

A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human–Robot Interaction

Dylan P. Losey

Department of Mechanical Engineering,
Rice University,
Houston, TX 77251
e-mail: dlosey@rice.edu

Craig G. McDonald

Department of Mechanical Engineering,
Rice University,
Houston, TX 77251
e-mail: Craig.G.McDonald@rice.edu

Edoardo Battaglia

Centro di Ricerca “E. Piaggio”,
University of Pisa,
Pisa 56126, Italy
e-mail: e.battaglia@centropiaggio.unipi.it

Marcia K. O’Malley

Professor
Fellow ASME
Mechatronics and Haptic Interfaces Laboratory,
Department of Mechanical Engineering,
Rice University,
Houston, TX 77251
e-mail: omalleym@rice.edu

As robotic devices are applied to problems beyond traditional manufacturing and industrial settings, we find that interaction between robots and humans, especially physical interaction, has become a fast developing field. Consider the application of robotics in healthcare, where we find telerobotic devices in the operating room facilitating dexterous surgical procedures, exoskeletons in the rehabilitation domain as walking aids and upper-limb movement assist devices, and even robotic limbs that are physically integrated with amputees who seek to restore their independence and mobility. In each of these scenarios, the physical coupling between human and robot, often termed physical human robot interaction (pHRI), facilitates new human performance capabilities and creates an opportunity to explore the sharing of task execution and control between humans and robots. In this review, we provide a unifying view of human and robot sharing task execution in scenarios where collaboration and cooperation between the two entities are necessary, and where the physical coupling of human and robot is a vital aspect. We define three key themes that emerge in these shared control scenarios, namely, intent detection, arbitration, and feedback. First, we explore methods for how the coupled pHRI system can detect what the human is trying to do, and how the physical coupling itself can be leveraged to detect intent. Second, once the human intent is known, we explore techniques for sharing and modulating control of the coupled system between robot and human operator. Finally, we survey methods for informing the human operator of the state of the coupled system, or the characteristics of the environment with which the pHRI system is interacting. At the conclusion of the survey, we present two case studies that exemplify shared control in pHRI systems, and specifically highlight the approaches used for the three key themes of intent detection, arbitration, and feedback for applications of upper limb robotic rehabilitation and haptic feedback from a robotic prosthesis for the upper limb. [DOI: 10.1115/1.4039145]

1 Introduction

The interaction between man and machine has changed considerably over the course of history. What started as a simple physical interaction with basic tools has transformed over time, with the tools becoming complex machines with sensors and sophisticated controls. This change has been even more evident in recent years, with modern robots developing increasing autonomy and capability to the point where they can actively interact with a human partner toward achieving a common task: this has led to human robot interaction (HRI) emerging as a new field of research aimed at maximizing the performance, efficiency, and applicability of coupled human–robot systems.

Human robot interaction can be defined as “a field of study dedicated to understanding, designing, and evaluating robotic systems for use by or with humans” [1]. This is a broad definition, and indeed HRI is a multifaceted discipline. For example, some applications are natural extensions of the industrial and manufacturing settings that originally supported automation via robotics [2,3]. In this domain, HRI systems enable high-level task planning and flexibility achievable with trained human operators compared

to preprogrammed industrial robots, while still leveraging the repeatability, precision, and load-carrying capacity of robots. Socially interactive robots are also gaining prominence [4], and are suited to applications where robotic systems are socially situated, embedded, and intelligent, with a focus on cognition, social behavior, and natural interactions (verbal, visual, and typically nonphysical) with human partners. HRI has also reached to the medical and healthcare sector, where we see robots used for minimally invasive surgery [5], and even being worn to improve mobility and independence [6,7].

As these applications imply, some forms of human robot interaction involve direct physical contact [3,7,8], often referred to as physical human robot interaction (pHRI). While much of the literature related to pHRI has traditionally had a strong focus on ensuring safety during the interaction between human and robot [9], we present this review from the viewpoint of the overall shared control architecture that is designed to achieve a desired, physically coupled, and cooperative pHRI task. First, we will provide our perspective on shared control in pHRI, and define a foundation for our review based on the themes of intent detection, arbitration, and feedback. Each theme is defined and reviewed in detail, interactions between themes are explained, and examples from the state of art will be presented to illustrate how similar concepts are present in applications from seemingly different

Manuscript received January 11, 2017; final manuscript received January 22, 2018; published online February 14, 2018. Editor: Harry Dankowicz.

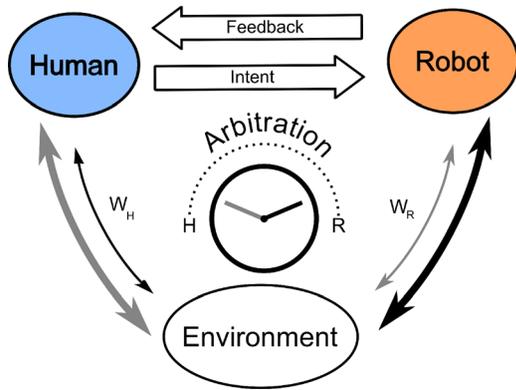


Fig. 1 Conceptual representation of the proposed framework: human and robot exchange information and interact with the environment according to what is decided by the arbitration (represented by the knob)

fields. Finally, two case studies from the authors' prior work are described, showing in detail how the framework can be applied in the design of two prototypical pHRI systems.

2 Overview

In this survey, we explore human-robot shared control over a collaborative task for applications where the human is physically coupled to, and cooperating with, the robotic device. While the traditional pHRI framework focuses strictly on applications where there is a direct physical contact between human and robot [9], here we extend our survey to also consider applications where physical interaction is mediated through a third object. This allows us to consider additional relevant applications, such as cooperative manipulation tasks, where an object is jointly manipulated by a human and a robot to achieve a common goal. This extension can also include bilateral teleoperation tasks where a human remotely controls a robot, with a haptic feedback channel conveying information to the human user regarding the physical interactions that are occurring between the robot and the remote environment. We present a general framework to describe the interaction process, with the aim of organizing design procedures from different subfields of pHRI.

We propose a framework for considering shared control between humans and physically coupled robots that features three key ideas. First, in each of our selected applications, the robot requires some knowledge of the human's goals and intents so that the robot behavior can be controlled accordingly. We term this *intent detection*, and will begin our survey by defining intent, and then exploring methods for measuring and interpreting intent in pHRI systems. Second, the interaction between human and robot and the way each affects the environment are regulated by *arbitration*, which we define as the mechanism that assigns control of the task to either the human or the robot. Finally, we posit that it is essential that the human be provided with information about the task and environment characteristics, and, where appropriate, suggested trajectories or task completion strategies that are developed by the robotic partner. Therefore, *feedback* from the robot to human is returned via some sensory channel, often haptic, so as to leverage the physical coupling that already exists between human and robot. We have illustrated this framework in Fig. 1. In this schematic model, arbitration is represented as a knob: when control is assigned primarily to the robot (darker shaded arrow), its energy exchange with the environment will be greater; conversely, if control is assigned primarily to the human, the energy exchanged between human and environment will be greater. The bilateral exchange between robot and human represents the robot detection of human intent, and the provision of feedback to the human.

These three elements (intent detection, arbitration, and feedback) can be used to model many applications of physical human robot interaction. In Ref. [10], for example, a cooperative manipulation task is presented where a human and a robot collaborate to move a bulky table. Intent detection was performed by using force/torque sensors and processing their measurements with a mathematical model of the task; arbitration was realized by controlling a role allocation parameter derived from task modeling; and feedback to the user was provided haptically through the cooperatively manipulated object. The framework can also be applied to the pHRI task of myoelectric control of a robotic upper limb prosthesis. Here, intent detection is achieved by monitoring surface electromyography (sEMG) signals; arbitration can be realized by directly mapping EMG activity to the actuators of the prosthesis to control grip pose, while maintaining automated low-level control of grip force to prevent an object from slipping from the prosthetic gripper's grasp [11]; and feedback can be provided to the human using haptic devices on the residual limb or embedded in the socket interface. In Secs. 3–5, we expand on each framework element, providing examples and implementation guidelines from the literature, and comparing approaches from different fields of pHRI. In particular, while the framework presented is general, we will focus on the context of rehabilitation in the rest of the paper.

Rehabilitation can be thought of in two contexts. First, pHRI applications in rehabilitation can be compensatory in nature, where human intent is detected to control a robotic device that replaces lost capabilities (e.g., myoelectric prostheses, or exoskeletons as mobility aids for paraplegics). In other scenarios, the objective is to promote partial or complete recovery from neurological injury such as stroke or spinal cord injury. These applications clearly require a distinct set of design requirements, as the objectives differ greatly. In the first case, we want to integrate a robot with the human control system, while, in the second case, we seek to promote neural recovery so that the participant can function independent of the robot after treatment is complete.

3 Intent Detection

We will define the problem of intent detection as the need for the robot to have knowledge of some aspect of the human's planned action in order for the robot to appropriately assist toward achieving that action. Therefore, the robot's ability to detect user intent relies directly on some channel of communication existing between the human and the robot. The structure of this section is to look at the three aspects of the unidirectional channel of communication of intent from user to robot (Fig. 2). First, the user's intention must be defined, and when referring to the many different forms in which intent can be defined, we will use the phrase *intent information*, or sometimes simply *intent*. Second, the modality by which intent information is measured by the robot must be decided, which we will refer to as the method of *intent measurement*. Finally, once the information reaches the robot, there is the more open-ended question of how this measurement is

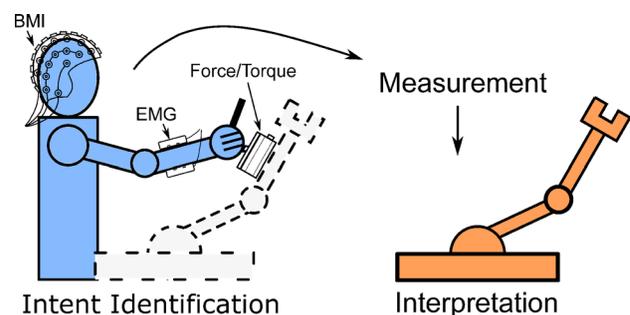


Fig. 2 The three steps for conveying of the human's intent to the robot: identification, measurement, and interpretation

to be understood by the robot as a representation of the intent information, and how it is to be incorporated into the robot control structure. We will refer to this aspect of the problem as *intent interpretation*. Consideration of approaches to intent interpretation begins to blur the line between intent detection and arbitration—a larger division that we have drawn in the shared control problem framework. Our discussion of intent detection will make some mention of robot control strategies, but will seldom descend into detail since arbitration will receive its own treatment in Sec. 4. Further, mechanisms of communication from the robot to the human will be discussed in the section on communication.

The system designer has significant freedom in how to approach all aspects of the intent detection problem; therefore, in our review of the intent detection literature, we aim to (1) expose the reader to some extent of the range of approaches that have been employed in the literature and (2) draw connections across seemingly disparate areas of research where similar strategies of intent detection have been employed. This review is not exhaustive, even within our narrowly defined list of applications. We do believe, however, that we have at least covered examples of the more common strategies for defining intent, measuring intent, and interpreting intent.

3.1 Defining Intent. We propose a unified definition of user intent. Human motor control is complex, involving activity in the central nervous system (CNS), the peripheral nervous system (PNS), the musculoskeletal system, and finally the environmental interactions that are being controlled. At each of these subsystems, there are measurable physical quantities, called state variables, which are manifestations of the human's intention. Additionally, we know that the state variables local to one subsystem are tightly linked to those of another subsystem. For instance, sensory afferent neurons in the PNS send information to the CNS, which guides motion planning in the motor cortex, while the motor cortex also sends neural commands back down to the PNS. Therefore, the user intent is most generally described by many different subsystem states containing different forms of intent information, which exist simultaneously and give rise to one another. Though only some information will be relevant for a specific application, the common characteristics of intent are that it can be represented by states that describe the human system and that this information has been deemed relevant to the task by the system designer.

For many applications, intent can be defined in a binary way. For the arbitration of effort between human and robot, it seems natural to ask, is the user actively trying to control the interaction or not? This type of intent—active versus passive—is defined in Ref. [12] for a hand shaking robotic application. In a more clinical setting, the movement of a robotic exoskeleton or prosthetic limb is often automated. This allows for the human intent to simply be defined as a trigger to initiate motion, often ascertained from a brain-machine interface (BMI), as seen in Ref. [13].

In other applications, there are more than two possible discrete states for intent. For instance, in order to manage the complexity of an upper limb robotic prosthesis, a user is often given predefined poses, grasps, or functions that the prosthesis can complete autonomously. In Refs. [14–17], upper limb prosthetics are controlled by the user selecting one of these predefined functions; therefore, the user intent is represented by a single categorical variable.

The control of lower limb exoskeletons and orthoses is another application where intent can often be reduced to selecting from a set of predefined motions. In Ref. [18], the wearer of the lower extremity assistive device is defined to have one of eight possibly intended motion states related to sitting, standing, or walking. The definition of user intent is similar in several of the devices reviewed by Yan et al. [19].

Intent can also be defined in terms of continuous variables. For those working with patients who are undergoing neurorehabilitation,

one of the most important questions is whether or not the user is actively engaged in completing the motor task at hand. Sarac et al. [20] extend a typical method for extracting a binary classification of move/rest to now output a continuously varying signal. They interpret this signal to be the “level of intention” of the user, which is then mapped to the speed of task execution.

In many examples, user intent is defined in terms of a velocity or position trajectory—the predicted forward path of the user and/or robot. In Ref. [21], an intelligent walker defines the intent of the user in the form of a predicted forward path over a short time horizon, while in Ref. [22], a cane robot uses similar methods to ascertain the user's direction of intended motion and the “magnitude” of the intention in that direction. Short-time-horizon forward paths can also be parameterized, as in Refs. [23] and [24], so that the parameter estimates serve as the user intent.

A slightly different form of user intent is a continuously time-varying desired position. It can be thought of as the reference trajectory for the control of a robotic manipulator, and it is often used as the input to a robot impedance controller. Ge and coworkers [25,26] define the motion intention in this way for a human and robot performing collaborative motions with shared end-effector position, as do Erden and Tomiyama [27].

The human's intent can be defined as a continuously time-varying force or torque. In the 2015 review by Yan et al., we see this as a common definition of intent for lower-limb exoskeletons. Pehlivan et al. [28] define the user's intent as the interaction force between the person and the wrist exoskeleton at the handle, and Lenzi et al. [29] look at the effective torque about the user's elbow within an elbow exoskeleton. In Ref. [28], the interaction force is estimated from position sensing and modeling of the robot dynamics, while in Ref. [29], the interaction torque is estimated from measurements of the user's muscle activation. Still, both applications could be said to have defined the same type of intent information.

Finally, we consider the example of a human wearing a seven degrees-of-freedom (DoF) upper-limb exoskeleton in work done by Kiguchi and Hayashi [30]. This is one example of how defining the human intent for a shared control problem requires us to consider multiple forms of intent information simultaneously. In Ref. [30], the exoskeleton is to provide powered assistance to the user as they perform unstructured tasks. The resulting controller structure defines signals such as the user's muscle activation signal, the estimated joint torque generated by the user, the force vector at the hand that results from the estimated joint torques, and the acceleration of the hand that should result from the hand force vector. All of these signals could possibly be under the control of the user, and therefore be a part of their intention. In theory, the designer should be able to measure any one of these signals and, with appropriate modeling techniques, reconstruct the others for the use of the robot controller.

3.2 Measuring Intent. In this section, the details of the selected applications will become clearer as we describe the specific methods by which researchers measure the different forms of intent information that were described in Sec. 3.1. The different methods of measurement—or, measurement modalities—that we discuss should be thought of as not entirely dependent on the intent information that is being measured. In other words, there is not a one-to-one mapping between an intent definition and a corresponding measurement modality.

There are a variety of neural methods for measuring intent information that have recently become available to us thanks to advancements in neuroscience. The first we will discuss is the technique of electroencephalography (EEG), which is an example of what is commonly referred to as a BMI. In brief, an array of electrodes measure electrical activity of the cortex at varying degrees of proximity to the surface, depending on the method. In order to measure user intent, researchers have used surface EEG, where the electrodes are placed noninvasively along the scalp

[20]. They have also made use of intracranial EEG, also called electrocorticography, where the electrodes have been placed on the surface of the brain as the result of a surgical procedure [13]. In both cases, the electrodes measure voltage relative to some reference, and the signal is referred to as the local field potential (LFP). We know that encoded within the LFP is information about the activity of the region of the brain local to the electrode. From there, depending on the intent information that we wish to extract, decoding the LFP can be carried out in a number of ways. This is reserved for Sec. 3.3.

Myography is the measurement of the activations of human muscles and their resulting contractile forces, and it can be accomplished in a number of ways. The most common method seen in the literature is EMG, which measures changing electrical potentials that result from the activation of motor units near electrodes. Surface EMG is more commonly used because it is completely noninvasive, though it is limited to measuring superficial muscles near the surface of the skin. We see examples of surface EMG being used to measure user intent in Refs. [14] and [31–33]. If greater specificity or deeper muscle recordings are needed, there is intramuscular EMG, which uses a fine needle electrode inserted into the muscle.

Recently, force myography (FMG), also known as topographic force mapping or muscle pressure mapping, has received attention as a possibly less expensive and, in some cases, more robust method of myography. FMG is also noninvasive; it infers muscle forces by detecting changes in muscle volume underneath tactile sensors placed on the surface of the limb. We see it used in Refs. [15–17] with success that certainly warrants further investigation.

