank test was, therefore, used. The differences for all components were weakly significant ($Z = -2.023$ and $p = 0.063$).

Case Study 2

The aim of the second case study was to attempt to evaluate the SDAC on a more realistic simulation of a robot arm. A

Figure 7. The information-theoretic metrics applied to the velocity components shown for case study 1, calculated over complete successful trajectories. Each box plot is based on the calculated metric for five participants and two sessions (ten measurements). The upper whisker represents the most extreme data point below the limit: 1.5 times the interquartile range beyond the third quartile; similarly for the lower whisker and the first quartile. (a) The metric for approximating controllability, $I(A_t; Z_t)$. Lower is better. (b) The metric for predictability of execution, $I(A_t; A_{t+1})$. Higher is better.

full assistive manipulator was simulated, based on the ASIBOT design [RPPPR, Figure 4(c)]. The simulated environment as seen by the participants is shown in Figure 5(c). The ASIBOT robot was simulated to be attached to the right-hand side of the user's wheelchair. The tasks performed involved moving the end-effector of the robot from an initial resting position to a pregrasp position around one of the five simulated cans in the virtual environment. The user could control four Cartesian velocities of the end-effector in the end-effector frame, $\vec{v}_{\text{robot}} = [v_x, v_y, v_{\text{pitch}}, v_{\text{yaw}}]^T$. The implementation had 68 proximity sensors in total. Eighteen were simulated as Vishay TCND5000, with a maximum sensed distance of 50 mm. These were all distributed over the end-effector, see Figure 4(b). The remaining was simulated as Sharp GP2D120, with a maximum sensed distance of 300 mm. All proximity sensors had a simulated 10° field of view. The voltage output of each proximity sensor, p_j , was simulated based on the distance measured and the calibration specifications of the two sensor types. In total, 229 individually distinguishable discrete collision sensors (on/off taxels) were simulated.

Simulated End Users

The simulation of the robot in the environment in Figure 5(c) was displayed on a 40-in (approximately 102-cm) display (Samsung 3-D TV, UE40D8000), at a distance of about 2 m. The simulation was displayed in 3-D, and the participants used active 3-D glasses. On a smaller display closer to the participant, the simulated view from the end-effector camera was shown, see insets in Figure 8. The input device used was a SpaceNavigator six-DoFs joystick. Noise was added to the user input to simulate a simplified physical disability, according to (3). This was Gaussian noise generated independently for each Cartesian component of the noise vector ($\vec{z} = [z_x, z_y, z_{\text{pitch}}, z_{\text{yaw}}]^T$) and low-pass filtered at 2 Hz. The

Figure 8. One attempt with shared control for case study 2. Cartesian x , z , pitch, and yaw components of velocities (in end-effector frame) are shown with time in seconds on the x axis. A discrepancy between the input ($\vec{v}_{\text{user}}$) and the output ($\vec{v}_{\text{robot}}$) velocities of the shared control means assistance is provided (from approximately 4 s onward for this example). The actual robot poses along the trajectory are shown, but the camera angle is altered for visualization. The corresponding view from the end-effector camera (as seen by the participant) is shown in the insets. The phases are described as follows. 0–4 s: highly coordinated gross movement, 4–8 s: adjustment of pitch during forward movement, 8–12 s: mainly yaw adjustments, 12–16 s: the final approach to the target.

magnitude of the translational velocities caused by the noise ($\vec{v}_{\text{noise, trans}}$) increased proportionally to the magnitude of the translational velocities commanded by the user ($\vec{v}_{\text{input, trans}}$), with some noise existing also when the user did not indicate movement (nonzero α_{noise}). This was similar for the rotational velocities ($\vec{v}_{\text{noise, rot}}$). See Figure 5(d) for example trajectories.

$$\vec{v}_{\text{user}} = \vec{v}_{\text{input}} + \vec{v}_{\text{noise}},$$

where

$$\begin{aligned} \vec{v}_{\text{noise, trans}} &= \vec{z}_{\text{trans}} (\alpha_{\text{noise}} + \beta_{\text{noise}} \|\vec{v}_{\text{input, trans}}\|), \\ \vec{v}_{\text{noise, rot}} &= \vec{z}_{\text{rot}} (\alpha_{\text{noise}} + \beta_{\text{noise}} \|\vec{v}_{\text{input, rot}}\|). \end{aligned} \quad (3)$$

Modeling the complete human-robot binomial, deriving closed-loop performance metrics from the models.

