ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

Embodied Communication: How Robots and People Communicate Through Physical Interaction

Aleksandra Kalinowska,[1] Patrick M. Pilarski,[2,3] and Todd D. Murphey[1,4]

[1]Department of Mechanical Engineering, Northwestern University, Evanston, Illinois, USA; email: t-murphey@northwestern.edu

[2]Department of Medicine, University of Alberta, Edmonton, Alberta, Canada

[3]Alberta Machine Intelligence Institute (Amii), Edmonton, Alberta, Canada

[4]Department of Physical Therapy and Human Movement Sciences, Northwestern University, Chicago, Illinois, USA

Annu. Rev. Control Robot. Auton. Syst. 2023. 6:205–32

First published as a Review in Advance on January 9, 2023

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-070122-102501

Keywords

human–robot collaboration, embodied communication, physical interaction, assistive robots

Abstract

Early research on physical human–robot interaction (pHRI) has necessarily focused on device design—the creation of compliant and sensorized hardware, such as exoskeletons, prostheses, and robot arms, that enables people to safely come in contact with robotic systems and to communicate about their collaborative intent. As hardware capabilities have become sufficient for many applications, and as computing has become more powerful, algorithms that support fluent and expressive use of pHRI systems have begun to play a prominent role in determining the systems' usefulness. In this review, we describe a selection of representative algorithmic approaches that regulate and interpret pHRI, describing the progression from algorithms based on physical analogies, such as admittance control, to computational methods based on higher-level reasoning, which take advantage of multimodal communication channels. Existing algorithmic approaches largely enable task-specific pHRI, but they do not generalize to versatile human–robot collaboration. Throughout the review and in our discussion of next steps, we therefore argue that emergent embodied

dialogue—bidirectional, multimodal communication that can be learned through continuous interaction—is one of the next frontiers of pHRI.

1. INTRODUCTION

As robots have become more capable over recent decades, physical human–robot interaction (pHRI; see the sidebar titled Physical Human–Robot Interaction) is no longer a nuisance but instead an opportunity for improved human–machine communication and closer collaboration (1). A major milestone in pHRI has been the introduction of specialized hardware approaches, including series-elastic actuators, cobots, and distributed sensing. Progress in perception has made robots more situationally aware of people during operation. Improved sensing, both novel (e.g., haptic skins) and traditional (e.g., improved force sensing at the joints), has made pHRI practical, particularly in terms of safety. Progress in interfaces for capturing intent [high-density electromyography (EMG), dry electroencephalography (EEG), context-sensitive joysticks, etc.] has made devices easier to use. In parallel, progress in computing power has enabled real-time execution of sophisticated algorithms. Although there is still plenty of room for hardware innovations, these developments have made algorithm design an important next stage in pHRI. With novel algorithmic approaches, available information channels have new potential utility as affordances for embodied communication—multimodal communication through actions, explicit interfaces, and physical contact (see the sidebar titled Embodied Communication). Thanks to these advances, we can start to design for facets of pHRI that were previously impractical, e.g., the subtle ways in which physical interaction creates an opportunity for dialogue between a robot and a person during use.

Historically, pHRI methods have focused on one-directional communication, typically a robot inferring something about a person through observation and providing assistance in response. Sometimes robots provide feedback about the environment or their internal state. This type of mostly unidirectional communication can be effective, but it does not take advantage of the collaborative potential of the human–robot pair—the ability to negotiate approaches, mutually adapt to each other and the environment, or coordinate a response to novelty. As a result, this approach does not generalize well to novel, unstructured interactions and therefore tends to be best suited to narrow, prespecified applications. If we want robots to be versatile collaborative partners, we need better communicative capacity (2).