Another form of myography, known as sonomyography, is based on ultrasound imaging and has only recently been introduced. Akhlaghi et al. [34] use a conventional clinical ultrasound system to generate images of a cross section of the forearm muscles, which are then used to train a database of ultrasound activity corresponding to different hand motions. A simple nearest neighbor classifier is then tested on real-time data to identify hand motions with a classification accuracy that lies within the range of reported results for similar sEMG systems. The main advantage of the sonomyography technique is that it can sense the activity of deep muscles, such as those in the forearm that control the fingers.

Traditional load cells can measure force and torque with a high accuracy in up to 6DoF. In the case where a human and a robot interact through a simple interface such as a handle, load cells provide the best measurement of the exchange of effort at the interaction point. They also come with a high cost and increased fragility. Nonetheless, Wang et al. use a 6DoF force/torque sensor at the robot end effector to monitor the human–robot interaction during a hand shaking task [12]. Huang et al. sensorize their intelligent walker with a load cell in each of its two handles [21]. Wakita et al. also use a 6DoF load cell in the interface of their omnidirectional-type cane robot as the primary source of sensory information [22]. By contrast, a popular force-sensing alternative is the use of compliant force sensors in series with the actuator and the user, commonly known as series elastic actuation. Pratt et al. demonstrate the integration of a series elastic actuator into a knee exoskeleton in Ref. [35], the principle being that by adding a linear spring into the actuator’s ball-screw transmission, measurements of the spring deflection can be converted directly into actuator output force by Hooke’s law.

At this point, we will note that the accessibility of a type of intent information, when paired with the chosen measurement modality, will determine the robustness of the signal measurement to environmental noise. For instance, the detailed activity of the CNS is generally inaccessible without invasive surgery, due to the fact that this intent information arises deep within the human body. Using EEG, one of the most direct, noninvasive techniques at our disposal, to study the motor cortex is much more sensitive to noise than simply measuring the resulting motion of the body. Likewise, surface EMG measurement of muscle excitation is

increasingly sensitive to noise the deeper the muscle is relative to the skin surface, while a direct measurement of the forces produced by that same part of the body in contact with a force/torque sensing load cell will be more reliable. However, in exchange for robustness, the external measurements of kinematics and kinetics lack any information regarding how the body is planned to achieve that outcome. The external signals also occur after some time delay in comparison to signals generated by the nervous system, which may or may not be acceptable depending on how quickly the robot should react to the human. Taking into account such tradeoffs, system designers may choose to measure one type of intent signal in order to estimate another, thanks to the flexibility that comes with intent interpretation.

3.3 Interpreting Intent. Now that we have covered what the intent information is that we are trying to measure, and what measurement tools we have at our disposal, the final component of the intent detection problem is our interpretation of the measurement. This deserves particular attention because for many applications, the intent information that we wish to know and the signal that we are able to measure are not necessarily the same thing, but they are related. It is at this point that we will make greatest use of our modeling of the system, both human and robotic, as well as simplifying assumptions that need to be made.

A common neural-based approach to intent interpretation for BMIs uses the neural intent to predict the user’s movement intention, with the output being simply a trigger that initiates the movement of a robot or prosthesis. McMullen et al. [13] use the method of decoding, or classification, of the inputs called linear discriminant analysis (LDA), and, along with support vector machines, it is one of the most common techniques for making inferences about neural data. Sarac et al. [20] make a clever extension of the LDA algorithm to obtain more than just a binary output of move/rest. The same two classes (move and rest) are used, but for each new data point, the posterior probabilities of each class are calculated providing a continuous output between zero and one. The authors map this continuous output to the task execution speed, which is a single parameter in the robot motion controller. This slightly artificial mapping is used to encourage patient engagement in therapy, and is an excellent example of the flexibility that exists in the interpretation step for the designer of the shared control architecture.

There are two main approaches in interpreting myographic signals: pattern recognition and mapping to continuous effort. Despite its complexity, we see an early example of using pattern recognition on the features of a single sEMG channel to control an upper-limb prosthesis in Ref. [31]. The pattern recognition approach is to map patterns of any number of signal features to desired prosthesis poses, grasps, or functions. The way in which control algorithms learn this mapping varies, but common approaches are LDA, support vector machines, and artificial neural networks. Simpler versions can use only a few features, such as the use of three electrode signals variances in Ref. [32], to learn to distinguish between a few discrete operating states. There has been much progress in this method over the years, with more current work involving many channels of EMG [14] or FMG [17]. In both Refs. [14] and [17], the authors also show improved robustness to potential disturbances such as limb position.

While pattern recognition approaches are appealing for their ability to learn arbitrary mappings between myographic signal features and intent information, they are most commonly used to select from a relatively small number of discrete control states. For applications such as neurorehabilitation, where sometimes the goal is only to trigger the appropriate robot motion at the appropriate time, this is acceptable. However, other applications, including other modes of neurorehabilitation, would benefit from extracting continuously time-varying information such as the user’s desired joint torques. A simple approach to this problem is to match the EMG signal amplitudes of agonist–antagonist muscle

pairs to antagonistic cable actuation systems for rotary exoskeleton joints, and then to hand tune a proportional gain from the EMG signal, postprocessing, to the assistive torque provided by the robot [33]. The assumption here is that subjects that have been weakened by neurological damage will benefit from a robot providing torque that is in the same direction and roughly proportional to the user's desired torque. Lenzi et al. take an interesting approach, which is to leverage the ability of the human CNS to adapt their EMG activation to minimize the error between their intended motion and observed robot motion [29] so that proportional control becomes a sufficient control strategy for the robot. While this approach has the potential to greatly simplify the design of control systems for powered exoskeletons, it has the significant drawback of disturbing a user's natural motor function at the neural level.

Other researchers have worked to develop control systems that explicitly learn the mapping between EMG signals and user's desired joint torque. An excellent example by Kiguchi and Hayashi is the use of an adaptive neuro-fuzzy modifier to learn the relationship between the root-mean-square of measured EMG signals and the estimated user's desired torques [30]. The approach is to use an error-backpropagation learning algorithm to modify the mapping, which is expressed as a weighting matrix. The neuro-fuzzy modifier takes as inputs the joint angle measurements provided by the robot in order to account for the effect of varying limb position on EMG signals.

Examples of lower-limb exoskeletons from previous decades, which are covered more extensively in Ref. [19], make use of ground reaction force sensing and lower-limb kinematics to estimate joint torques [35]. Once again, the estimated joint torque at the knee has a hand-tunable scaling factor applied to produce the commanded actuator torque, under the assumption that the user's desired torque is reflected accurately in the torque they are able to generate. Such assumptions must be applied carefully to situations where there is user impairment or other environmental factors that limit the user's ability to tightly control the measured force/torque at the interaction point.

When designing controllers for mechanical systems, the variable to be controlled is often position and the variable representing the controller effort is force. It is then no surprise that in applications of human-robot shared control systems, the user intent is often defined as the force generated by the user. Consequently, many examples of user intent detection revolve around estimating the user contribution to the interaction force. Pehlivan et al. use an inverse dynamics model of the robot, along with knowledge of the actuator commands and a predefined motion trajectory, to estimate the user force applied to the robot from the robot encoder measurements [28]. This user-intended force is then subtracted from the robot controller effort so that it assists the user minimally in achieving the predefined trajectory.

Instead of measuring robot position to estimate interaction force, one can measure interaction force between the human and the robot at the end effector and estimate the desired human position. Ge et al. [25] extract the user's desired position by assuming a model of the user's control strategy—in this case, a linear impedance (mass-spring-damper). The user's intended position is then assumed to be the equilibrium point of the spring. The authors use a radial basis function (RBF) neural network, which has the property of universal functional approximation, to learn an estimate of the human dynamics. Li and Ge [26] extend their previous work so that the synaptic weight vector of the neural network can be updated in real time to respond to changing human impedance.

Just as researchers have used interaction force measurements to estimate a desired position of the human [25,26], the measured interaction force can be used to estimate other forms of motion intention. For example, in Ref. [21], the interaction forces measured in the handles of an intelligent walker are fed into a model of the nonholonomic walker dynamics to predict the walker's forward path over a short time horizon. Force/torque sensing in

the handle of an assistive cane was used by Wakita et al. to ascertain the hidden walking state of the user [22]. The user's state is a discrete variable representing possible walking modes, e.g., "go straight forward" or "turn to the right," as in Ref. [18]. The detection of the walking state is paired with the use of a Kalman filtering (KF) technique—based on the forward dynamics of the cane robot—to estimate the direction and magnitude of the user's desired acceleration. Finally, Erden and Tomiyama [27] present a unique interpretation of human intent obtained from the measured interaction force between a human hand and a HapticMaster robot. The robot is under impedance control; thus, using the principle of preservation of momentum, it follows that the integral of the controller force applied by the robot in the time period between two stable resting states is equal to the total momentum delivered by the human interaction. Therefore, the authors use the integral of the robot controller force—which can also be called the impulse—as a measurement of the human intention and define a user's desired change in set point position of the robot as being proportional to the impulse by some tunable scaling factor. This final relationship is based solely on an intuitive understanding of the load dynamics and the ways in which humans tend to manipulate objects. It is a convenient substitution, since relating the impulse to the desired set point position means the intent can easily be given as an input to the robot impedance controller.

Another paradigm for intent interpretation involves the use of robot position measurement to predict future motion. Corteville et al. [24] assume the human to be in control, so the motion of the human-robot interaction point is assumed to result entirely from the human intention. Position sensing is used to estimate a minimum jerk, bell-shaped profile of the user's desired speed that is continuously updated. The minimum-jerk speed trajectory has been in used by many researchers as a model for human movements [36]. Corteville et al. have allowed the robot to assume that at any instant, the intended human velocity is to have this bell shape; therefore, estimation of the curve parameters is equated with estimation of the human's motion intent. Under the same paradigm, Brescianini et al. make use of a combination of simple position and force sensing embedded in a pair of augmented crutches [23]. From these signals, the authors extract gait parameters such as stride length, height difference, direction, and operation mode. These values are then used to generate motion trajectories for a lower-limb exoskeleton that is being worn along with the crutches.

The final example of intent interpretation we will examine is the use of force and position to estimate human impedance. Wang et al. [12] have a 10DoF hand-shaking robot with a force/torque sensor mounted in the end effector. The robot end effector is simply a metal rod that a human may grasp as if it is the other partner in a hand shake. An impedance relationship can be defined between the measured position and orientation of the robot end effector and the resulting forces and torques measured at the end effector, resulting from interaction with the human. The human is then modeled as a linear impedance with three parameters—mass, damping, and stiffness. Using the recursive least squares algorithm for online parameter estimation, the current human impedance parameters are heuristically classified as being "low" or "high" and are then used as the inputs to a hidden Markov model (HMM) to decide if the person intends to be "active" or "passive" in the handshake interaction with the robot. The authors have made extensive use of their model of the human control, i.e., the consistent relationship between a hidden user state and their resulting impedance parameters, and the linear impedance control employed by the human—to infer a more abstract definition of user intent from mechanical sensing at a lower level.

4 Arbitration

Arbitration here refers to the division of control among agents when attempting to accomplish some task. More specifically, during physical human-robot interaction with shared control,

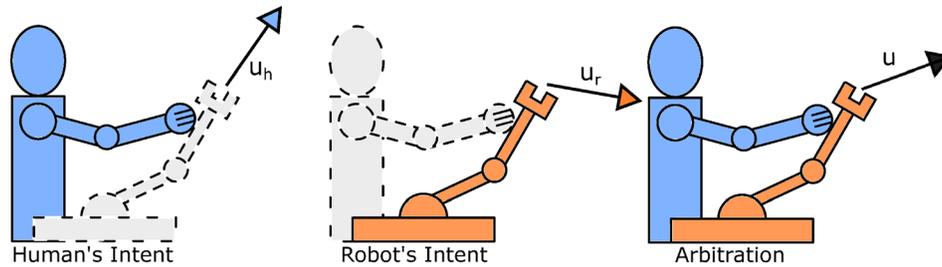


Fig. 3 A simple arbitration between human and robot, where together the human and robot are sharing control of the position of the robot's end effector. On the left panel, the robot uses intent detection in order to infer the human's desired motion, u_h . In the middle panel, we show the robot's intended direction of motion, u_r , which is tangent to the desired trajectory. Finally, on the right panel, we arbitrate between the human and robot intents, and thus the robot's end effector moves in a direction u , which compromises between u_h and u_r .

arbitration determines how control is divided between the human and robot. Many different types of arbitration are possible; for instance, the human might be responsible for controlling the position of the robot's end effector, while the robot controls the end effector's orientation. Alternatively, both the human and robot could be jointly responsible for the position and orientation of the robot's end effector, but have different relative levels of influence. A simple example of this kind of arbitration for shared control is shown in Fig. 3.

Clearly, whenever the human and robot are both actively working together to accomplish some task, arbitration of some sort is either implicitly or explicitly present. Recalling the trio of features, which Bratman [37] has theorized, are essential to shared cooperative activity, we observe that shared control requires a "commitment to mutual support." Within shared control, each agent has an objective, and, if either the human or robot needs help, they should be able to expect the other agent to provide assistance toward achieving that objective. Without arbitration, it would be unclear how (or when) the other agent is meant to intervene and provide this "mutual support"—if the human communicates an unexpected intent, for example, should the robot (a) attempt to correct that intent, or (b) defer to the human? An understanding of arbitration within physical human–robot interaction is therefore necessary for shared control.

In fact, arbitration—in some form or fashion—has already been described as an integral part of human–robot interaction, even outside the field of shared control. In their 2007 review, Goodrich and Schultz [1] describe the "level and behavior of autonomy" as one of the five key attributes that affect interactions between humans and robots. Our use of arbitration is roughly analogous to these "levels of autonomy"; the first level of autonomy, where the robot does nothing and the human completes the task, and, at the other end of the spectrum, the tenth level of autonomy, where the robot completes the task while ignoring human intent, are simply two very different instantiations of arbitration. On the one hand, to reduce the human's burden, we generally want the robot to be as autonomous as possible. On the other hand, however, we note that autonomy in shared control is limited by Bratman's commitment to mutual support, since, if the completely autonomous robot ignores the human's intent altogether, there is no opportunity to work collaboratively or offer assistance. So as to resolve this conflict and better understand arbitration within our context of shared control, we briefly turn to recent studies of physical human–human dyads, where two humans are working together to accomplish the same task.

Reed and Peshkin [38] examined physical human–human interaction during a simple 1DoF task, and focused on haptic communication of information (see Fig. 4). They found—like other researchers—that dyads of humans working together completed the task more quickly than a single individual working alone, and, of particular interest, they discovered that humans naturally assume different roles during task execution. For instance, one

human might take a "leader" role, and actively move the co-manipulated object, while the other human could take a "follower" role, and passively resist motion. Although this phenomenon is not yet fully understood, subsequent work by Ueha et al. [39] explored not only the tangential forces, which Reed and Peshkin [38] had reported, but also the radial forces during a similar 1DoF rotational task. Here, the authors found that one human naturally assumed control of the tangential forces, which are related to larger motions, while the second human took control of the radial forces, which are related to finer positioning; again, a natural arbitration of roles emerges.

Finally, Feth et al. [40] provided further experimental evidence for the existence of roles in physical human–human interaction based on the asymmetric energy flow between agents. Like in Ref. [38], one human is active, injecting energy, and the second is passive, dissipating energy. Moreover, these studies of human–human dyads found that roles varied dynamically over time so that a human who once served as leader could become a follower, or, by the same process, the human who had assumed a follower role could transition into leadership. Hence, in order to make human–robot interaction more similar to human–human interaction, arbitration in shared control should assign dynamic "roles," i.e., provide a framework that allows both agents to contribute, and meaningfully change their type of contributions over time.

Viewed together, these concepts of arbitration in cooperative activity and studies of arbitration within human–human dyads suggest two fundamental questions for arbitration and shared

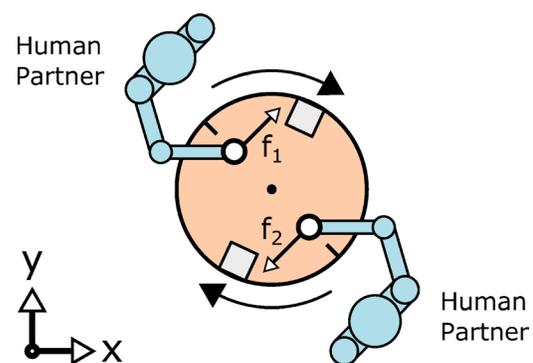


Fig. 4 Experimental setup used by Reed and Peshkin [38]. Two human partners are working together to rotate a crank to a desired orientation (gray boxes). During these experiments, the vision of the partners was occluded, and only haptic communication was allowed. It was found that human partners naturally assume different roles, and that performance improves when the task is performed by human dyads, as opposed to a single human operator.

control: (a) how should roles be allocated and (b) how should these roles be dynamically updated? In Secs. 4.1 and 4.2, we will review the ways in which other research groups have addressed these questions when considering physical human–robot interaction with shared control. This review is not meant to be comprehensive, and will almost certainly omit several exciting works; the works we have included, however, are meant to provide the reader with a sense of the complexities and benefits associated with the different types of static and dynamic role arbitration.