Experimental Protocol

The testing was performed over two days for each participant, with one hour of commitment per day. Informed

consent was obtained from each participant. The first day, the participants were given three sessions for training, followed by two sessions for establishing a benchmark. The second day, the participants were introduced to the shared control and were first given a maximum of two training sessions with the adaptation (learning) of the shared control activated. Then followed two sessions used to establish the performance with the (static) learned neural network weights. Each session consisted of three repetitions of each of the five target can locations, for 15 trials in total. For a given trial, the target can was red, while the remaining were blue. A trial was automatically judged as completed when the two fingers of the robot end-effector were positioned around the thickest part of the can, stopped or with a small remaining velocity magnitude. The arm was reset to the initial position if any part of the robot collided with the environment. The same instructions were given to the participants as in case study 1. See Figure 8 for an example execution of task 3.

For the data collection and conditioning, the same approach as in the “Experimental Protocol” section was followed. However, here the time series was the x , z , pitch, and yaw Cartesian components of the robot velocity ($\vec{v}_{\text{robot}}$) and the noise velocity ($\vec{v}_{\text{noise}}$). The data for two of the participants were not used for the time series analysis, as it was found that the noise was incorrect for some of their attempts. For this case study, attempts with collisions were not repeated, and the time series analyses were, therefore, conducted

over the last 6 s of the trajectories only (most attempts were above 6 s). All three attempts for all five tasks were used.

Results

Figure 6(b) shows the completion time distributions for case study 2. All participants had a reduction in the MT with the shared control, ranging from 5.3 to 59.9%. There was a 32.5% reduction in the average metric over participants. A Shapiro–Wilks test on the distributions in Figure 6(b) indicated nonnormality, with $W = 0.800$, $p < 0.001$ for the benchmark condition, and $W = 0.722$, $p < 0.001$ for the shared control condition. A Wilcoxon signed-rank test was, therefore, used to test for statistical difference. The MT over all tasks and attempts was used, giving one sample for each of the eight participants. A significant result was found, with $Z = 2.521$ and $p = 0.008$.

The results for the controllability for case study 2 can be seen in Figure 9(a). There was a decrease in the mutual information, and, thus, an increase in the controllability, for all velocity components with the shared control. The distributions were found to possibly be nonnormal, and the Wilcoxon signed-rank test was, therefore, used. The differences for z , pitch, and yaw were significant ($Z = 2.201$ and $p = 0.031$). The results for the predictability of execution for case study 2 can be seen in Figure 9(b). There was an increase for all velocity components with the shared control. The distributions were found to possibly be nonnormal, and the Wilcoxon signed-rank test was, therefore, used. The differences for all components were significant ($Z = -2.201$ and $p = 0.031$).

General Discussion

The three main metrics defined showed improvements in performance for both case studies. This indicates that the approach followed for modeling and benchmarking the system has captured some important features of the effect of the shared control. The consistent improvement also gives an indication of the potential of the SDAC approach to shared control of assistive robot manipulators, although there is room for improvement in the experimental method. One particular issue stems from the fact that the shared control was performed after the benchmarking. This was driven by the overall time limitations and the need to gradually train the participants. Future experiments should attempt to minimize any unwanted skill-transfer effects by performing training, benchmarking, and shared control on separate days, with counterbalancing. A larger set of participants would also be desirable to explore the statistical significance of the results seen here.