Notably, effective human–human communication relies on a bidirectional, multimodal dialogue. In his book *Speaking Our Minds* (3), Thom Scott-Phillips distinguishes two ways of describing human communication—code model and ostensive–inferential communication—and builds a case for why code model communication cannot explain people's incredibly flexible communication capabilities. Code model communication is achieved through pairs of association: one

PHYSICAL HUMAN–ROBOT INTERACTION

The term pHRI, as used in this review, spans multiple forms of physical interaction, ranging from haptics and forceful touch to less traditional interfaces, such as vibrotactile sleeves, on-body sensors, and haptic joysticks. As a consequence, facilitating pHRI plays a key role in the design of any robot that works in close proximity with its human partners, including tightly coupled exoskeletons that move synchronously with the wearer as well as stand-alone robot arms that experience intermittent contact with their human collaborators.

EMBODIED COMMUNICATION

From Unidirectional Human–Machine Communication . . .

Human–robot teams rely on embodied communication—a multimodal exchange of information through actions, touch, forceful interaction, biometric signals, verbal cues, and more. In its current form, communication is often unidirectional and static over time: The robot is a passive observer of human intent, providing limited feedback to its human partner, and the communicative conventions do not evolve over time. The robot's feedback, if any, usually does not depend on the human partner's mental state or the communication history.

. . . Toward Emergent Embodied Dialogue

Effective human–robot collaboration can benefit from flexible, bidirectional dialogue which relies on communicative conventions that (*a*) emerge through interaction and (*b*) account for the partners' mental states and communication history. In this viewpoint, the human and the robot are collaborative agents, refining their communicative capacity over time and using communication to negotiate approaches to physical tasks.

between a state of the world and a signal, and the other between the signal and a response. Morse code is an example of this type of communication—it is defined by unambiguous symbols that can be broadcast into the ether and interpreted reliably with a decoder. The code enables an effective exchange of information, but it does not allow fluent collaboration. Current human–machine communication at best aspires to code model communication.

Ostensive–inferential communication is different—it relies on people's ability to express and recognize intentions (3). Ostensive–inferential communication depends strictly on one agent influencing the mental state of another agent, and so any behavior can, in principle, be used communicatively, so long as it influences the mental state of the other agent in the intended manner. In this viewpoint, natural languages function to make ostensive–inferential communication more precise and more expressive than it otherwise would be, but they are only one tool in the communicative toolbox. In addition to communicating verbally, people rely largely on nonverbal cuing during joint action, wherein they model each other's activity and contributions to a shared goal and enact voluntary nonverbal signaling in service of the shared goal (4–10). Human–machine collaboration stands to benefit from an ostensive–inferential communication approach.

Finally, much research and development effort is spent on making human–machine communication as intuitive as possible. This includes designing the robot to interpret either anthropomorphic signaling (e.g., hand gestures and natural language) or intuitive intent signaling through biological signals (e.g., EMG and EEG) to influence robot behavior. These forms of communicative conventions are often insufficient for fluent/expressive device use and may not even be optimal for human–machine communication—anthropomorphic signaling evolved and works well for human–human communication, while measurable biological signals are noisy and only somewhat informative. Importantly, people are good at adapting and developing communication protocols. How can we put more emphasis on the emergence of embodied communication *through* pHRI, rather than replicating familiar communication conventions *for* pHRI?

In this review, we explore how embodied dialogue can create a path for more natural collaborative workflows and how it can facilitate safer, more effective interaction. In Section 2, we describe existing algorithmic approaches to facilitating human–machine communication, from algorithms based on physical analogies, such as admittance and impedance control, to computational methods capable of task-level reasoning and adaptation. In Section 3, we discuss the