4.1 Types of Role Arbitration. An excellent taxonomy for the different types of role arbitration in shared control was recently published by Jarrassé et al. [41], and is based on concepts from neuroscience and game theory. The authors argue that both the human agent and the robotic agent have an inherent cost function, which consists of a sum of *error* and *effort* components, and each agent naturally attempts to minimize their individual cost function at a Nash equilibrium. By *error*, we here mean a difference in position or orientation with respect to the agent's desired trajectory or goal pose, and by *effort*, we mean the amount of force, torque, or muscle activation, which an agent applies during interaction. In what follows, we will classify the types of role allocation using the same general convention introduced by Jarrassé et al. [41]: co-activity, master–slave, teacher–student, and collaboration. These four types of role allocation are distinguished by differences in the robotic cost functions; the human is always assumed to minimize their own perceived error and effort.

Within co-activity the task is divided into subtasks, and the human and robot are assigned unique subtasks which can be completed independently. In this case, the cost associated with an agent is a function of that agent's own error and effort, and when one agent changes his or her error or effort, it does not directly alter the cost of the other agent. By contrast, in a master–slave role allocation, both agents are attempting to complete the same task, and the cost of the robot is defined to be the sum of the human's error and effort. Therefore, the robot will here exert as much effort as possible to minimize the human's error and effort without any regard for the robot's own error and effort, i.e., the human is the “master” and the robot is the “slave.” Next, within a teacher–student role allocation, the robot again seeks to minimize the error of the human, but does this while encouraging human involvement. Accordingly, the robot considers its own effort, and gradually attempts to reduce its effort as the human performs more of the task independently, i.e., the human is the “student” and the robot is the “teacher.” Finally, in collaborative or partner–partner role allocation, the robot and human are assumed a priori to be equals, and together complete the same task while considering their individual errors and efforts. This role allocation for pHRI is most similar to the human–human dyads previously discussed. We will now attempt to classify examples of arbitration in shared control based on these four types of role allocation, noting that some research is difficult to place in a single category because it contains elements of multiple role allocation types.

4.1.1 Co-Activity. To begin, we consider co-activity, or a division of subtasks, which is particularly applicable when there are aspects of the task that one agent is unable to perform. Perhaps, the most illustrative instances of this can be found in powered prosthetics, where the human is necessarily unable to actuate the prosthesis herself, and must instead rely on the prosthesis correctly carrying out their communicated intent. Work by Varol et al. [42] considers this problem for lower-limb prosthetics. These authors assume that the human wants their prosthesis to be in one of three different states—either sitting, standing, or walking—and the authors have also developed automated transitions between these states. Given the robot's interpretation of the human's intent to sit, stand, or walk, the robot transitions to or remains in the most appropriate state. Hence, the roles are allocated such that the human's subtask is to decide on the desired

state, and the robot's subtask is to execute the motions relevant to that state. Wakita et al. [22] likewise use stop, going forward, and turning as discrete states for a cane-like mobile robot, which assists and supports human walking. It is assumed in Ref. [22] that the human cannot walk without help from the cane, but the human can communicate intents related to their desired state; as before, the robot is allocated the subtask of stabilizing the human's movement.

Myoelectric control of upper-limb powered prostheses likewise relies on the robot performing the human's desired motions. In this case, the main difficulty is leveraging signals generated by the human's muscles in order to control artificial hands that often possess a high number of actuated DoF. Typically, this problem is resolved with pattern recognition techniques [43,44]. While classification results can be reasonably accurate, users have reported that the prosthesis is still difficult to control [45], and so this process remains an open avenue for research. Interestingly, while upper-limb prostheses generally try to leave full control to the user, they can also include some levels of autonomy, such as for lower level slip prevention tasks [11].

Along these same lines, we should also quickly mention brain controlled wheelchairs [46–48], where the human again has a higher level decision task, and the robot is responsible for a lower level navigation subtasks. Philips et al. [46] suggest that the robotic system should be responsible for collision avoidance, obstacle avoidance, and/or orientation recovery while the human communicates EEG signals, which correspond to moving the wheelchair left, right, or straight ahead. A similar scheme is presented by Carlson and Millan [47], where by default, the wheelchair moves forward while avoiding obstacles, and the human communicates the arbitrated intention to move right or left. The subtask of the robot can further be expanded to include holistic navigation; in Rebsamen et al. [48], the user specifies a goal location from a set of discrete options, and the robot's subtask involves generating and following a trajectory to reach that goal.

Co-activity in shared control, however, is not limited simply to applications where the human is physically unable to carry out certain aspects of the task. Aarno et al. [49] developed an approach for teleoperation and co-manipulation where the desired task is segmented into several states, and each of these states has an associated virtual fixture. It may be possible for the human to perform the teleoperation or co-manipulation task completely alone, but the inclusion of these subtasks, and a probabilistic estimation of the current subtask, was found to improve the human's performance. Indeed, since the human is involved in all aspects of the task's execution, this research combines components of both co-activity (the delegation of subtasks) and master–slave arbitration (virtual fixtures along those subtasks).

4.1.2 Master–Slave. The master–slave role allocation, with human masters and robotic slaves, is likely the most traditional and ubiquitous type of role arbitration in shared control. We contend that virtual fixtures and impedance/admittance controllers are classical instances of the master–slave role allocation, since for both virtual fixtures and impedance/admittance control, the robot's error is the same as the human's error—i.e., the deviation from the desired trajectory—and the robot has no explicit cost associated with its effort. Although the term “master–slave” might imply that the robot has little autonomy, in fact this role allocation scheme can encourage the robot to do as much of the task as possible. Utilizing impedance control, for instance, a robot can follow the desired trajectory without any human participation, while still responding naturally to external perturbations [9]. Moreover, we would point out that in some sense, the (human) master and (robot) slave arbitration should always be present within shared control, because, for safety purposes, the human must always retain final authority during situations where the human and robot are in conflict [9].

Abbott et al. [50] and Bowyer et al. [51] provide surveys of recent work on virtual fixtures, which are synonymously referred

to as active constraints. Although Refs. [50] and [51] describe a wide variety of different virtual fixtures, we can generally state that virtual fixtures attenuate or nullify the human's intent in cases where performing this intent would lead to undesirable outcomes, such as increasing error or decreasing performance. Virtual constraints have primary applications in surgical teleoperation and co-manipulation, and can deal with situations where certain areas of the workspace are out of bounds (forbidden-region virtual fixtures) or where the human seeks to follow a desired trajectory (guidance virtual fixtures). Impedance/admittance controllers can be thought of as guidance virtual fixtures.

Of special interest for our study of arbitration are the unique works on virtual fixtures by Yu et al. [52] and Li and Okamura [53]. Yu et al. [52] augment the role of the robotic slave to include more autonomy so that the robot is responsible not only for identifying the human's overarching intent but also for defining corresponding virtual fixtures to satisfy that intent. Li and Okamura [53] provide a methodology for the robot to discretely switch the virtual fixture on or off, depending on the human's communicated intent. Practically, turning off the virtual fixture provides humans the freedom to leave the desired trajectory when seeking to avoid unexpected obstacles. Theoretically, removing the virtual fixture amounts to shifting from a master-slave role allocation to a role allocation where the human completes the task alone, in other words, the first level of autonomy as described by Goodrich and Schultz [1].

Interestingly, this switch between discrete master-slave and single-agent role allocations is also prevalent in rehabilitation robotics studies; for example, see Mao and Agrawal [54]. These authors implement a virtual tunnel around the desired trajectory, within which only a constant tangential force is applied to help keep the human moving (see Fig. 5). If the human accidentally moves outside of the tunnel, however, a master-slave role allocation is invoked, and the robot uses impedance control to correct the human's positional error. Another combination of master-slave and single-agent role allocation for rehabilitation applications is offered by Duschau-Wicke et al. [55]. Here, the human is completely responsible for the timing of their motions—without robotic assistance—but the robot uses impedance control to constrain positional errors with respect to the given path. Thus, as pointed out by Jarrassé et al. [41], the master-slave role arbitration can lead to an unanticipated contradiction; because control

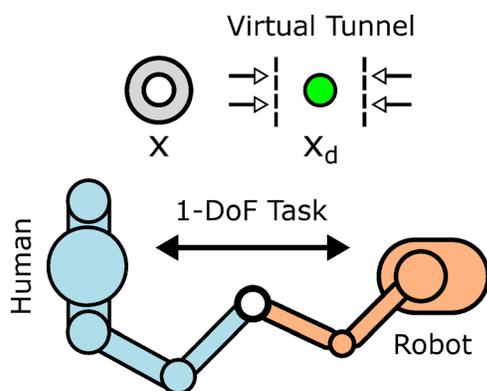


Fig. 5 Using a virtual tunnel to arbitrate roles between human and robot, the human and robot are attempting to follow a desired trajectory during a 1DoF task. The current position, x , is given by the torus, and the desired position, x_d , is given by the sphere. When the current position is within the virtual tunnel, the robot does not provide the human any assistance. When the current position is outside of the virtual tunnel, like shown, the robot provides haptic feedback (arrows) to guide the human back toward the desired trajectory. This combines both master-slave (human master, robotic slave) and teacher-student (robotic teacher, human student) role arbitrations.

over so much of the task is arbitrated to the slave—even to the extent where the slave performs the task autonomously—the master has little or no incentive to participate in cooperative shared control. Fortunately, this is not an issue for surgical applications, where the human's involvement is required due to safety concerns [51]. On the other hand, master-slave role allocations can undesirably de-incentivize human participation during rehabilitation applications [56], which, in turn, leads to both the combinations of master-slave and single-agent arbitration that we have discussed, as well as teacher-student role allocations.

4.1.3 Teacher-Student. The teacher-student role arbitration is well suited for situations where we are attempting to train humans using robotic platforms [57], which, considering the application areas focused on in this review, primarily entails robotic rehabilitation. The teacher-student role allocation is distinguished from discrete combinations of master-slave and single-agent arbitration, since teacher-student role arbitrations constantly attempt to reduce the amount of robotic effort. As explained by Blank et al. [56], extensive research in the field of rehabilitation robotics has argued that increasing the patient's level of engagement is important to improving neural plasticity, and is therefore a means to facilitate recovery from stroke or spinal cord injury. Shared control strategies, which employ the teacher-student role arbitration in the field of rehabilitation robotics, are typically referred to as assist-as-needed (AAN) controllers.

Assist-as-needed controllers balance the cost of decreasing the human's error with the cost of increasing the robot's effort; the objective of AAN controllers can therefore be posed as maintaining a suitable level of “challenge” such that the human always remains actively engaged in their therapy. In other words, as argued by Wolbrecht et al. [58] and Pehlivan et al. [28], AAN controllers should strive to provide the minimal amount of assistance, which guarantees that the human completes his rehabilitation exercise without an unreasonable amount of error—as determined by the human's capabilities—but the human is ultimately responsible for further reducing their error. Since, as experimentally shown by Emken et al. [59], the human motor system learns new motions while greedily minimizing kinematic errors and muscle activation, AAN controllers intentionally allow or introduce errors, and these, in turn, desirably motivate the human's muscle activation.

In practice, teacher-student role arbitration is often effected by starting with a master-slave role arbitration, and then reducing the robot's effort whenever possible. Our group has employed this technique in Ref. [28], where we decreased the gains of an impedance controller if the human's performance satisfied some threshold, or, conversely, increased the gains if the human was unable to complete the task successfully. Hence, as the human grows more adept at the task over time, the initial master-slave impedance controller can gradually become a single-agent role allocation, where the human is responsible for performing the task alone. A similar instantiation of the teacher-student role allocation may also be achieved using forgetting factors, such as in Wolbrecht et al. [58] and Emken et al. [60], which introduce a scheduled and exponential decrease in the robot's contributed effort. Again, the arbitration here gradually shifts from master-slave—where the robot can complete the task even with an unskilled or passive human operator—to a single-agent arbitration—where the human must complete the task without any assistance.

It should be understood, however, that teacher-student role allocation does not monotonically shift from master-slave to single agent; the mentioned works [28,58,60] all incorporate features that can increase the robot's (teacher's) assistance when the human (student) is regressing. Extending this line of reasoning, Rauter et al. [61] developed an AAN scheme that shifts among master-slave, single-agent, and antagonistic allocations, where in the antagonistic role arbitration, the robot expends effort to intentionally increase the human's error. Perhaps the teacher-student role arbitration, particularly in AAN or rehabilitation applications,

is best characterized as a dynamic role arbitration, where the robot's role continuously adjusts from more helpful (slave) to less helpful (uninvolved or antagonistic) depending on the human's motor learning and participation.

4.1.4 Collaboration. Collaborative arbitration for shared control is thus similar to teacher–student arbitration, because the equitable relationship between human and robot implied by collaboration also leads to very changeable, or dynamic, roles. Unlike teacher–student arbitration, however, which we discovered to be applied primarily within rehabilitation, collaborative arbitration has more general application areas. The majority of papers surveyed below [10,62–65] use collaborative arbitration for co-manipulation tasks, where the human and robot are both grasping a real or virtual object, and together are attempting to move that object along a desired trajectory, or place that object in a desired goal pose. We might imagine, for example, a human and robot moving a table together. Alternatively, work by Dragan and Srinivasa [66], which includes a brief review of arbitration, atypically develops a collaborative arbitration architecture for teleoperation applications. Here, the human's inputs are captured using motion tracking, from which the robotic system probabilistically estimates the human's desired goal via minimum entropy inverse reinforcement learning. The robot then arbitrates between the inputs of the human and its own prediction of the human's goal in order to choose the manipulator's motion. We note that for both co-manipulation and teleoperation applications of collaborative arbitration, the robot is theoretically meant to act as a human-like partner—i.e., an equal partner—and therefore the arbitration of roles should emulate what is naturally found in physical human–human collaborative interaction.

As we have previously discussed, roles naturally develop during physical human–human interaction [38–40]. These roles can be generally classified as an active leader role and a passive follower role [38,40], where both human participants can dynamically take and switch between the complementary roles. Extending these experimental results into human–robot interaction, researchers such as Li et al. [64] argue that the robot ought to actively reduce the amount of human effort, and accordingly assume a leader role by default. Furthermore, at times when the human and robot disagree, the robot should yield control back to the human, and quickly transition into a follower role.

Consider, for instance, the work conducted by Evrard and Kheddar [62]; these authors provide a simple mathematical framework for the robot to interpolate between leader and follower roles. The robot simultaneously maintains a leader controller, which minimizes errors from the desired trajectory (high-impedance), and a follower controller, which reduces the forces felt by the human (zero-impedance). The robot then continuously switches between these two controller outputs based on externally assigned leader/follower role arbitration. Mörzl et al. [10] further formalize the concept of leader and follower role allocations for co-manipulation applications by using redundant degrees-of-freedom—where both the human and robot can contribute forces—and nonredundant degrees-of-freedom—where the actions of the human and robot are uniquely defined. When dividing forces within the redundant (voluntary) degree-of-freedom, the human can provide more of the voluntary effort, making the robot a follower, or the robot can perform more of the voluntary effort, thereby taking a leader role.

We might wonder, however, how the robot should behave when this correct arbitration between collaborative leader and follower roles is not externally provided. Works by Medina et al. [63], Li et al. [64], and Thobbi et al. [65] address this question by introducing a level of confidence in the robot's predictions. When the robot is confident that the human will behave a certain way—for instance, follow a known trajectory [63,65]—the robot assumes a leader role in order to reduce the human's effort. On the other hand, when the robot is confident that the human wants to deviate from the robot's current trajectory—i.e., the human is persistently

applying strong forces on the robot's end-effector [64]—the robot takes a follower role in order to resolve this conflict and reduce the human's effort.

Although the leader and follower roles appear to be the most prevalent form of collaborative arbitration, research by Kucukylmaz et al. [67] and Dragan and Srinivasa [66] combine aspects of co-activity with collaborative arbitration. Here, the human takes the leader role during larger scale motions, which direct the robot toward its goal pose, while the robot takes the leader role during finer, smaller scale motions, i.e., precise positioning of the end effector when approaching the goal pose. This form of collaborative arbitration, which was also found in human–human dyads [39], is somewhat akin to co-activity, where the human's subtask might entail larger, less constrained movements, and the robot is tasked with the smaller, intricate motions.

4.2 Dynamic Changes in Role Arbitrations. In the preceding discussion of the different types of role arbitration, we have already seen indications that role arbitrations can dynamically change during task execution. These changes could occur within the same type of role allocation—such as switching between leader and follower roles during collaborative arbitration—or between two different types of role allocation—such as gradually transitioning between master–slave and single-agent roles during teacher–student arbitration. In general, however, dynamic changes in role arbitration are meant to either increase the robot's level of autonomy at the expense of the human's authority, or, conversely, increase the human's control over the shared cooperative activity at the expense of the robot's autonomy. Referring back to Fig. 3, the arbitrated movement (panel 3) can shift to become more like the human's intent (panel 1) or the robot's intent (panel 2). In what follows, we will outline the two predominant tools employed within the shared control literature to determine when to change arbitration: machine learning and performance metrics.