By including the effect of collisions in the time to complete the task, it is clear to both the participant and experimenter what the speed–accuracy tradeoff is for the task. This makes the MT with collisions included a reasonable metric for overall task performance. However, in the current setup, collisions occurring later in the movement are, in effect, penalized the strongest. This could be improved upon, for example, by using a fixed time penalty or one related to

Figure 9. The information-theoretic metrics applied to the velocity components shown for case study 2, calculated over the last 6 s of trajectories. Each box plot is based on the calculated metric for six participants and two sessions (12 measurements). The upper whisker represents the most extreme data point below the limit: 1.5 times the interquartile range beyond the third quartile; similarly for the lower whisker and the first quartile. (a) The metric for approximating controllability, $I(A_i; Z_i)$. Lower is better. (b) The metric for predictability of execution, $I(A_i; A_{t+1})$. Higher is better.

the severity of the collision. The controllability and predictability metrics can complement this information and give the experimenter confidence about the inner workings of the complex human–robot system.

Conclusions

We proposed a methodology for the replication of experiments with closely coupled human–robot systems. In addition to the robot software, hardware, and data components, we believe such an experimental article must include also the complete experimental method and experimental setup relevant to the human participants. That is, we need a sufficient description of the complete system to allow for an equivalent developmental trajectory in the two adaptive agents involved, the robot and the human. An information-theoretic approach to modeling and benchmarking was presented and applied to a distributed and adaptive shared control of assistive robot manipulators. We hope to produce experiments that can unambiguously answer whether, but also in what aspect, the shared control system has had a positive effect. For the SDAC approach, this translated into a 23.1–32.5% average improvement in the MTs achieved (case study 1 and 2), and consistent improvements in the controllability from the user’s perspective and the predictability of execution.

We hope to produce experiments that can unambiguously answer whether, but also in what aspect, the shared control system has had a positive effect.

The Future

In the future, we would like to extend the evaluations to real end users on physical implementations and to promote replication of the results obtained. We would also encourage other researchers in the field to help make replicable experiments on similar human-robot system available, through their own channels or on <http://throughput.sourceforge.net>.

We explore a complementary low-level collision-limitation behavior, dubbed SDAC.

Acknowledgments

This research received funding from the ARCADIA project DPI2010-21047-C02-0, thanks to the MINECO ministry of Spain, and the ROBOCITY2030 II project S2009/DPI-1559, thanks to Comunidad de Madrid and EU structural funds.