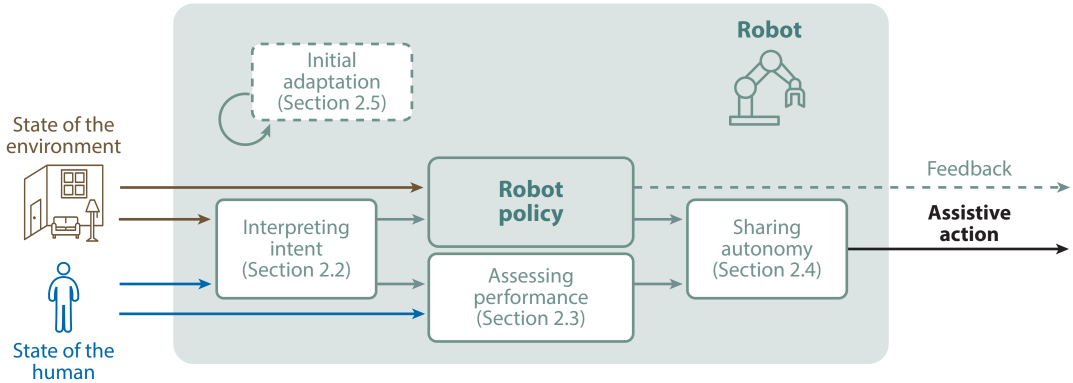

Figure 1

Integral components of existing algorithmic approaches for physical human–robot interaction (pHRI), as discussed in Section 2. At each time step, the robot predicts intent based on the state of the human—i.e., motion, biometric signals, etc.—and the state of the environment (Section 2.2). Using a metric of human performance (Section 2.3), a shared control paradigm (Section 2.4), and its task-based policy, the robot decides on an assistive action. Many approaches undergo an initial period of adaptation (Section 2.5), and some enable the robot to take additional communicative actions to provide feedback for the human collaborator.

state-of-the-art communication paradigms for different applications of pHRI and comment on how emergent human–machine dialogue could improve performance. In Section 4, we conclude with an outline of future directions for developing embodied communication in pHRI systems. We discuss the need to algorithmically facilitate bidirectional human–robot communication, where both the human and robot can play a role in shaping the communication protocol. As a key position statement established by this review, we argue that to achieve versatile human–machine collaboration, we need algorithmic methods that enable emergent embodied dialogue based on an ostensive–inferential model of communication.

2. ALGORITHMIC APPROACHES TO EMBODIED COMMUNICATION

Effective robotic assistance depends on successful communication of collaborative intent from both the human and the machine. In pHRI, the simplest form of embodied communication is through movement and physical contact, which can be regulated on an action-by-action basis (11). With additional sensing, the robot can more reliably reason about the person's intent and assess their performance (1); depending on the goal of the interaction, it can adjust assistance, often putting its control strategy into the context of a task. The interaction benefits if the robot can adapt its communication protocol and assistive strategy to its collaborative partner (2, 12, 13). In this section, we describe the current approaches to embodied communication in pHRI, ranging from task-agnostic communication based on compliant physical contact to adaptive task-oriented communication strategies. **Figure 1** illustrates how the algorithmic approaches work together to generate robot actions.

2.1. Admittance and Impedance

Admittance and impedance control were arguably the first algorithmic approaches for regulating pHRI that enabled reasonably safe human–machine interaction. Admittance control dates back to the late 1970s (14), when it was used to respond to hard contact in industrial applications. Impedance control drew academic interest in the 1980s (11) and was one of the first algorithms used more broadly to control physical interaction with a robot.

Fundamentally, both admittance and impedance control imitate mechanical properties of contact between nonrigid systems, modulating elasticity in the collision. In impedance control, motion is detected and converted into interaction forces through an internal model. Impedance control is good at rendering low inertia but has difficulty with stiff virtual surfaces. In robot-assisted rehabilitation, it can be used to implement active assistance or active resistance proportional to participant movement (15) without the hardware requirement of a force sensor. Admittance control (16), on the other hand, converts measured forces into motion, requiring sensors to measure the applied forces. It enables rendering of stiff virtual surfaces but struggles with constrained motion, e.g., interacting with real surfaces (17). In human–robot collaboration, admittance control can be used to account for unintended collisions (18) or to jointly manipulate an object (19). It is commonly used to reduce the inertia of bulky devices with a payload (20) or to ease movement during rehabilitation (15). The HapticMaster (21) is an example of a commercially available admittance-controlled end-effector robot.