By “machine learning,” we here refer to techniques such as HMMs [49,52,53,68], Gaussian mixture models [42], and RBFs [58]. These data-driven approaches typically require a supervised training phase, where the human practices communicating intents with known classifications; after the model is trained, it can be applied to occasionally change role arbitrations in real time. Works by Li and Okamura [53], Yu et al. [52], and Aarno et al. [49] leverage HMMs in order to change role arbitrations for teleoperation and co-manipulation applications. Using the robot's measurements of the human's intent, as well as precomputed transition and emission matrices, these HMMs probabilistically determine which role arbitration “state” the human is currently attempting to occupy. For instance, in Ref. [53], HMMs are used to determine whether or not the human wants to follow a given trajectory—if the human is attempting to follow that trajectory, the stiffness of the virtual fixture increases, and thus our arbitration shifts toward the robot; on the other hand, if the human is attempting to leave the given trajectory, the compliance of the virtual fixture increases, and arbitration shifts toward the human. Similarly, in Ref. [52], the states include (a) following the trajectory, (b) avoiding obstacles, and (c) aligning the end effector, while in Ref. [49], the states consist of virtual fixtures along different line segments; just like before, once a state is detected, the role arbitration shifts toward the human or robot as appropriate.

Another interesting approach for adjusting arbitration using HMMs was recently proposed by Kulic and Croft [68]. In their research, the HMM attempts to estimate the human's affective (emotional) state from measured physiological signals including heart rate, skin conductance, and facial muscle contractions. First, the human and robot are placed in close proximity, and the human's affective state is estimated in response to the behavior of the robotic manipulator. Next, if the human is “alarmed” by the robot's movement, arbitration shifts toward the human, and the robot begins to move more slowly and/or away from the human. Though this application with affective states is not particularly

common, it certainly provides a natural way to convey a sense of the human's "confidence" or "trust" in the robot's actions, which can then be used to intuitively update the role arbitration.

Indeed, a 2011 meta-analysis by Hancock et al. [69], which examined the different factors that can affect human-robot interaction, found that robotic performance has the largest and most identifiable influence on trust in HRI. It seems reasonable, therefore, to employ performance metrics as a means to dynamically change role arbitrations between human and robot. A straightforward performance metric for this purpose could simply be the amount of force or torque applied by the human; both Kucukylmaz et al. [67] and Li et al. [64] implement this method for co-manipulation applications. In essence, when the human applies larger efforts, these authors argue that the human is actively attempting to take control of the task, and hence arbitration should shift toward the human. Conversely, when the human is passive, and not significantly interacting with the robot, arbitration switches to grant the robot a larger portion of the shared control. A related performance metric was developed by Thobbi et al. [65] and Mörzl et al. [10]; here, when the human consistently applies forces in the direction of the robot's motion, the robot becomes more confident in its prediction, and assumes a greater arbitration role. Alternatively, when the human interacts with the robot in a manner at odds with the robot's internal predictions, the robot returns control to the human and begins to develop new predictive models.

More generally, we can imagine that the human has a reward function, which the robot can learn from human-robot interactions [66,70,71]. As before, the robot lets the human take control during interactions, and then resumes autonomous behavior after the human stops interacting. Next, based on how the human interacted, the robot updates its estimate of the human's reward function—i.e., what behavior is optimal—and then replans the rest of the task in accordance with this new reward function. Dragan and Srinivasa [66] have applied this concept to robotic teleoperation systems which are unsure of the human's goal position: when the human inputs new commands into the teleoperation interface, the robot updates its estimate of the desired goal. Once the robot is quite confident that the human is trying to reach a particular goal, then the robot becomes more dominant in the shared control; when the robot is unsure, however, the human moves with little robotic assistance. Works by Losey and O'Malley [70] and Bajcsy et al. [71] have extended this online adaptation to learn the robot's desired trajectory based on physical interactions between the human and robot, leading to changing role allocations during the current task. At the other end of the spectrum, we can also use measured outcomes from the previous task to update role allocations for the next task—in Pehlivan et al. [28], the human's cumulative error with respect to the desired trajectory is used to adjust the impedance gains for subsequent trials.

Updating between trials is best suited for applications where the human and robot will be performing the same task for multiple iterations [56], just like the optimization approach proposed by Medina et al. [63]. Within their work, Medina et al. explicitly consider how to deal with uncertainty in human behavior. The human is recorded performing the task multiple times, and, by incorporating the Mahalanobis distance, arbitration shifts to favor the human over portions of the trajectory where large motion uncertainty is present: for instance, should the robot and human go around an obstacle in a clockwise or counterclockwise direction? Next, risk-sensitive optimal control is introduced to determine how the robot will respond to conflicts with the human: should the robot yield leadership, or take a more aggressive, dominant role? To conclude, we summarize the discussed performance metrics [10,28,63–67] as fundamentally derived from the concept of trust between human and robot. As the human comes to trust that the robot will behave as expected, and, simultaneously, as the robot better learns what the human wants to accomplish, arbitration can dynamically change to increase the robot's level of shared control.

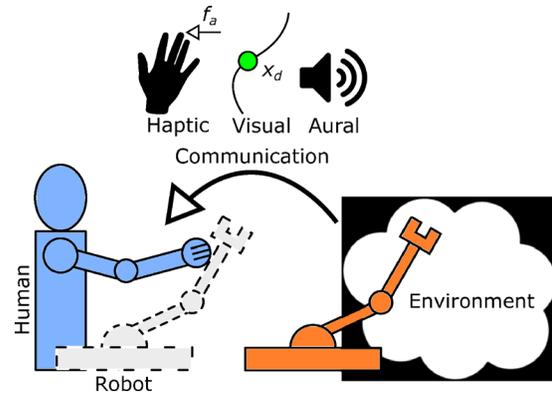


Fig. 6 Simplified schematic of communication between the human and robot during shared control. Three different modalities of communication are shown: haptic, visual, and aural feedback. Feedback is based on the robot's interaction with the environment, on right, where the environment could be virtual (such as in rehabilitation) or physical (such as for prosthetics). Thus, the kinesthetic haptic feedback force, f_a , emulates virtual or real robot-environment interaction forces. The visual feedback depicts the desired trajectory, as well as the desired position at the current time, x_d . Aural feedback can provide information on errors or instructions to assist the human operator.

Once a type of role arbitration has been determined, and the current role allocation is decided, the robot can provide feedback to the user based on both the environment and this arbitration. From this feedback, the user can better infer the current arbitration strategy, and understand their role within the interaction.

5 Communication

When a human and robot are physically coupled and sharing control of a task, such as an amputee using an advanced prosthetic device, the user depends on the robotic system to not only replace the function of the missing limb, detect their intent, and arbitrate control of the task, but also to communicate back to the human operator the properties of the environment. A similar situation occurs in bilateral telemanipulation, where force cues that arise between the remote tool and environment are relayed to the human operator at the master manipulator. In applications where a robotic device held or worn by the human operator is intended to instruct or assist with task performance, such as would be the case for surgical simulators, motor learning platforms, and exoskeletons for gait rehabilitation or upper limb reaching, to name a few, it is necessary to convey not just task or environment forces either real or virtual, but also the desired actions and behaviors that the human should execute. Further still, one can picture scenarios where the human user should be informed of the intent or future actions of the robot. In this section, examples of the methods employed by robotic systems to communicate with the human operator in shared control scenarios are surveyed using example applications.

5.1 Modalities of Sensory Feedback. The communication mechanism between human and robot in a coupled shared control system typically relies on the sensory channels available for information conveyance. For example, feedback can be provided visually, aurally, or haptically. For applications of physical human-robot interaction, the haptic channel is of particular interest because the force-motion coupling between action on the environment and resultant forces and actions can be leveraged in much the same way that our own body uses sensors embedded in the muscles to modulate the forces that we impose on the environment. A depiction of these different types of sensory feedback can be seen in Fig. 6.

5.1.1 Haptic Feedback. Haptic feedback, which is a general reference to cues that are referred via our sense of touch, can be subdivided as kinesthetic feedback (forces and torque applied to the human body and sensed at the muscles and joints), and cutaneous or tactile (forces and sensations sensed through the mechanoreceptors in our skin). Kinesthetic feedback requires complex, custom haptic devices unique to a particular task to be trained (for example, multi-degree-of-freedom devices to simulate rowing [72,73] or tennis swings [74]). Some devices are used to convey forces to the upper limb moving on a planar working surface [75] using an end effector-based design. Alternatively, exoskeletal-based robotic designs aim to provide prescribed feedback to the joints of the upper [76] or lower [77–79] limbs. Such devices must convey large forces and torques, and therefore tend to be heavy and expensive. Further, for applications such as haptic guidance for training, kinesthetic-based haptic feedback devices have been ineffective when it comes to demonstrating retention of skill or transfer to a similar task [57,80,81], despite their success at enhancing performance when the individual is coupled to the device [82]. These results provide further support for the guidance hypothesis [83], which advises that augmented feedback, such as the haptic guidance used in these studies, can be detrimental to learning if relied upon for more than just reducing errors. In these haptic guidance studies, we suspect that subjects were additionally depending on the guidance to develop strategies for task completion. Work by Winstein et al. showed that frequent physical guidance in a target reaching task resulted in poor retention and skill transfer [84], which provides explanation for the observed findings in these haptic guidance studies.

Given these drawbacks to kinesthetic type haptic feedback systems, our recent work has focused on the development of prototype wearable haptic feedback devices that could, for example, be used to provide sensory information to an amputee [85–89]. These devices use a variety of haptic feedback modalities, including vibration, skin stretch (resulting from shear force on the skin), and pressure to encode haptic information about the state of a prosthetic device. We have demonstrated, in a few focused studies, the potential of these devices to improve object manipulation in prosthesis control in human subject experiments focused on grasp and lift tasks, which are common in activities of daily living and therefore of interest to prosthesis users. A grasp and lift task is an appealing choice to investigate touch feedback in dexterous manipulation because it involves coordinating grip force and load force with object weight [90]. It is a planned movement and requires an internal model of the object's properties. Healthy individuals can use touch sensations to develop this internal model, but upper limb amputees rely primarily on vision since their prosthetic devices lack the provision of touch feedback. Therefore, it is imperative to investigate the effect of haptic feedback via sensory substitution in human performance of grasp and lift tasks. These studies will illustrate the value of providing haptic feedback, both to understand the impact on task performance, and on the participant's ability to maintain and update their internal model of the objects and task.

We have conducted a series of experiments using a simplified grasp and hold task in a virtual environment and haptic feedback to explore the effectiveness of such feedback for improving performance in object manipulation, specifically for grasping and lifting objects without slipping. In these experiments, a user interacts with a virtual environment via a SensAble Phantom, an off-the-shelf haptic interface device. The user controls the Phantom to hold a fragile virtual object against a wall, with the goal of keeping the object from breaking or slipping. This task is simple enough that we can fully model the interactions and control the feedback available to the user, yet complex enough to preserve the most interesting parts of a real grasp and lift task: the coordination of grasp and load force and the tradeoff between risks of dropping the object and damaging the object.

In our experiments, the different feedback combinations correspond to different real-world scenarios. With only visual feedback, the situation is similar to that experienced by a user of a typical myoelectric prosthesis; the user is forced to carefully watch the interaction with the object to get any information about grip forces and object slip. The addition of haptic feedback provides extra information via sensory substitution that can be used to supplement visual information; this case corresponds to an ideal advanced prosthesis with haptic feedback. When visual feedback is turned off, the situation is similar to prosthesis use without careful visual attention; this case corresponds to real-world scenarios of a prosthesis user being unable to watch the prosthesis move or being distracted and giving less than full attention. We are particularly interested in these no-vision cases, because these are the cases in which we expect haptic feedback to be most necessary. Without vision, the case of no haptic feedback corresponds to typical prosthesis use without vision, and the cases with haptic feedback correspond to prosthesis use with an advanced system that includes extra sensory information.

Analysis of the no-vision results from these experiments showed two main findings. First, vibrotactile feedback of gross slip velocity and skin stretch feedback of incipient slip cues considerably improved performance; when participants could not see the virtual object slipping, they were able to rely on the slip cues to recover the virtual object from slips much more frequently than with force feedback alone. These results, coupled with previous results from our collaborators showing that force feedback reduced the user's likelihood to damage an object [91], clearly indicate the importance of haptic feedback for object manipulation. Second, users rated ease of use highest for conditions with more feedback types active, suggesting that incorporating all three types of haptic feedback (force feedback, vibrotactile feedback, and skin stretch feedback) into a prosthetic limb could be beneficial. These results are consistent with prior studies showing that multisensory feedback does not have any negative effects on performance in haptic control tasks [92,93].

5.1.2 Multimodality Feedback. Though there is some evidence that combining modalities of sensory feedback can result in worse performance than a using single modality [94], other studies have shown that using haptic feedback combined with visual feedback has no negative effects on performance [92,93], and multichannel sensory information feedback can surpass single modality interfaces in some cases [95]. Although the previously mentioned findings were for relatively simple tasks, using multimodal feedback has also proven effective during complex motor tasks. Indeed, incorporating haptic guidance with visual feedback during complex motor tasks has been shown in Ref. [96] to improve both motor learning and performance, while reducing the human's perceived workload. Leveraging a multimodal feedback strategy for complex tasks is also preferred by users over unimodal interactions [97] and well suited for complex tasks with high workload in one modality, so as to prevent cognitive overload and in turn enhance motor learning [98].

5.2 Conveying Environment Characteristics. Bilateral teleoperation systems enable a human operator to manipulate a remote environment, and in such shared control systems, there is typically a faithful mapping of the remote environment forces back to the master manipulator. Most master-slave systems employ kinematically similar (though sometimes scaled) robotic devices, and force sensors at the remote end effector capture interaction forces which are then mapped to the master and displayed by commanding forces via the actuators on the master robotic device. Such systems typically relay kinesthetic type forces from the remote environment.

Upper-limb prostheses are prototypical pHRI systems that enable amputees to manipulate their environment. These devices provide a context for discussion of methods to convey environment characteristics and forces that arise during manipulation with

objects and the environment in a broader sense. Despite rapid advances in mechanism design and control leading to multi-DOF robotic hands with dexterity capabilities comparable to able-bodied humans, advanced prosthetic devices still lack the touch feedback necessary for dexterous manipulation of objects [99,100]. The result of this critical absence is that many prosthesis users opt to abandon their devices [101–103]. The challenge of providing an amputee access to sensory feedback from these devices is significant. Consider the numerous types of sensory cues available to an individual interacting with an object such as a coffee cup. The person feels various haptic feedback sensations: the object weight, the grasping force, the object texture, and many other sensations [90,104–106]. These sensations allow us to manipulate objects almost effortlessly, even without visually attending to the task. The importance of such feedback should not be underestimated; people with impaired sensory feedback have difficulty interacting with objects, often dropping or damaging objects in their grasp [107–109]. Thus, if a highly articulated prosthetic limb fails to offer sensory feedback to the amputee, they must watch carefully when interacting with objects, lest they damage or drop something. Alternatives to this constant, inefficient, and cognitively taxing visual monitoring are needed to reduce the mental effort required for interactions between the prosthetic and the environment, and prosthesis users often express a desire for more sensory feedback from their devices [110].

In an effort to recreate natural touch sensations for amputees, many methods have been explored [100], ranging from invasive techniques (e.g., peripheral nerve stimulation [111,112]) to non-invasive sensory substitution methods (e.g., encoding grip force, hand/arm configuration, or other information in vibration patterns applied to a user's skin [92,113–117], stretching of the user's skin [118,119], or pressure/shear forces [87,91,94,120], see Ref. [121] for a review). Targeted re-innervation, involving the surgical re-implantation of nerve bundles to alternate muscle sites, appears to hold promise for chronic use. A muscle such as the pectoralis serves to amplify peripheral nerve signals associated with activation of the phantom limb, and transduction to electrodes takes place across the skin using myoelectric technology. Sensory feedback that is referred to the phantom limb is an added bonus, as afferent fibers implant on cutaneous sensors in the skin of the targeted area [122]. However, targeted re-innervation involves a substantial surgical intervention and seems to be indicated only for a small population of amputees [123].

An alternative approach for conveying sensory feedback to the amputee is necessary, and an approach that is gaining traction is to use sensory substitution-based haptic feedback designs optimized for translation to the amputee population. Typically, haptic devices based on sensory substitution are designed as modular devices to be worn on other parts of the user's body [100]. Our own recent work includes several wearable devices, used to provide information about grip force, gripper aperture, and object slip to the user of a prosthesis performing a manipulation task [85–89]. Many of these devices have been shown to help users manipulate objects [85,88,89,100]. Indeed, the matched modality of sensory feedback in these examples may contribute to these positive outcomes, and can be explained by the recent finding that when a prosthesis interface feeds back the mechanical response from the environment to the muscle that activated that response, then the brain seems to adopt that prosthesis as an extension of the body [124].

5.3 Conveying Desired Actions of the Human. The addition of haptic feedback to virtual environment simulations and telerobotic systems is known to provide benefits over visual-only displays, such as reduced learning times, improved task performance quality, increased dexterity, and increased feelings of realism and presence [125–131]. Haptic feedback in virtual environments also enables a wider range of applications, including manipulation and

assembly tasks where force cues are necessary, and medical applications, such as training for palpation, needle insertion, minimally invasive surgery, and rehabilitation [132]. When a virtual environment is to be used for training, it can be augmented with additional feedback that can convey the key strategies for successful task completion [133,134]. We have shown that several mechanisms for displaying fundamental movement strategies are beneficial to enhancing performance [135,136], including visual representation, verbal instructions, and haptic assistance. Other groups provide additional strong evidence to support the addition of haptic cues to assist with task completion for performance enhancement [98,137,138]. Haptic augmentation can improve performance in dynamic tasks [82], tactile cueing systems have been extensively studied to determine appropriate methods for guiding wrist rotation movements [139], and motion guidance has been effectively conveyed through tactile cues both vibrational [140] and skin stretch [141]. There is strong consensus in the literature to suggest that vibrotactile cueing is a useful tool for guiding human movements [140,142–146] though a meta-analysis of 45 studies indicated that vibrotactile cues should provide redundant information or supplement another modality such as vision [147].