References

- [1] F. Bonsignorio, J. Hallam, A. P. del Pobil, Eds. (2008). GEM Guidelines, Euron GEM SIG Report. [Online]. Available: <http://www.heronrobots.com/EuronGEMSig/>
- [2] F. Bonsignorio, J. Hallam, and A. P. del Pobil, "Good experimental methodologies in robotics: State of the art and perspectives," in *Proc. Workshop Performance Evaluation Benchmarking Intelligent Robots Systems, IEEE/RSJ Int. Conf. Intelligent Robots Systems*, San Diego, CA, 2007.
- [3] A. Bonarini, W. Burgard, G. Fontana, M. Matteucci, D. G. Sorrenti, and J. D. Tardos, "RAWSEEDS: Robotics advancement through web-publishing of sensorial and elaborated extensive data sets," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems Workshop Benchmarks Robotics Research*, Beijing, China, 2006.
- [4] M. F. Stoelen, V. F. Tejada, A. Jardón, F. Bonsignorio, and C. Balaguer, "Benchmarking shared control for assistive manipulators: From controllability to the speed-accuracy trade-off," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Vilamoura, Portugal, 2012, pp. 4386–4391.
- [5] M. F. Stoelen, V. F. Tejada, J. G. Victores, A. Jardón, F. Bonsignorio, and C. Balaguer, "Adaptive collision-limitation behavior for an assistive manipulator," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Tokyo, Japan, 2013, pp. 1143–1148.
- [6] M. F. Stoelen, V. F. Tejada, A. Jardón, F. Bonsignorio, and C. Balaguer, "Distributed sensing, learning and control in an assistive manipulator," in *Proc. IEEE Multi-Conf. Systems Control*, Antibes, France, 2014, pp. 1348–1353.
- [7] T. Carlson and Y. Demiris, "Collaborative control for a robotic wheelchair: Evaluation of performance, attention, and workload," *IEEE Trans. Syst. Man, Cybern. B*, vol. 42, no. 3, pp. 876–888, 2012.
- [8] D. Vanhooydonck, E. Demeester, A. Huntemann, J. Philips, G. Vanacker, H. Brussel, and M. Nuttin, "Adaptable navigational assistance for intelligent wheelchairs by means of an implicit personalized user model," *Robot. Auton. Syst.*, vol. 58, no. 8, pp. 963–977, 2010.
- [9] T. L. Chen, M. Ciocarlie, S. Cousins, P. M. Grice, K. Hawkins, K. Hsiao, C. C. Kemp, C.-H. King, D. A. Lazewatsky, A. E. Leeper, H. Nguyen, A. Paepcke, C. Pantofaru, W. D. Smart, and L. Takayama, "Robots for humanity: Using assistive robotics to empower people with disabilities," *IEEE Robot. Automat. Mag.*, vol. 20, no. 1, pp. 30–39, 2013.
- [10] P. F. M. J. Verschure, B. J. A. Krose, and R. Pfeifer, "Distributed adaptive control: The self-organization of structured behavior," *Robot. Auton. Syst.*, vol. 9, no. 3, pp. 181–196, 1993.
- [11] E. Cheung and V. J. Lumelsky, "Proximity sensing in robot manipulator motion planning: System and implementation issues," *IEEE Trans. Robot. Autom.*, vol. 5, no. 5, pp. 740–751, 1989.
- [12] H. Touchette and S. Lloyd, "Information-theoretic approach to the study of control systems," *Phys. A, Stat. Mech. Applicat.*, vol. 331, nos. 1–2, pp. 140–172, 2004.
- [13] K. P. Valavanis and G. N. Saridis, "Information-theoretic modeling of intelligent robotic systems," *IEEE Trans. Syst. Man, Cybern.*, vol. 18, no. 6, pp. 852–872, 1988.
- [14] C. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.
- [15] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, "Methods for quantifying the informational structure of sensory and motor data," *Neuroinformatics*, vol. 3, no. 3, pp. 243–262, 2005.
- [16] P. M. Fitts, "The information capacity of the human motor system in controlling the amplitude of movement," *J. Exp. Psychol.*, vol. 47, no. 6, pp. 381–391, 1954.
- [17] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nature Rev. Neurosci.*, vol. 9, no. 4, pp. 292–303, 2008.
- [18] P. H. McCrea and J. J. Eng, "Consequences of increased neuromotor noise for reaching movements in persons with stroke," *Exp. Brain Res.*, vol. 162, no. 1, pp. 70–77, 2005.
- [19] J. Mankoff, H. Fait, and R. Juang, "Evaluating accessibility by simulating the experiences of users with vision or motor impairments," *IBM Syst. J.*, vol. 44, no. 3, pp. 505–517, 2005.
- [20] W. Bialek, I. Nemenman, and N. Tishby, "Predictability, complexity, and learning," *Neural Comput.*, vol. 13, no. 11, pp. 2409–2463, 2001.

Martin F. Stoelen, Centre for Robotics and Neural Systems, University of Plymouth, United Kingdom. E-mail: martin.stoelen@plymouth.ac.uk and mfstolen@gmail.com.

Virginia Fernández de Tejada, Universidad Carlos III de Madrid, Spain. E-mail: vtejada@ing.uc3m.es and virginia.ftejada@gmail.com.

Alberto Jardón Huete, Universidad Carlos III de Madrid, Leganes, Spain. E-mail: ajardon@ing.uc3m.es and ajardonster@gmail.com.

Carlos Balaguer, Universidad Carlos III de Madrid, Leganes, Spain. E-mail: balaguer@ing.uc3m.es.

Fabio Paolo Bonsignorio, The Biorobotics Institute, Scuola Superiore Santa Anna, Pontedera, Italy, and HeronRobots, Genoa, Italy. E-mail: fabio.bonsignorio@sss.it and fabio.bonsignorio@heronrobots.com.

RA