Both admittance and impedance control assume a static model that maps human actions to robot actions based on mechanical analogies. The model's parameters can be tuned, using intuition about the mechanical analogies as a guide, but given that the robot is not responsible for any high-level reasoning about a task goal, the complexity of the interaction is limited. Communication is constrained to negotiating real-time actions without the possibility of task-level coordination of movement. Algorithmically, this is a conservative attitude to facilitating pHRI, restricting the physical interactions to those that have mechanical properties [even if not always mechanically plausible, such as in work by Patton et al. (22)]. With richer sensory inputs and onboard computation, the robot can begin to reason about the task, in terms of both its goal and a plausible solution. As a result, admittance and impedance control are often combined with online algorithms that reason about human intent and enable more effective task-oriented assistance, as described in the sections below.

2.2. Predicting Intent

In human–robot collaboration, the human and the robot often have complementary roles: The person has a high-level understanding of the task, while the robot has physical capabilities that can help the person accomplish the task. One way of framing the communication that occurs between a person and a robot is to view the robot as an observer, tasked with inferring the person's intent. Once the autonomy has an understanding of high-level intent, it can provide assistance that is not simply proportional to a person's motion or forces (as is the case in admittance and impedance control). Below, we describe different ways of inferring intent based on available sensory information. Example sensors for pHRI are shown in **Figure 2**.

Figure 2

Example devices for capturing human intent: (*a*) motion-capture suit from Motion Workshop, (*b*) high-density electromyograph from OT Bioelettronica, (*c*) wet electroencephalography cap from TMSi, (*d*) joystick from MERU, and (*e*) sip-n-puff from Therafin Corporation. All panels reproduced with permission from their respective companies.

2.2.1. From body motion and eye gaze. Motion data can be captured passively (without requiring explicit human input) through a range of sensors on the body, including inertial measurement units, motion-capture markers, cameras, or encoders on back-drivable electric motors. With recent advances in computer vision, it is possible to reliably estimate pose from RGB images (23), as well as to track eye gaze during real-time interactions (24). For stand-alone robots, camera-based monitoring of motion offers a significant advantage over most alternatives, because it does not require sensors to be placed on the human body.

However, even with information about motion history, it is challenging to anticipate intended human motion. Over the years, many methods have been proposed for recognizing motion (25). One type of reasoning classifies motion into primitives or modes, where the mode sequences can be predicted based on a task objective (26, 27), while motion trajectories within modes can be anticipated based on an approximation of motion dynamics (28). Other studies have proposed velocity-based position projection (29), general value functions trained using temporal-difference learning (30), or a combination of multiple approaches (29). Inferring intent from motion assumes that the kinematics of motion are sufficiently rich to capture a person's goals. In most scenarios, motion trajectories are dependent on contextual information, and therefore improving predictions requires inferring the situational context and/or obtaining additional sensory inputs. As an example, some studies augment kinematics data with EMG signals to improve predictions of motor intent (13, 30).

Motion can also be used to purposefully communicate intent, e.g., through gestures (31), motion cues (32), facial movements (33), shoulder shrugs (170), or gaze (34, 35). Substantial progress has been made in sign language recognition that could be used in human–robot collaboration (36). While these studies offer promising results, it is important to note that explicitly prescribing a set of communicative conventions for each motion type is difficult and impractical because of the scope of possible communicative symbols. In Section 4.3, we describe the need for algorithmic solutions that would enable communicative conventions to emerge from interaction.

2.2.2. From electromyography. Electrical activity from skeletal muscles, known as EMG, is an appealing source of intent prediction, particularly for assistive devices that are directly coupled with the human body, such as prosthetics and exoskeletons. There are many documented approaches to EMG-based control (13, 30, 37, 38), including a growing body of work on the use of advanced machine learning methods (37). Even so, two of the most commonly studied control methods continue to be (*a*) proportional control, where the EMG signals are directly mapped onto low-level trajectories (39, 40), and (*b*) mode-based control, where the EMG signals are classified into modes based on pattern recognition techniques and used in combination with an autonomous low-level controller (41, 42). Proportional control gives the user direct control over robot motion but requires precise decoding of EMG signals. As a result, multijoint motions are typically difficult to achieve. Mode-based control expects sparser high-level commands from the human operator and has been shown to achieve better performance (43). However, studies highlight the need for more research on the benefits of pattern recognition for mode-based control (44, 45). Combining EMG signals with other control inputs has been noted by the community as a promising path forward (13). One example is EMG combined with EEG (46, 47), which has shown encouraging initial potential.