In some applications of virtual environment training, it is desirable to realize both the performance gained with the addition of cues to guide successful task completion, and the ability of the human to transfer that skill to an unassisted task. Consider surgical skill training or sports applications where ultimately the individual will need to complete the task independently. In these cases, the desired strategies or methods of task completion must be acquired by the trainee during those interactions with the augmented virtual training environment, and retained when performing the task without assistance from the robotic device. Some types of kinesthetic haptic assistance, while beneficial for enhancing performance, have been ineffective when it comes to demonstrating retention of skill or transfer to a similar task [57,80,81]. In these cases, tactile feedback has the potential to be widely applied for the training of complex movements in later stages when task execution strategies need to be refined, where subjects are already familiar with the basics of completing a particular task, but lack the dexterity to do so efficiently and repeatedly. Using visual and vibrotactile feedback to teach users a desired oar trajectory was shown in Ref. [148] to result in slightly better learning than teaching with only visual feedback or only vibrotactile feedback. Studies on drawing different shapes [149] and on handwriting [150] have demonstrated an improvement in movement fluidity by the addition of haptic feedback during training.

5.4 Haptic Feedback Architectures. Some thought must be given to the architecture of the haptic communication channels, particularly in the case where the robotic system is attempting to guide the human operator to a particular task completion strategy, such as is the case with shared control haptic guidance. Imagine a surgical trainee using a virtual reality simulator to rehearse a procedure, with the simulator providing haptic cues that suggest the most appropriate trajectories that should be followed through the anatomy. The defining characteristic of robot-mediated training is that guidance is administered physically to a novice subject via a haptic interface.

In Ref. [57], we proposed a taxonomy for the various approaches that have been used to convey this duality of feedback. The taxonomy is based on three factors. The primary factor differentiating guidance paradigms is whether the additional haptic cues (beyond environment force feedback) are assisting or resisting the novice in completing the task. Next, the paradigm must resolve the co-presentation of task and guidance forces by some means of separation (temporal, spatial, or multimodal), or in the absence of such a strategy, will settle for summation of task and guidance forces, which we term gross assistance. Third, shared control haptic guidance paradigms may feature a mechanism for adjusting the amount of assistance provided by the robotic device. We summarize the taxonomy categories here:

- (1) *Type of guidance*: guidance cues are intended to assist in task performance (reducing difficulty), or the task is increased in difficulty with perturbing or resistive forces;
- (2) *Reconciliation of co-presentation of cues*: the mechanism by which the paradigm conveys task cues versus guidance cues to the novice (e.g., temporally separated, spatially separated, or multiple modalities of feedback); and
- (3) *Progressive guidance*: reducing the amount of guidance based on some factor (e.g., time or performance).

In our prior work, we have exposed difficulties in successful implementation of gross assistance since guidance cues can interfere with task performance [57,151]. When different haptic modalities are employed, such as kinesthetic cues for task forces and tactile cues for guidance forces, these guidance cues have been shown to enhance dexterous performance of tasks [89]. This is an example of spatial separation of haptic guidance and task-inherent forces, which also has been demonstrated with kinesthetic-only devices [57,152]. Temporal separation of haptic guidance [57,153] can be realized by using terminal feedback after completed trials, and is suggested to reduce user reliance on the feedback [80,138], as per the guidance hypothesis [83,84]. Across all of these applications, it is vital to properly match a guidance paradigm to a task's dynamic characteristics in order to achieve high efficacy and low cognitive workload for the trainee [57].

6 Case Studies

In Secs. 3–5, we have separately focused on intent detection, arbitration, and communication within shared control, defining each formally and briefly providing examples taken from the state of the art. In this section, we will emphasize two case studies of shared control within the intent detection, arbitration, and communication framework, while providing more detail on how our research groups incorporated and implemented these three fundamental aspects of shared control. Speaking in general, and recalling the concept introduced in Sec. 2 as well as the scheme presented in Fig. 1, arbitration can be seen as an overseer that decides how to distribute control between human and robot, while intent and feedback are the means through which the human and the robot exchange information. From this point of view, intent detection and communication feedback can be seen as being symmetric entities, with the important difference that feedback is generally used as a way to convey information to the human, while intent can be directly used as part of the arbitration process. We argue that the following case studies exemplify these concepts, and illustrate our intent detection, arbitration, and communication framework.

6.1 Case Study: Rehabilitation Robotics With the Minimal Assist-as-Needed Controller. Within the mechatronics and haptic interfaces lab at Rice University, our research group has applied concepts from shared control to upper-limb rehabilitation robotics. We consider situations where a human and robot are working together to successfully complete repetitive motions; the robot cooperates with the human by correcting the human's actions whenever necessary. In particular, our minimal assist-as-needed (mAAN) controller, recently presented in Pehlivan et al. [28], employs specific intent detection, arbitration, and communication methods to better help subjects recover from neurological injuries. The following discussion highlights the intent detection, arbitration, and communication aspects of this mAAN controller, and demonstrates their underlying influence on our shared control application.

6.1.1 The Minimal Assist-as-Needed Controller. Robots provide an effective means for rehabilitation after stroke because they can cooperate with humans to provide consistent, repetitive, and intense therapeutic interactions. Indeed, a 127 person clinical

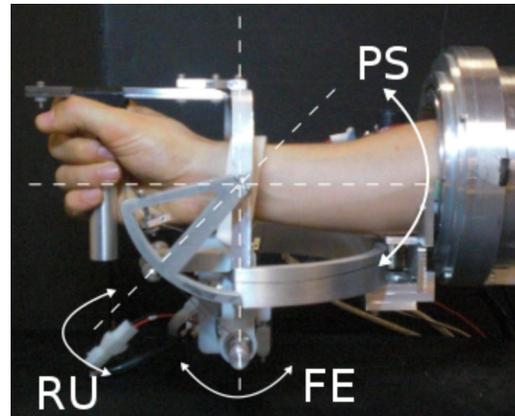


Fig. 7 RiceWrist-S wrist-forearm exoskeleton, with labeled joints for pronation/supination (PS), flexion/extension (FE), and radial/ulnar deviation (RU). The human and robot share control of the handle position during trajectory following tasks, with applications in upper-limb rehabilitation (Reproduced with permission from Pehlivan et al. [28]. Copyright 2016 by IEEE).

trial conducted over a 6 month period found that there was no significant difference in the improvements of motor function for stroke subjects trained by robots as compared to those receiving traditional human-mediated therapy [154]. To further improve the performance of rehabilitation robots, researchers have developed control strategies, which focus on promoting subject engagement [56]. The problem of maximizing human engagement—i.e., encouraging the human to complete as much of the motion as possible—can also be posed as minimizing robotic assistance—i.e., helping the human as little as the task requires. We must remember, however, that the ability of neurologically impaired individuals to perform these desired motions varies in a highly nonlinear and unpredictable fashion, in part due to the effects of movement disorders [155]. Hence, within our research, we argue that the robot can only provide minimal assistance to the human if the human's ability is measured in real time (intent detection). Using this estimate of the human's current ability, the mAAN controller then offers therapists a simple and intuitive way to tune the amount of trajectory error which the subject is allowed (arbitration). Finally, we experimentally realize this controller on the RiceWrist-S [156], a 3DoF wrist-forearm exoskeleton, which renders desired haptic environments to the subject (communication). An image of the RiceWrist-S can be seen in Fig. 7.

6.1.2 Intent Detection: Sensorless Force Estimation. Although the subject's capability could be measured using force/torque sensors on the exoskeleton's handle, we instead elected to utilize sensorless force estimation so as to reduce the overall system cost. On the one hand, we can use our dynamic model of the robot, as well as the known controller torques, to predict the robot's expected state (joint position and velocity) when given an estimate of the human's force input. On the other hand, we have rotary encoders to measure the actual position of each joint, and we can take the derivative of these measurements to estimate the actual joint velocities. We combined these two sources of information—predicted and actual—using a KF. By then applying Lyapunov stability analysis in conjunction with the KF, we were able to find adaption laws that, when integrated, provided an estimate of the human's current force input. Moreover, these adaption laws ensured that the system was stable in the sense of Lyapunov, and therefore both the estimated system states and the estimated human inputs desirably converged to the moving average of their actual values.

From our perspective, the most interesting aspect of this intent detection method was that it did not require the subject's abilities to consistently follow some underlying pattern. Prior work by

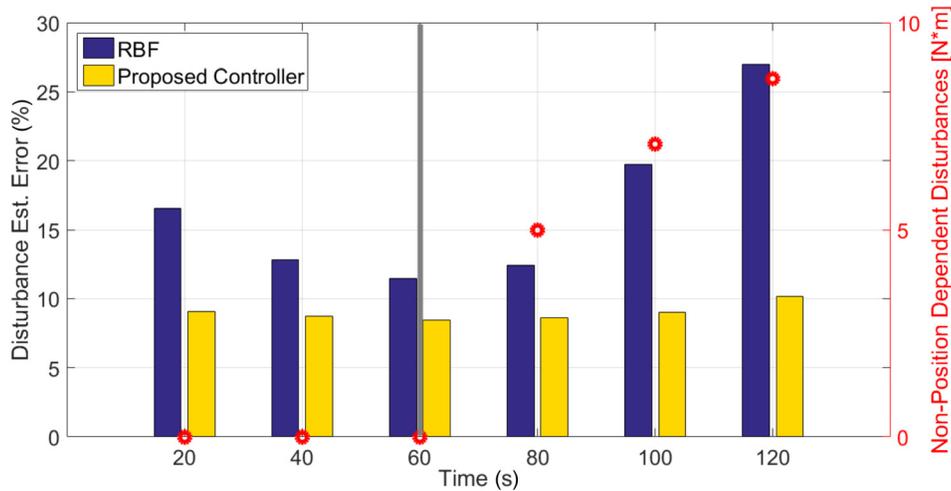


Fig. 8 Comparison of pre-existing RBF intent detection scheme (dark bar) with our proposed KF intent detection approach (light bar). The normalized error between the actual and estimated human intents for both methods is plotted over 20 s intervals (disturbance estimation error). After 60 s, nonposition-dependent human inputs of increasing magnitude were applied (circles) (Reproduced with permission from Pehlivan et al. [28]. Copyright 2016 by IEEE).

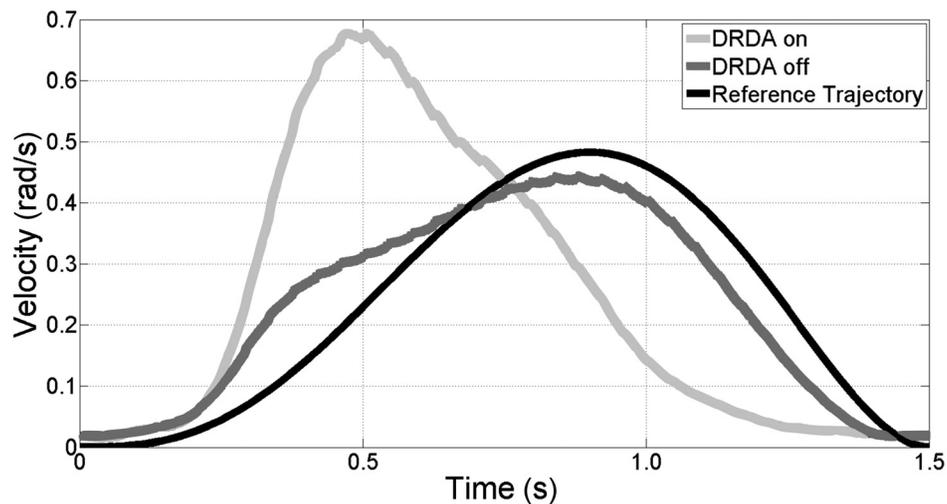


Fig. 9 Average trajectory velocities with our decay algorithm (left) and without our decay algorithm (middle). When the decay algorithm was present, arbitration is dynamically shifted toward able subjects, and users were allowed to reach the goal more quickly than the reference trajectory (right) (Reproduced with permission from Pehlivan et al. [28]. Copyright 2016 by IEEE).

Wolbrecht et al. [58] had employed Gaussian RBFs to learn the human’s force inputs as a function of position; each location in the shared workspace was thus associated with a certain intent. When comparing our KF approach to this state-of-the-art RBF method, we found that the described KF intent detection scheme was faster and more accurate (see Fig. 8). More precisely, during cases where the human’s input was nonposition dependent—as might occur with neurologically impaired individuals [155]—the performance of RBF intent detection degrades, while our time-dependent KF approach provides relatively constant performance.

6.1.3 Arbitration: Subject-Adaptive Algorithms. Like other applications of shared control for rehabilitation robotics, our mAAN controller employs a teacher–student role arbitration. The robot ensures that the human follows the desired trajectory by incorporating an impedance controller together with a “disturbance” rejection term, where this disturbance is just the human force input as estimated by our KF intent detection

approach. Applying Lyapunov stability analysis again, we discovered that the human’s trajectory tracking error is uniformly ultimately bounded when using the mAAN control law, and the bounds on the human’s allowable error can be manipulated by changing the gain of the impedance controller. Hence, we achieve a teacher–student role arbitration by decreasing the amount of allowable error when unskilled subjects are attempting to learn the desired motion, and increasing the amount of allowable error when these subjects demonstrate proficiency.

This arbitration of shared control between human and robot is dynamically changed by two subject-adaptive algorithms: an error bound modification algorithm and a decay algorithm. As previously discussed, the error bound modification algorithm updates the gain of the impedance controller between trials based on the human’s total trajectory error during the previous trial. Next, the decay algorithm permits able subjects who are capable of reaching the goal faster than the given trajectory to exceed, or surpass, that reference trajectory (as shown in Fig. 9). Viewed together, these

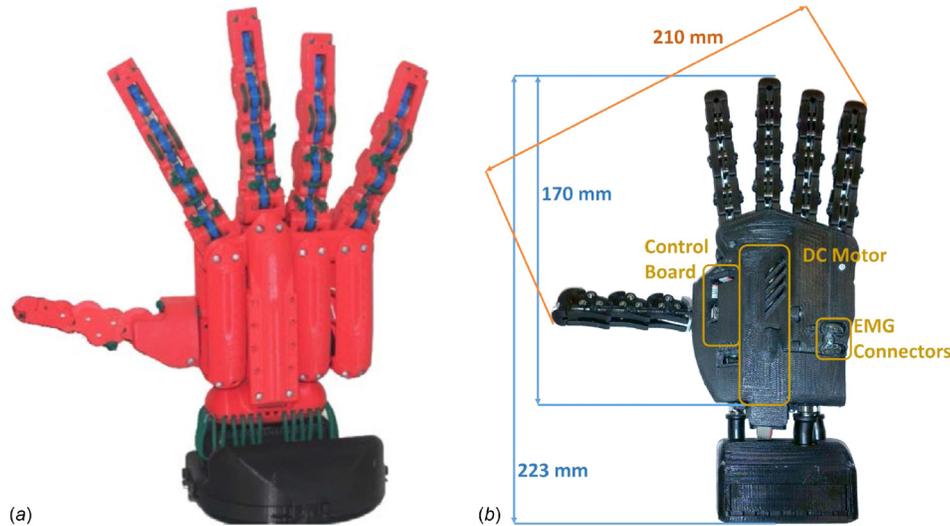


Fig. 10 First implementation of the Pisa/IIT SoftHand (left—Reproduced with permission from Catalano et al. [157]. Copyright 2014 by Sage Ltd.) and a recent release of the SoftHand Pro (right—Reproduced with permission from Fani et al. [162]. Copyright 2016 by authors, including Antonio Bicchi and Marco Santello.)

algorithms shift arbitration toward the human when the human can perform the desired motions, but shift arbitration toward the robot for less capable subjects. Importantly, our teacher–student role arbitration ensures that error is always present when the human is not actively involved in the task, thereby encouraging subject involvement.

6.1.4 Communication: Feedback With the RiceWrist-S. Shared control was communicated between the human and robot using both haptic and visual feedback. Haptic feedback consisted of the assistive torques output by the mAAN controller, which was then perceived by the human grasping our RiceWrist-S exoskeleton. In order to more accurately convey the desired virtual environment to the human, the RiceWrist-S reduced mechanical friction and backlash by means of cable drive transmissions; the device was also designed to have a low apparent inertia, sufficient workspace, and adequate torque outputs for wrist rehabilitation [156]. Visual feedback was provided in real time on a desktop monitor, which the subject viewed while performing their trial motions. This visual interface showed the current position of the coupled human and robot, as well as the desired trajectory and goal positions. Accordingly, by combining haptic and visual feedback, the human simultaneously “felt” corrective forces and observed their current error; without this communication, the robot would have been unable to assist the human, and the human would have been unaware of their mistakes.

6.2 Case Study: From the Pisa/IIT SoftHand to the SoftHand Pro. Another example application of shared control in physical HRI is the transition of the Pisa/IIT SoftHand (SH) from a robotic manipulator to the prosthetic hand SoftHand Pro (SH-P). While the design process and evolution of the hand through the years was not explicitly developed in the intent-arbitration-communication approach proposed in this paper, it is possible to cast it within that framework. In this section, we will describe the use of this hand as an upper-limb prosthesis, and show how each element fits into our shared control approach.

6.2.1 The Pisa/IIT SoftHand. Inspired by mechanisms naturally found in human hands, the Pisa/IIT SoftHand design [157] was developed as an intuitive, easy-to-use robotic hand for manipulation tasks. In particular, we know from neuroscientific studies that humans control their own hands—not by independently actuating each degree-of-freedom—but rather by leveraging

coordinated co-activations, which can be referred to as synergies [158,159]. Previous work [160] has introduced an interaction strategy based on “soft” synergies: in this framework, postural synergies serve as a reference posture for the hand, and the hand’s actual posture when grasping an object is determined by both (a) the shape of the object and (b) the stiffness of the grasp.

Based on this research, the Pisa/IIT SoftHand (Fig. 10) was designed to achieve a single synergy with human-like grasp stiffness. In particular, the Pisa/IIT SoftHand is a 19DoF artificial hand, where the thumb has 3DoF, and each other finger has 4DoF. Revolute joints and rolling contact joints with elastic ligaments are used so that the robotic hand’s motions are both physiologically accurate and compliant. All of the joints are connected with a single tendon; by actuating this tendon, the fingers flex and adduct along the path of a single predefined human grasp synergy from Ref. [158]. Using this synergy reduces the mechanical complexity of the device—only one DC motor is needed to actuate the entire 19DoF hand. Moreover, since the fingers are compliant, the Pisa/IIT SoftHand has a flexible grasp pattern, which allows the robot to intuitively manipulate many different objects. Under the current iteration, the hand can apply a maximum force of 130 N perpendicular to its palm. To see more—including a model of the Pisa/IIT SoftHand as well as the design of its accompanying electronic board—visit the Natural Machine Motion Initiative Website.¹

6.2.2 Intent Detection and Arbitration: Myoelectric Control for Prosthetic Use. The intelligence of this hand is embedded in its hardware since complex manipulation is achieved by deforming and adapting the primitive synergistic shape, instead of using a complex control strategy. This greatly simplifies the intent detection and arbitration components, since a simple opening-closure signal is all that is necessary to control the hand. For this reason, adaptation of the hand for prosthetic use was relatively straightforward, at least in as far as the control strategy is concerned.

Electromyography control was obtained with two electrodes placed on the proximal forearm [161], and two different myoelectric controllers were tested: a standard controller in which the EMG signal is used only as a position reference, and an impedance controller that determines both position and stiffness references from the EMG input. Grasp performance was similar under the two control modes, but in questionnaires, subjects reported

¹<http://www.naturalmachinemotioninitiative.net/>

that impedance control was easier to use. This result was confirmed by lower muscles activation measured by the EMG sensors.

Although EMG seems a promising approach for detecting user intent, it is not necessarily clear how that intent should be used to control the human's prostheses. Accordingly, a comparison of three possible myoelectric control strategies is described by Fani et al. [162]. Within this experiment, subjects performed reach-to-grasp movements using their native hand, while their EMG signals were recorded. These EMG signals were simultaneously used as inputs to the SoftHand Pro, which was controlled by one of the following three strategies: differential control, first-come-first-served (FCFS), or FCFS-Advanced. For the differential controller, the difference between the EMG signals was used to control the robot; for the FCFS controller, the first EMG signal which exceeded a threshold was used to control the robot; for the FCFS-Advanced controller, an additional requirement was added to prevent involuntary direction changes. Assessments were used to determine how well the SoftHand Pro mimicked the native human hand using each controller type. Based on the results in Ref. [162], differential control leads to the most natural behavior, and appears to be a promising method of arbitration for myoelectric prostheses.

6.2.3 Communication: Haptic Feedback Devices. Despite their high level of technology, myoelectric prostheses do not have a high level of acceptance among users, and there is indication that a reason for this is the inherent lack of sensory feedback in EMG controlled prostheses [163]. Apart from being helpful during grasp and manipulation of objects, haptic feedback has potential for increasing embodiment of the prosthetic hand. For these reasons, parallel work regarded design of haptic devices to be integrated with the Pisa/IIT SoftHand, to convey feedback to the user, together with investigation of the effectiveness of such devices in improving grasp quality. Two different approaches have been presented so far.

Communication to the human operator has been explored in a couple of ways. First, the clenching upper-limb force feedback (CUFF) was introduced [164], which is a device consisting of two motors commanded in position that are used to control a piece of fabric so that both normal and tangential forces can be applied on the user's arm (the device is shown in Fig. 11). In Ref. [164], it was shown that the device can be used successfully to delivering information on the grasp force of the SoftHand, which also helps the user discriminate softness. In the later work [165], effect of

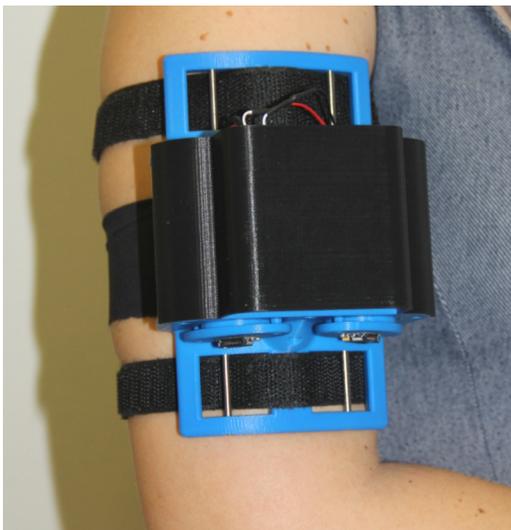


Fig. 11 Clenching upper-limb force feedback haptic device (Reproduced with permission from Casini et al. [164]. Copyright 2015 by IEEE.)

haptic feedback from the CUFF on grip force modulation was investigated: results suggested an overall reduction of grasp force when using the CUFF compared to use of the SoftHand Pro alone, which, however, did not reach the effect of statistical significance.

Another approach is vibrotactile feedback, which was studied in both Refs. [161] and [166]. In particular, Godfrey et al. [161] report that for healthy subjects, the perceived cognitive load from using the SoftHand is lower when using vibrotactile feedback, while in Ref. [166], subjects using vibrotactile feedback were shown to be able to correctly discriminate between textures. Interestingly, Ajoudani et al. [166] also proposed a mechano-tactile feedback device similar to the CUFF. During their experiments, this device was found to effectively convey meaningful texture information by itself, as well as to improve blind grasping.

7 Conclusion

We have presented a unified view of shared control systems that feature pHRI, with a focus on applications in healthcare that critically depend on collaborative and cooperative actions between man and machine. Three key themes were explored in the review. First, we surveyed methods of intent detection, focused on defining, measuring, and interpreting the human's intent in the shared control scenario. Second, we reviewed arbitration, the modulation of control authority over a task between human and robot. Here, we defined types of arbitration, and the more advanced topic of dynamically changing role arbitration in shared control tasks. Finally, we discussed the role of feedback in shared control scenarios. We presented a survey of types of feedback (with an emphasis on the haptic sensory channels which are prominent in pHRI), and examples of feedback both of environment characteristics and methods of guiding or instructing the human operator to execute desired task completion strategies or trajectories. We described taxonomy for the approaches that are typically used for providing feedback in shared control systems. The paper also provides an illustration of shared control in pHRI through two case studies (rehabilitation robotics and robotic prosthetics). In these case studies, we elucidate the realization of intent detection, arbitration, and feedback in such prototypical applications.

Funding Data

- Directorate for Computer and Information Science and Engineering (Grant No. NSF IIS-1065497).
- National Science Foundation (Grant No. NSF DGE-1250104).
- Seventh Framework Programme (Grant No. 601165).

References

- [1] Goodrich, M. A., and Schultz, A. C., 2007, "Human-Robot Interaction: A Survey," *Found. Trends Hum.-Comput. Interact.*, 1(3), pp. 203–275.
- [2] Heyer, C., 2010, "Human-Robot Interaction and Future Industrial Robotics Applications," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct. 18–22, pp. 4749–4754.
- [3] Peshkin, M. A., Colgate, J. E., Wannasupphrasit, W., Moore, C. A., Gillespie, R. B., and Akella, P., 2001, "Cobot Architecture," *IEEE Trans. Rob. Autom.*, 17(4), pp. 377–390.
- [4] Fong, T., Nourbakhsh, I., and Dautenhahn, K., 2003, "A Survey of Socially Interactive Robots," *Rob. Auton. Syst.*, 42(3), pp. 143–166.
- [5] Okamura, A. M., 2009, "Haptic Feedback in Robot-Assisted Minimally Invasive Surgery," *Curr. Opin. Urol.*, 19(1), p. 102.
- [6] Pons, J. L., 2008, *Wearable Robots: Biomechanical Exoskeletons*, Wiley, Hoboken, NJ.
- [7] Jamwal, P. K., Xie, S. Q., Hussain, S., and Parsons, J. G., 2014, "An Adaptive Wearable Parallel Robot for the Treatment of Ankle Injuries," *IEEE/ASME Trans. Mechatronics*, 19(1), pp. 64–75.
- [8] Haddadin, S., Suppa, M., Fuchs, S., Bodenmüller, T., Albu-Schäffer, A., and Hirzinger, G., 2011, "Towards the Robotic Co-Worker," *Robotics Research*, Springer, Berlin, pp. 261–282.
- [9] De Santis, A., Siciliano, B., De Luca, A., and Bicchi, A., 2008, "An Atlas of Physical Human–Robot Interaction," *Mech. Mach. Theory*, 43(3), pp. 253–270.
- [10] Mörtl, A., Lawitzky, M., Kucukylmaz, A., Sezgin, M., Basdogan, C., and Hirche, S., 2012, "The Role of Roles: Physical Cooperation Between Humans and Robots," *Int. J. Rob. Res.*, 31(13), pp. 1656–1674.

- [11] Engeberg, E. D., and Meek, S. G., 2013, "Adaptive Sliding Mode Control for Prosthetic Hands to Simultaneously Prevent Slip and Minimize Deformation of Grasped Objects," *IEEE/ASME Trans. Mechatronics*, **18**(1), pp. 376–385.
- [12] Wang, Z., Peer, A., and Buss, M., 2009, "An HMM Approach to Realistic Haptic Human-Robot Interaction," Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (*World Haptics 2009*), Salt Lake City, UT, Mar. 18–20, pp. 374–379.
- [13] McMullen, D. P., Hotson, G., Kataly, K. D., Wester, B. A., Fifer, M. S., McGee, T. G., Harris, A., Johannes, M. S., Vogelstein, R. J., Ravitz, A. D., Anderson, W. S., Thakor, N. V., and Crone, N. E., 2014, "Demonstration of a Semi-Autonomous Hybrid Brain-Machine Interface Using Human Intracranial EEG, Eye Tracking, and Computer Vision to Control a Robotic Upper Limb Prosthetic," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **22**(4), pp. 784–796.
- [14] Radmand, A., Scheme, E., and Englehart, K., 2014, "A Characterization of the Effect of Limb Position on EMG Features to Guide the Development of Effective Prosthetic Control Schemes," 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (*EMBC*), Chicago, IL, Aug. 26–30, pp. 662–667.
- [15] Yap, H. K., Mao, A., Goh, J. C., and Yeow, C.-H., 2016, "Design of a Wearable FMG Sensing System for User Intent Detection During Hand Rehabilitation With a Soft Robotic Glove," Sixth IEEE International Conference on Biomedical Robotics and Biomechanics (*BioRob*), Singapore, June 26–29, pp. 781–786.
- [16] Cho, E., Chen, R., Merhi, L.-K., Xiao, Z., Pousett, B., and Menon, C., 2016, "Force Myography to Control Robotic Upper Extremity Prostheses: A Feasibility Study," *Front. Bioeng. Biotechnol.*, **4**, p. 18.
- [17] Rasouli, M., Chellamuthu, K., Cabibihan, J.-J., and Kukreja, S. L., 2016, "Towards Enhanced Control of Upper Prosthetic Limbs: A Force-Myographic Approach," Sixth IEEE International Conference on Biomedical Robotics and Biomechanics (*BioRob*), Singapore, June 26–29, pp. 232–236.
- [18] Shen, B., Li, J., Bai, F., and Chew, C.-M., 2013, "Motion Intent Recognition for Control of a Lower Extremity Assistive Device (LEAD)," IEEE International Conference on Mechatronics and Automation (*ICMA*), Takamatsu, Japan, Aug. 4–7, pp. 926–931.
- [19] Yan, T., Cempini, M., Oddo, C. M., and Vitiello, N., 2015, "Review of Assistive Strategies in Powered Lower-Limb Orthoses and Exoskeletons," *Rob. Auton. Syst.*, **64**, pp. 120–136.
- [20] Sarac, M., Koyas, E., Erdogan, A., Cetin, M., and Patoglu, V., 2013, "Brain Computer Interface Based Robotic Rehabilitation With Online Modification of Task Speed," IEEE International Conference on Rehabilitation Robotics (*ICORR*), Seattle, WA, June 24–26, pp. 1–7.
- [21] Huang, C., Wasson, G., Alwan, M., Sheth, P., and Ledoux, A., 2005, "Shared Navigational Control and User Intent Detection in an Intelligent Walker," AAAI Fall 2005 Symposium, Arlington, VA, Nov. 4–6.
- [22] Wakita, K., Huang, J., Di, P., Sekiyama, K., and Fukuda, T., 2013, "Human-Walking-Intention-Based Motion Control of an Omnidirectional-Type Cane Robot," *IEEE/ASME Trans. Mechatronics*, **18**(1), pp. 285–296.
- [23] Brescianini, D., Jung, J.-Y., Jang, I.-H., Park, H. S., and Rienen, R., 2011, "INS/EKF-Based Stride Length, Height and Direction Intent Detection for Walking Assistance Robots," IEEE International Conference on Rehabilitation Robotics (*ICORR*), Zurich, Switzerland, June 29–July 1, pp. 1–5.
- [24] Corteville, B., Aertbelien, E., Bruyninckx, H., De Schutter, J., and Van Brussel, H., 2007, "Human-Inspired Robot Assistant for Fast Point-to-Point Movements," IEEE International Conference on Robotics and Automation (*ICRA*), Rome, Italy, Apr. 10–14, pp. 3639–3644.
- [25] Ge, S. S., Li, Y., and He, H., 2011, "Neural-Network-Based Human Intention Estimation for Physical Human-Robot Interaction," Eighth International Conference on Ubiquitous Robots and Ambient Intelligence (*URAI*), Incheon, South Korea, Nov. 23–26, pp. 390–395.
- [26] Li, Y., and Ge, S. S., 2014, "Human-Robot Collaboration Based on Motion Intention Estimation," *IEEE/ASME Trans. Mechatronics*, **19**(3), pp. 1007–1014.
- [27] Erden, M. S., and Tomiyama, T., 2010, "Human-Intent Detection and Physically Interactive Control of a Robot Without Force Sensors," *IEEE Trans. Rob.*, **26**(2), pp. 370–382.
- [28] Pehlivan, A. U., Losey, D. P., and O'Malley, M. K., 2016, "Minimal Assist-as-Needed Controller for Upper Limb Robotic Rehabilitation," *IEEE Trans. Rob.*, **32**(1), pp. 113–124.
- [29] Lenzi, T., De Rossi, S. M. M., Vitiello, N., and Carrozza, M. C., 2012, "Intention-Based EMG Control for Powered Exoskeletons," *IEEE Trans. Biomed. Eng.*, **59**(8), pp. 2180–2190.
- [30] Kiguchi, K., and Hayashi, Y., 2012, "An EMG-Based Control for an Upper-Limb Power-Assist Exoskeleton Robot," *IEEE Trans. Syst., Man, Cybern., Part B*, **42**(4), pp. 1064–1071.
- [31] Hudgins, B., Parker, P., and Scott, R. N., 1993, "A New Strategy for Multifunction Myoelectric Control," *IEEE Trans. Biomed. Eng.*, **40**(1), pp. 82–94.
- [32] Au, S., Berniker, M., and Herr, H., 2008, "Powered Ankle-Foot Prosthesis to Assist Level-Ground and Stair-Descent Gaits," *Neural Networks*, **21**(4), pp. 654–666.
- [33] Song, R., Tong, K.-Y., Hu, X., and Li, L., 2008, "Assistive Control System Using Continuous Myoelectric Signal in Robot-Aided Arm Training for Patients After Stroke," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **16**(4), pp. 371–379.
- [34] Akhlaghi, N., Baker, C., Lahlou, M., Zafar, H., Murthy, K., Rangwala, H., Kosecka, J., Joiner, W., Pancrazio, J., and Sikdar, S., 2016, "Real-Time Classification of Hand Motions Using Ultrasound Imaging of Forearm Muscles," *IEEE Trans. Biomed. Eng.*, **63**(8), pp. 1687–1698.
- [35] Pratt, J. E., Krupp, B. T., Morse, C. J., and Collins, S. H., 2004, "The Roboknee: An Exoskeleton for Enhancing Strength and Endurance During Walking," IEEE International Conference on Robotics and Automation (*ICRA '04*), New Orleans, LA, Apr. 26–May 1, pp. 2430–2435.
- [36] Burdet, E., and Milner, T. E., 1998, "Quantization of Human Motions and Learning of Accurate Movements," *Biol. Cybern.*, **78**(4), pp. 307–318.
- [37] Bratman, M. E., 1992, "Shared Cooperative Activity," *Philos. Rev.*, **101**(2), pp. 327–341.
- [38] Reed, K. B., and Peshkin, M. A., 2008, "Physical Collaboration of Human-Human and Human-Robot Teams," *IEEE Trans. Haptics*, **1**(2), pp. 108–120.
- [39] Ueha, R., Pham, H. T., Hirai, H., and Miyazaki, F., 2009, "A Simple Control Design for Human-Robot Coordination Based on the Knowledge of Dynamical Role Division," IEEE/RSJ International Conference on Intelligent Robots and Systems (*IROS*), St. Louis, MO, Oct. 10–15, pp. 3051–3056.
- [40] Feth, D., Groten, R., Peer, A., Hirche, S., and Buss, M., 2009, "Performance Related Energy Exchange in Haptic Human-Human Interaction in a Shared Virtual Object Manipulation Task," IEEE Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (*World Haptics 2009*), Salt Lake City, UT, Mar. 18–20, pp. 338–343.
- [41] Jarrassé, N., Charalambous, T., and Burdet, E., 2012, "A Framework to Describe, Analyze and Generate Interactive Motor Behaviors," *PLoS One*, **7**(11), p. e49945.
- [42] Varol, H. A., Yap, F., and Goldfarb, M., 2010, "Multiclass Real-Time Intent Recognition of a Powered Lower Limb Prosthesis," *IEEE Trans. Biomed. Eng.*, **57**(3), pp. 542–551.
- [43] Englehart, K., and Hudgins, B., 2003, "A Robust, Real-Time Control Scheme for Multifunction Myoelectric Control," *IEEE Trans. Biomed. Eng.*, **50**(7), pp. 848–854.
- [44] Huang, Y., Englehart, K. B., Hudgins, B., and Chan, A. D., 2005, "A Gaussian Mixture Model Based Classification Scheme for Myoelectric Control of Powered Upper Limb Prostheses," *IEEE Trans. Biomed. Eng.*, **52**(11), pp. 1801–1811.
- [45] Chadwell, A., Kenney, L., Thies, S., Galpin, A., and Head, J., 2016, "The Reality of Myoelectric Prostheses: Understanding What Makes These Devices Difficult for Some Users to Control," *Front. Neurobotics*, **10**, p. 7.
- [46] Phillips, J., Millán, J. D. R., Vanacker, G., Lew, E., Galán, F., Ferrez, P. W., Van Brussel, H., and Nuttin, M., 2007, "Adaptive Shared Control of a Brain-Actuated Simulated Wheelchair," IEEE Tenth International Conference on Rehabilitation Robotics (*ICORR*), Noordwijk, The Netherlands, June 13–15, pp. 408–414.
- [47] Carlson, T., and Millan, J. d. R., 2013, "Brain-Controlled Wheelchairs: A Robotic Architecture," *IEEE Rob. Autom. Mag.*, **20**(1), pp. 65–73.
- [48] Rebsamen, B., Guan, C., Zhang, H., Wang, C., Teo, C., Ang, M. H., and Burdet, E., 2010, "A Brain Controlled Wheelchair to Navigate in Familiar Environments," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **18**(6), pp. 590–598.
- [49] Aarno, D., Ekvall, S., and Kragic, D., 2005, "Adaptive Virtual Fixtures for Machine-Assisted Teleoperation Tasks," IEEE International Conference on Robotics and Automation (*ICRA*), Barcelona, Spain, Apr. 18–22, pp. 1139–1144.
- [50] Abbott, J. J., Marayong, P., and Okamura, A. M., 2007, "Haptic Virtual Fixtures for Robot-Assisted Manipulation," *Robotics Research*, Springer, Berlin, pp. 49–64.
- [51] Bowyer, S. A., Davies, B. L., and Baena, F. R., 2014, "Active Constraints/Virtual Fixtures: A Survey," *IEEE Trans. Rob.*, **30**(1), pp. 138–157.
- [52] Yu, W., Alqasemi, R., Dubey, R., and Peralte, N., 2005, "Telemat Manipulation Assistance Based on Motion Intention Recognition," IEEE International Conference on Robotics and Automation (*ICRA*), Barcelona, Spain, Apr. 18–22, pp. 1121–1126.
- [53] Li, M., and Okamura, A. M., 2003, "Recognition of Operator Motions for Real-Time Assistance Using Virtual Fixtures," 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (*HAPTICS 2003*), Los Angeles, CA, Mar. 22–23, pp. 125–131.
- [54] Mao, Y., and Agrawal, S. K., 2012, "Design of a Cable-Driven Arm Exoskeleton (CAREX) for Neural Rehabilitation," *IEEE Trans. Rob.*, **28**(4), pp. 922–931.
- [55] Duschau-Wicke, A., von Zitzewitz, J., Caprez, A., Lunenburger, L., and Rienen, R., 2010, "Path Control: A Method for Patient-Cooperative Robot-Aided Gait Rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **18**(1), pp. 38–48.
- [56] Blank, A. A., French, J. A., Pehlivan, A. U., and O'Malley, M. K., 2014, "Current Trends in Robot-Assisted Upper-Limb Stroke Rehabilitation: Promoting Patient Engagement in Therapy," *Curr. Phys. Med. Rehabil. Rep.*, **2**(3), pp. 184–195.
- [57] Powell, D., and O'Malley, M. K., 2012, "The Task-Dependent Efficacy of Shared-Control Haptic Guidance Paradigms," *IEEE Trans. Haptics*, **5**(3), pp. 208–219.
- [58] Wolbrecht, E. T., Chan, V., Reinkensmeyer, D. J., and Bobrow, J. E., 2008, "Optimizing Compliant, Model-Based Robotic Assistance to Promote Neurorehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **16**(3), pp. 286–297.
- [59] Emken, J. L., Benitez, R., Sideris, A., Bobrow, J. E., and Reinkensmeyer, D. J., 2007, "Motor Adaptation as a Greedy Optimization of Error and Effort," *J. Neurophysiol.*, **97**(6), pp. 3997–4006.
- [60] Emken, J. L., Bobrow, J. E., and Reinkensmeyer, D. J., 2005, "Robotic Movement Training as an Optimization Problem: Designing a Controller That Assists Only as Needed," IEEE Ninth International Conference on Rehabilitation Robotics (*ICORR 2005*), Chicago, IL, June 28–July 1, pp. 307–312.
- [61] Rauter, G., Sigrist, R., Marchal-Crespo, L., Vallery, H., Rienen, R., and Wolf, P., 2011, "Assistance or Challenge? Filling a Gap in User-Cooperative Control," IEEE/RSJ International Conference on Intelligent Robots and Systems (*IROS*), San Francisco, CA, Sept. 25–30, pp. 3068–3073.
- [62] Evrard, P., and Kheddar, A., 2009, "Homotopy Switching Model for Dyad Haptic Interaction in Physical Collaborative Tasks," IEEE Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (*World Haptics 2009*), Salt Lake City, UT, Mar. 18–20, pp. 45–50.