A significant appeal of using EMG signals is their intuitiveness. However, while people have some control over the EMG signal by voluntarily generating muscle contractions, they are limited by their physiology in the variety and expressivity of patterns that they can generate. As a result, some work has explored EMG control using nonbiological mappings, where the person generates maximally differentiated muscle contractions (e.g., internal/external rotation of the forearm) and these contractions are mapped onto useful configurations of a robot (e.g., different grips of a hand prosthesis) (48). These approaches show promising performance but can be difficult for a person to learn. Emergent interfaces, discussed in Section 4.3, might offer a solution.

EMG-based methods are in many ways similar to motion-based approaches, in that they infer intent based on a planned execution of motion. They have the advantage of being able to anticipate motion before it happens or, as in the case of amputees, detect motion intent. However, neither motion-based nor EMG-based methods attempt to infer the cognitive intent of a person. EEG-based techniques, discussed next, try to use high-level signaling in the brain to infer cognitive aspects of intent.

2.2.3. From electroencephalography. Reliably capturing and interpreting electrical activity of the brain would enable thought-controlled devices, or at least devices that are explicitly dependent on what a person is thinking. If successful, EEG paves the path for intuitive interfaces to control multidimensional robots. Thus far, however, surface EEG (sEEG) has proven difficult to use as a control interface for robots (49) because of the noisiness and low spatial resolution of sEEG signals. Electrical brain signals are much stronger when measured inside the brain through an implant, rather than noninvasively on the surface of the skin; this is a promising avenue of research but has received limited attention because of the risks associated with its invasiveness (50, 51).

Algorithmic approaches for interpreting EEG signals are ample, ranging from using combinations of prespecified features (e.g., amplitudes of signal within frequency bands in the motor cortices) (52) to training deep neural networks on the raw signal (53). State-of-the-art algorithms enable reaching a target in a 2D plane by interpreting the sEEG signal as one of a handful of mental states, corresponding to possible target locations of the robot's end effector (49). As an example, EEG-predicted mental states in combination with a context-aware controller enabled a person to navigate a mobile robot through a network of connected rooms (54). Given the low spatial resolution of sEEG signals, researchers have had difficulty identifying more than three classes of distinct mental activities (49). In more recent studies, the error potential has been used to successfully correct the 3D motion of a robot arm (47, 55). This is an interesting approach that takes advantage of a biological interpretation of sEEG signals—the error potential is detectable when a person perceives error.

The key outcome of these studies is that sEEG provides a way to record an array of signals from the brain and, for now, to interpret them in a small number of ways. One of the points we make later is that these interpretations are prespecified—e.g., they are trained to mean *right*, *left*, or *straight*—when they could be an unlabeled range of mental states that the robot learns to interpret over time.

2.2.4. From controller signals. There are a range of designated interfaces that can provide human-generated input for an assistive robot. The most common are control devices that can be operated by the hand, such as a joystick or remote controller with a range of buttons. Hand-operated control devices are currently in commercial use, and although they require effort (compared with passively measured biological signals), their control signal is more informative and less noisy—they already achieve adequate performance for the control of powered wheelchairs to warrant wide adoption.

In addition to hand-operated devices, there exist a range of control interfaces that are designed to be used by individuals who are paralyzed below the neck. These include a head array (which enables button presses with the head) and a sip-n-puff (which records a 1D continuous signal from the person exhaling and inhaling air into a sensorized straw) (56). Like the hand-operated controllers, these interfaces are already in use outside of research laboratories. However, signals from these interfaces are lower-dimensional, posing a challenge even when controlling a 2D powered wheelchair.