- [63] Medina, J. R., Lorenz, T., and Hirche, S., 2015, "Synthesizing Anticipatory Haptic Assistance Considering Human Behavior Uncertainty," *IEEE Trans. Rob.*, **31**(1), pp. 180–190.
- [64] Li, Y., Tee, K. P., Chan, W. L., Yan, R., Chua, Y., and Limbu, D. K., 2015, "Continuous Role Adaptation for Human–Robot Shared Control," *IEEE Trans. Rob.*, **31**(3), pp. 672–681.
- [65] Thobbi, A., Gu, Y., and Sheng, W., 2011, "Using Human Motion Estimation for Human-Robot Cooperative Manipulation," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, San Francisco, CA, Sept. 25–30, pp. 2873–2878.
- [66] Dragan, A. D., and Srinivasa, S. S., 2013, "A Policy-Blending Formalism for Shared Control," *Int. J. Rob. Res.*, **32**(7), pp. 790–805.
- [67] Kucukyilmaz, A., Sezgin, T. M., and Basdogan, C., 2013, "Intention Recognition for Dynamic Role Exchange in Haptic Collaboration," *IEEE Trans. Haptics*, **6**(1), pp. 58–68.
- [68] Kulic, D., and Croft, E. A., 2007, "Affective State Estimation for Human–Robot Interaction," *IEEE Trans. Rob.*, **23**(5), pp. 991–1000.
- [69] Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y., De Visser, E. J., and Parasuraman, R., 2011, "A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction," *Hum. Factors*, **53**(5), pp. 517–527.
- [70] Losey, D. P., and O'Malley, M. K., 2018, "Trajectory Deformations From Physical Human–Robot Interaction," *IEEE Trans. Rob.*, **34**(1), pp. 126–138.
- [71] Bajcsy, A., Losey, D. P., O'Malley, M. K., and Dragan, A. D., 2017, "Learning Robot Objectives From Physical Human Interaction," *Conference on Robot Learning (CoRL)*, Mountain View, CA, Nov. 13–15, pp. 217–226.
- [72] Rauter, G., von Zitzewitz, J., Duschau-Wicke, A., Vallery, H., and Riener, R., 2010, "A Tendon-Based Parallel Robot Applied to Motor Learning in Sports," *Third IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob)*, Tokyo, Japan, Sept. 26–29, pp. 82–87.
- [73] von Zitzewitz, J., Wolf, P., Novaković, V., Wellner, M., Rauter, G., Brunschweiler, A., and Riener, R., 2008, "Real-Time Rowing Simulator With Multimodal Feedback," *Sports Technol.*, **1**(6), pp. 257–266.
- [74] Marchal-Crespo, L., van Raaij, M., Rauter, G., Wolf, P., and Riener, R., 2013, "The Effect of Haptic Guidance and Visual Feedback on Learning a Complex Tennis Task," *Exp. Brain Res.*, **231**(3), pp. 277–291.
- [75] Volpe, B., Krebs, H., Hogan, N., Edelman, L., Diels, C., and Aisen, M., 2000, "A Novel Approach to Stroke Rehabilitation Robot-Aided Sensorimotor Stimulation," *Neurology*, **54**(10), pp. 1938–1944.
- [76] Gupta, A., and O'Malley, M. K., 2006, "Design of a Haptic Arm Exoskeleton for Training and Rehabilitation," *IEEE/ASME Trans. Mechatronics*, **11**(3), pp. 280–289.
- [77] Veneman, J. F., Kruidhof, R., Hekman, E. E., Ekkelenkamp, R., Van Asseldonk, E. H., and Van Der Kooij, H., 2007, "Design and Evaluation of the Lopes Exoskeleton Robot for Interactive Gait Rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **15**(3), pp. 379–386.
- [78] Banala, S. K., Agrawal, S. K., and Scholz, J. P., 2007, "Active Leg Exoskeleton (ALEX) for Gait Rehabilitation of Motor-Impaired Patients," *IEEE Tenth International Conference on Rehabilitation Robotics (ICORR)*, Noordwijk, The Netherlands, June 13–15, pp. 401–407.
- [79] Lo, H. S., and Xie, S. Q., 2012, "Exoskeleton Robots for Upper-Limb Rehabilitation: State of the Art and Future Prospects," *Med. Eng. Phys.*, **34**(3), pp. 261–268.
- [80] Li, Y., Huegel, J. C., Patoglu, V., and O'Malley, M. K., 2009, "Progressive Shared Control for Training in Virtual Environments," *IEEE Third Joint Euro-Haptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (World Haptics 2009)*, Salt Lake City, UT, Mar. 18–20, pp. 332–337.
- [81] Crespo, L. M., and Reinkensmeyer, D. J., 2008, "Haptic Guidance Can Enhance Motor Learning of a Steering Task," *J. Mot. Behav.*, **40**(6), pp. 545–557.
- [82] O'Malley, M. K., Gupta, A., Gen, M., and Li, Y., 2006, "Shared Control in Haptic Systems for Performance Enhancement and Training," *ASME J. Dyn. Syst. Meas. Control*, **128**(1), pp. 75–85.
- [83] Schmidt, R. A., 1991, "Frequent Augmented Feedback Can Degrade Learning: Evidence and Interpretations," *Tutorials in Motor Neuroscience*, Springer, Dordrecht, The Netherlands, pp. 59–75.
- [84] Winstein, C. J., Pohl, P. S., and Lewthwaite, R., 1994, "Effects of Physical Guidance and Knowledge of Results on Motor Learning: Support for the Guidance Hypothesis," *Res. Q. Exercise Sport*, **65**(4), pp. 316–323.
- [85] Blank, A., Diogenes, H., and O'Malley, M., 2015, "Characterizing Haptic Interference to Enable Prosthesis Socket Integration of Haptic Feedback Via Sensory Substitution," *Third ASU Rehabilitation Robotics Workshop*, Tempe, AZ, Feb. 13–14.
- [86] Liang, X., Makatura, C. R., Schubert, M., Solomon, B. H., Walker, J. M., Blank, A. A., and O'Malley, M. K., 2014, "Skin-Stretch Proprioceptive Feedback for a Robotic Gripper," *IEEE Haptics Symposium (HAPTICS)*, Houston, TX, Feb. 23–26, p. 1.
- [87] Treadway, E., Gillespie, B., Bolger, D., Blank, A., O'Malley, M., and Davis, A., 2015, "The Role of Auxiliary and Referred Haptic Feedback in Myoelectric Control," *IEEE World Haptics Conference (WHC)*, Evanston, IL, June 22–26, pp. 13–18.
- [88] Walker, J. M., Blank, A. A., Shewokis, P. A., and O'Malley, M. K., 2014, "Tactile Feedback of Object Slip Improves Performance in a Grasp and Hold Task," *IEEE Haptics Symposium (HAPTICS)*, Houston, TX, Feb. 23–26, pp. 461–466.
- [89] Walker, J. M., Blank, A. A., Shewokis, P. A., and O'Malley, M. K., 2015, "Tactile Feedback of Object Slip Facilitates Virtual Object Manipulation," *IEEE Trans. Haptics*, **8**(4), pp. 454–466.
- [90] Johansson, R. S., and Flanagan, J. R., 2009, "Coding and Use of Tactile Signals From the Fingertips in Object Manipulation Tasks," *Nat. Rev. Neurosci.*, **10**(5), pp. 345–359.
- [91] Brown, J. D., Paek, A., Syed, M., O'Malley, M. K., Shewokis, P. A., Contreras-Vidal, J. L., Davis, A. J., and Gillespie, R. B., 2013, "Understanding the Role of Haptic Feedback in a Teleoperated/Prosthetic Grasp and Lift Task," *IEEE World Haptics Conference (WHC)*, Daejeon, South Korea, Apr. 14–17, pp. 271–276.
- [92] Christiansen, R., Contreras-Vidal, J. L., Gillespie, R. B., Shewokis, P. A., and O'Malley, M. K., 2013, "Vibrotactile Feedback of Pose Error Enhances Myoelectric Control of a Prosthetic Hand," *IEEE World Haptics Conference (WHC)*, Daejeon, South Korea, Apr. 14–17, pp. 531–536.
- [93] Sergi, F., Accoto, D., Campolo, D., and Guglielmelli, E., 2008, "Forearm Orientation Guidance With a Vibrotactile Feedback Bracelet: On the Directionality of Tactile Motor Communication," *Second IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob)*, Scottsdale, AZ, Oct. 19–22, pp. 433–438.
- [94] Kim, K., and Colgate, J. E., 2012, "Haptic Feedback Enhances Grip Force Control of sEMG-Controlled Prosthetic Hands in Targeted Reinnervation Amputees," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **20**(6), pp. 798–805.
- [95] D'Alonzo, M., Dosen, S., Cipriani, C., and Farina, D., 2014, "HyVE: Hybrid Vibro-Electrotactile Stimulation for Sensory Feedback and Substitution in Rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **22**(2), pp. 290–301.
- [96] Brickman, B. J., Hettinger, L. J., Roe, M. M., Lu, L., Repperger, D. W., and Haas, M. W., 1996, "Haptic Specification of Environmental Events: Implications for the Design of Adaptive, Virtual Interfaces," *IEEE Virtual Reality Annual International Symposium*, Santa Clara, CA, Mar. 30–Apr. 3, pp. 147–153.
- [97] Oviatt, S., Coulston, R., and Lunsford, R., 2004, "When Do We Interact Multimodally?: Cognitive Load and Multimodal Communication Patterns," *ACM Sixth International Conference on Multimodal Interfaces (ICMI)*, State College, PA, Oct. 13–15, pp. 129–136.
- [98] Sigrist, R., Rauter, G., Riener, R., and Wolf, P., 2013, "Augmented Visual, Auditory, Haptic, and Multimodal Feedback in Motor Learning: A Review," *Psychon. Bull. Rev.*, **20**(1), pp. 21–53.
- [99] Antfolk, C., D'Alonzo, M., Rosén, B., Lundborg, G., Sebelius, F., and Cipriani, C., 2013, "Sensory Feedback in Upper Limb Prosthetics," *Expert Rev. Med. Dev.*, **10**(1), pp. 45–54.
- [100] Schofield, J. S., Evans, K. R., Carey, J. P., and Hebert, J. S., 2014, "Applications of Sensory Feedback in Motorized Upper Extremity Prosthesis: A Review," *Expert Rev. Med. Dev.*, **11**(5), pp. 499–511.
- [101] Biddiss, E. A., and Chau, T. T., 2007, "Upper Limb Prosthesis Use and Abandonment: A Survey of the Last 25 Years," *Prosthetics Orthotics Int.*, **31**(3), pp. 236–257.
- [102] Lundborg, G., and Rosen, B., 2001, "Sensory Substitution in Prosthetics," *Hand Clin.*, **17**(3), pp. 481–488.
- [103] Wright, T. W., Hagen, A. D., and Wood, M. B., 1995, "Prosthetic Usage in Major Upper Extremity Amputations," *J. Hand Surg.*, **20**(4), pp. 619–622.
- [104] Kourtis, D., Kwok, H. F., Roach, N., Wing, A. M., and Praamstra, P., 2008, "Maintaining Grip: Anticipatory and Reactive EEG Responses to Load Perturbations," *J. Neurophysiol.*, **99**(2), pp. 545–553.
- [105] Wiertelowski, M., Endo, S., Wing, A. M., and Hayward, V., 2013, "Slip-Induced Vibration Influences the Grip Reflex: A Pilot Study," *IEEE World Haptics Conference (WHC)*, Daejeon, South Korea, Apr. 14–17, pp. 627–632.
- [106] Zhang, W., Gordon, A. M., Mclsaac, T. L., and Santello, M., 2011, "Within-Trial Modulation of Multi-Digit Forces to Friction," *Exp. Brain Res.*, **211**(1), pp. 17–26.
- [107] Nowak, D. A., Glasauer, S., and Hermsdörfer, J., 2004, "How Predictive is Grip Force Control in the Complete Absence of Somatosensory Feedback?," *Brain*, **127**(1), pp. 182–192.
- [108] Nowak, D. A., and Hermsdörfer, J., 2003, "Selective Deficits of Grip Force Control During Object Manipulation in Patients With Reduced Sensibility of the Grasping Digits," *Neurosci. Res.*, **47**(1), pp. 65–72.
- [109] Nowak, D. A., Hermsdörfer, J., Glasauer, S., Philipp, J., Meyer, L., and Mai, N., 2001, "The Effects of Digital Anaesthesia on Predictive Grip Force Adjustments During Vertical Movements of a Grasped Object," *Eur. J. Neurosci.*, **14**(4), pp. 756–762.
- [110] Biddiss, E., Beaton, D., and Chau, T., 2007, "Consumer Design Priorities for Upper Limb Prosthetics," *Disability Rehabil. Assist. Technol.*, **2**(6), pp. 346–357.
- [111] Dhillon, G. S., and Horch, K. W., 2005, "Direct Neural Sensory Feedback and Control of a Prosthetic Arm," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **13**(4), pp. 468–472.
- [112] Dhillon, G. S., Lawrence, S. M., Hutchinson, D. T., and Horch, K. W., 2004, "Residual Function in Peripheral Nerve Stumps of Amputees: Implications for Neural Control of Artificial Limbs," *J. Hand Surg.*, **29**(4), pp. 605–615.
- [113] Blank, A., Okamura, A. M., and Whitcomb, L. L., 2012, "User Comprehension of Task Performance With Varying Impedance in a Virtual Prosthetic Arm: A Pilot Study," *Fourth IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob)*, Rome, Italy, June 24–27, pp. 500–507.
- [114] Chaubey, P., Rosenbaum-Chou, T., Daly, W., and Boone, D., 2014, "Closed-Loop Vibratory Haptic Feedback in Upper-Limb Prosthetic Users," *JPO: J. Prosthetics Orthotics*, **26**(3), pp. 120–127.
- [115] Cheng, A., Nichols, K. A., Weeks, H. M., Gurari, N., and Okamura, A. M., 2012, "Conveying the Configuration of a Virtual Human Hand Using Vibrotactile Feedback," *IEEE Haptics Symposium (HAPTICS)*, Vancouver, BC, Canada, Mar. 4–7, pp. 155–162.

- [116] Cipriani, C., D'Alonzo, M., and Carrozza, M. C., 2012, "A Miniature Vibrotactile Sensory Substitution Device for Multifingered Hand Prosthetics," *IEEE Trans. Biomed. Eng.*, **59**(2), pp. 400–408.
- [117] Rombokas, E., Stepp, C. E., Chang, C., Malhotra, M., and Matsuoka, Y., 2013, "Vibrotactile Sensory Substitution for Electromyographic Control of Object Manipulation," *IEEE Trans. Biomed. Eng.*, **60**(8), pp. 2226–2232.
- [118] Bark, K., Wheeler, J., Shull, P., Savall, J., and Cutkosky, M., 2010, "Rotational Skin Stretch Feedback: A Wearable Haptic Display for Motion," *IEEE Trans. Haptics*, **3**(3), pp. 166–176.
- [119] Gurari, N., and Okamura, A. M., 2014, "Compliance Perception Using Natural and Artificial Motion Cues," *Multisensory Softness*, Springer, London, pp. 189–217.
- [120] Erwin, A., and Sup, F. C., IV, 2015, "A Haptic Feedback Scheme to Accurately Position a Virtual Wrist Prosthesis Using a Three-Node Tactor Array," *PLoS One*, **10**(8), p. e0134095.
- [121] Shull, P. B., and Damian, D. D., 2015, "Haptic Wearables as Sensory Replacement, Sensory Augmentation and Trainer—A Review," *J. Neuroeng. Rehabil.*, **12**(1), p. 59.
- [122] Marasco, P. D., Schultz, A. E., and Kuiken, T. A., 2009, "Sensory Capacity of Reinnervated Skin After Redirection of Amputated Upper Limb Nerves to the Chest," *Brain*, **132**(6), pp. 1441–1448.
- [123] Stubblefield, K. A., Miller, L. A., Lipschutz, R. D., and Kuiken, T. A., 2009, "Occupational Therapy Protocol for Amputees With Targeted Muscle Reinnervation," *J. Rehabil. Res. Dev.*, **46**(4), pp. 484–488.
- [124] Ohnishi, K., Weir, R. F., and Kuiken, T. A., 2007, "Neural Machine Interfaces for Controlling Multifunctional Powered Upper-Limb Prostheses," *Expert Rev. Med. Dev.*, **4**(1), pp. 43–53.
- [125] Massimino, M. J., and Sheridan, T. B., 1994, "Teleoperator Performance With Varying Force and Visual Feedback," *Hum. Factors*, **36**(1), pp. 145–157.
- [126] Richard, P., and Coiffet, P., 1995, "Human Perceptual Issues in Virtual Environments: Sensory Substitution and Information Redundancy," Fourth IEEE International Workshop on Robot and Human Communication (RO-MAN'95), Tokyo, Japan, July 5–7, pp. 301–306.
- [127] Fabiani, L., Burdea, G. C., Langrana, N. A., and Gomez, D., 1996, "Human Interface Using the Rutgers Master II Force Feedback Interface," IEEE Virtual Reality Annual International Symposium, Santa Clara, CA, Mar. 30–Apr. 3, pp. 54–59.
- [128] Meech, J., and Solomonides, A., 1996, "User Requirements When Interacting With Virtual Objects," IEE Colloquium on Virtual Reality—User Issues (Digest No: 1996/068), London, Mar. 25, p. 3.
- [129] Adams, R. J., Klowden, D., and Hannaford, B., 2001, "Virtual Training for a Manual Assembly Task," *Elec. J. Haptics Res.*, **2**(2), pp. 1–7.
- [130] Williams, L. E., Loftin, R. B., Aldridge, H. A., Leiss, E. L., and Bluethmann, W. J., 2002, "Kinesthetic and Visual Force Display for Telerobotics," IEEE International Conference on Robotics and Automation (ICRA'02), Washington, DC, May 11–15, pp. 1249–1254.
- [131] O'Malley, M. K., Hughes, K. J., Magruder, D. F., and Ambrose, R. O., 2003, "Simulated Bilateral Teleoperation of Robonaut," AIAA Paper No. 2003-6272.
- [132] Burdea, G. C., and Brooks, F. P., 1996, *Force and Touch Feedback for Virtual Reality*, Wiley, New York.
- [133] Huegel, J. C., Lynch, A. J., and O'Malley, M. K., 2009, "Validation of a Smooth Movement Model for a Human Reaching Task," IEEE International Conference on Rehabilitation Robotics (ICORR), Kyoto, Japan, June 23–26, pp. 799–804.
- [134] Todorov, E., Shadmehr, R., and Bizzi, E., 1997, "Augmented Feedback Presented in a Virtual Environment Accelerates Learning of a Difficult Motor Task," *J. Mot. Behav.*, **29**(2), pp. 147–158.
- [135] Huegel, J. C., and O'Malley, M. K., 2010, "Progressive Haptic and Visual Guidance for Training in a Virtual Dynamic Task," IEEE Haptics Symposium, Waltham, MA, Mar. 25–26, pp. 343–350.
- [136] Huegel, J. C., and O'Malley, M. K., 2014, "Workload and Performance Analyses With Haptic and Visually Guided Training in a Dynamic Motor Skill Task," *Computational Surgery and Dual Training*, Springer, New York, pp. 377–387.
- [137] Hale, K. S., and Stanney, K. M., 2004, "Deriving Haptic Design Guidelines From Human Physiological, Psychophysical, and Neurological Foundations," *IEEE Comput. Graph. Appl.*, **24**(2), pp. 33–39.
- [138] Marchal-Crespo, L., and Reinkensmeyer, D. J., 2009, "Review of Control Strategies for Robotic Movement Training After Neurologic Injury," *J. Neuroeng. Rehabil.*, **6**(1), p. 1.
- [139] Stanley, A. A., and Kuchenbecker, K. J., 2012, "Evaluation of Tactile Feedback Methods for Wrist Rotation Guidance," *IEEE Trans. Haptics*, **5**(3), pp. 240–251.
- [140] Bark, K., Hyman, E., Tan, F., Cha, E., Jax, S. A., Buxbaum, L. J., and Kuchenbecker, K. J., 2015, "Effects of Vibrotactile Feedback on Human Learning of Arm Motions," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **23**(1), pp. 51–63.
- [141] Norman, S. L., Doxon, A. J., Gleeson, B. T., and Provancher, W. R., 2014, "Planar Hand Motion Guidance Using Fingertip Skin-Stretch Feedback," *IEEE Trans. Haptics*, **7**(2), pp. 121–130.
- [142] Rotella, M. F., Guerin, K., He, X., and Okamura, A. M., 2012, "HAPI Bands: A Haptic Augmented Posture Interface," IEEE Haptics Symposium (HAPTICS), Vancouver, BC, Canada, Mar. 4–7, pp. 163–170.
- [143] Alahakone, A. U., and Senanayake, S. A., 2010, "A Real-Time System With Assistive Feedback for Postural Control in Rehabilitation," *IEEE/ASME Trans. Mechatronics*, **15**(2), pp. 226–233.
- [144] Wall, C., III, and Kentala, E., 2010, "Effect of Displacement, Velocity, and Combined Vibrotactile Tilt Feedback on Postural Control of Vestibulopathic Subjects," *J. Vestibular Res.*, **20**(1–2), pp. 61–69.
- [145] Sienko, K. H., Balkwill, M. D., Oddsson, L., and Wall, C., 2008, "Effects of Multi-Directional Vibrotactile Feedback on Vestibular-Deficient Postural Performance During Continuous Multi-Directional Support Surface Perturbations," *J. Vestibular Res.*, **18**(5–6), pp. 273–285.
- [146] Sienko, K. H., Balkwill, M. D., Oddsson, L. I., and Wall, C., 2013, "The Effect of Vibrotactile Feedback on Postural Sway During Locomotor Activities," *J. Neuroeng. Rehabil.*, **10**(1), pp. 65–73.
- [147] Prewett, M. S., Elliott, L. R., Walvoord, A. G., and Covert, M. D., 2012, "A Meta-Analysis of Vibrotactile and Visual Information Displays for Improving Task Performance," *IEEE Trans. Syst., Man, Cybern., Part C*, **42**(1), pp. 123–132.
- [148] Ruffaldi, E., Filippeschi, A., Frisoli, A., Sandoval, O., Avizzano, C. A., and Bergamasco, M., 2009, "Vibrotactile Perception Assessment for a Rowing Training System," Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (World Haptics 2009), Salt Lake City, UT, Mar. 18–20, pp. 350–355.
- [149] Buteau, J., Coquillart, S., Payan, Y., and Gentaz, E., 2008, "Haptic Guidance Improves the Visuo-Manual Tracking of Trajectories," *PLoS One*, **3**(3), p. e1775.
- [150] Palluel-Germain, R., Bara, F., de Boisferon, A. H., Hennion, B., Gougout, P., and Gentaz, E., 2007, "A Visuo-Haptic Device-Telematique-Increases Kindergarten Children's Handwriting Acquisition," IEEE Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07), pp. 72–77.
- [151] Li, Y., Patoglu, V., and O'Malley, M. K., 2009, "Negative Efficacy of Fixed Gain Error Reducing Shared Control for Training in Virtual Environments," *ACM Trans. Appl. Percept.*, **6**(1), p. 3.
- [152] Gillespie, R. B., O'Modhrain, M., Tang, P., Zaretzky, D., and Pham, C., 1998, "The Virtual Teacher," ASME International Mechanical Engineering Congress and Exposition, Anaheim, CA, Nov. 15–20, pp. 171–178.
- [153] Endo, T., Kawasaki, H., Kigaku, K., and Mouri, T., 2007, "Transfer Method of Force Information Using Five-Fingered Haptic Interface Robot," IEEE Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07), Tsukuba, Japan, Mar. 22–24, pp. 599–600.
- [154] Lo, A. C., Guarino, P. D., Richards, L. G., Haselkorn, J. K., Wittenberg, G. F., Federman, D. G., Ringer, R. J., Wagner, T. H., Krebs, H. I., Volpe, B. T., Bever, C. T., Bravata, D. M., Duncan, P. W., Corn, B. H., Maffucci, A. D., Nadeau, S. E., Conroy, S. S., Powell, J. M., Huang, G. D., and Peduzzi, P., 2010, "Robot-Assisted Therapy for Long-Term Upper-Limb Impairment After Stroke," *N. Engl. J. Med.*, **362**(19), pp. 1772–1783.
- [155] Handley, A., Medcalf, P., Hellier, K., and Dutta, D., 2009, "Movement Disorders After Stroke," *Age Ageing*, **38**(3), pp. 260–266.
- [156] Pehlivan, A. U., Sergi, F., Erwin, A., Yozbatiran, N., Francisco, G. E., and O'Malley, M. K., 2014, "Design and Validation of the RiceWrist-S Exoskeleton for Robotic Rehabilitation After Incomplete Spinal Cord Injury," *Robotica*, **32**(8), pp. 1415–1431.
- [157] Catalano, M. G., Grioli, G., Farnioli, E., Serio, A., Piazza, C., and Bicchi, A., 2014, "Adaptive Synergies for the Design and Control of the Pisa/IIT Soft-Hand," *Int. J. Rob. Res.*, **33**(5), pp. 768–782.
- [158] Santello, M., Flanders, M., and Soechting, J. F., 1998, "Postural Hand Synergies for Tool Use," *J. Neuroscience*, **18**(23), pp. 10105–10115.
- [159] Thakur, P. H., Bastian, A. J., and Hsiao, S. S., 2008, "Multidigit Movement Synergies of the Human Hand in an Unconstrained Haptic Exploration Task," *J. Neurosci.*, **28**(6), pp. 1271–1281.
- [160] Bicchi, A., Gabbicini, M., and Santello, M., 2011, "Modelling Natural and Artificial Hands With Synergies," *Philos. Trans. R. Soc. B*, **366**(1581), pp. 3153–3161.
- [161] Godfrey, S. B., Ajoudani, A., Catalano, M. G., Grioli, G., and Bicchi, A., 2013, "A Synergy-Driven Approach to a Myoelectric Hand," 13th International Conference on Rehabilitation Robotics (ICORR), Seattle, WA, June 24–26, pp. 1–6.
- [162] Fani, S., Bianchi, M., Jain, S., Pimenta Neto, J. S., Boege, S., Grioli, G., Bicchi, A., and Santello, M., 2016, "Assessment of Myoelectric Controller Performance and Kinematic Behavior of a Novel Soft Synergy-Inspired Robotic Hand for Prosthetic Applications," *Front. Neurobotics*, **10**, p. 11.
- [163] Pylatiuk, C., Schulz, S., and Döderlein, L., 2007, "Results of an Internet Survey of Myoelectric Prosthetic Hand Users," *Prosthetics Orthotics Int.*, **31**(4), pp. 362–370.
- [164] Casini, S., Morvidoni, M., Bianchi, M., Catalano, M., Grioli, G., and Bicchi, A., 2015, "Design and Realization of the CUFF-Clenching Upper-Limb Force Feedback Wearable Device for Distributed Mechano-Tactile Stimulation of Normal and Tangential Skin Forces," IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hamburg, Germany, Sept. 28–Oct. 2, pp. 1186–1193.
- [165] Godfrey, S. B., Bianchi, M., Bicchi, A., and Santello, M., 2016, "Influence of Force Feedback on Grasp Force Modulation in Prosthetic Applications: A Preliminary Study," IEEE 38th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC), Orlando, FL, Aug. 16–20, pp. 5439–5442.
- [166] Ajoudani, A., Godfrey, S. B., Bianchi, M., Catalano, M. G., Grioli, G., Tsagarakis, N., and Bicchi, A., 2014, "Exploring Teleimpedance and Tactile Feedback for Intuitive Control of the Pisa/IIT SoftHand," *IEEE Trans. Haptics*, **7**(2), pp. 203–215.