In commercially available products, the control interfaces (i.e., joysticks, head arrays, and sip-n-puff devices) are treated like many interfaces used in nonrobotic applications (e.g., video games)—the interface mapping of control inputs onto the action space is designed ahead of time, and it is up to the person to learn and adjust to the prespecified mapping. Notably, research has shown that the physical interaction with the control interface is nonnegligible and should be taken into account when designing the mapping from interface inputs to control outputs (57). Although intentionally generated, the signal is often imperfect; recent work has highlighted the effects of impairment and interface type on the acquired control signal, including the timing, transient noise, and accuracy of a signal (58). When designing mappings for these control interfaces, there is a need to enable adaptation to the person's unique physical and cognitive capabilities—effective communication can give the person increased agency over their interaction with the robot.

2.2.5. Through natural language. With recent progress in natural language processing (59, 60), researchers have incorporated voice commands into human–robot interaction (61). Verbal communication can be a useful tool for high-level task alignment and will likely remain an important research direction. However, control using natural language has three main limitations in pHRI: (*a*) It is nonpragmatic for many applications (e.g., a robot assisting with eating), (*b*) it is socially inconvenient (e.g., walking in a prosthesis and talking), and (*c*) it requires high levels of effort for continuous control (i.e., guiding low-level movement). In the context of pHRI, the most that voice commands can do is align the robot with the person's high-level intent; actual robot actions need to be planned in the context of the body, environment, and continuous-time, continuous-space decision-making. Voice commands will likely be used in combination with other interfaces to provide high-level directions or corrective feedback.

2.3. Metrics of Motion for Adjusting Assistance

Given a prediction of high-level intent, a robot requires a mathematical specification of performance that enables continuous evaluation of human actions. In combination with decision-making techniques (e.g., linear control, optimal control, or reinforcement learning), metrics of motion enable computational synthesis of real-time assistance.

2.3.1. Motion quality. One way to assess performance is by quantifying motion quality. As an example, energy minimization has been shown to explain the dynamics of human motion, and so energy is a universally used assessment of motion quality (62). Motion has also been assessed in a task-independent manner using the norm of mechanical jerk—the third time derivative of state (63). Jerk is a way to quantify motion smoothness, with the underlying assumption that smooth motion is desirable.

2.3.2. Task performance. Another intuitive way to assess human actions is with respect to a reference solution to a task goal. If one defines a task solution as a trajectory in time, one can use error with respect to a predefined trajectory to provide real-time assistance (63). This approach works well if there exists an optimal solution to a given task but overspecifies the task solution when many trajectories can, for practical purposes, be considered equally correct. For most tasks, there exist many correct ways to execute an action sequence and achieve a high-level goal (64).

As an alternative, new metrics of task performance have been proposed based on decomposing trajectories in ways that ignore the solution's evolution in the time domain. Solution quality can be assessed based on the relative positioning of joints with respect to one another (65). It can also be measured using distribution-based metrics that evaluate the state statistics of a motion trajectory relative to the statistics of a task (64, 66), or using frequency-domain assessments that quantify motion bandwidth (67, 68). Trajectory-independent ways of quantifying task performance allow one to equally evaluate multiple approaches to solving a task and do not constrain the solution as a function of time.

2.3.3. Motion predictability.

Learned forecasts of movement and movement anomalies can also be used as a metric to modulate pHRI. As one key example, researchers have shown that general value functions can be used to build predictions of anticipated motion and to make control decisions based on the motion's deviation from the learned prediction (69, 70). Similarly, researchers have introduced and made use of metrics that quantify surprise. As an example, a study illustrated that unexpected perturbations can be characterized through the process of continual temporal-difference learning, and that the resulting metrics, such as the unexpected demon error, can be used to effectively modulate assistance (71).

2.3.4. Effort.

Lastly, in physical assistance, there is often interest in assessing effort exerted by the individual (72). Effort can be measured directly using EMG (73), heart rate variability (74), or respiratory data on the flow rates of oxygen and carbon dioxide (75, 76). As an example, one can adjust the parameters of an admittance or impedance controller based on EMG readings to maintain a desired level of physical effort (77).

Task effort can also be approximated via cognitive load. Researchers have attempted to measure cognitive load directly through EEG signal (78); other studies have quantified cognitive effort via proxy metrics, such as gaze trajectory and pupil size (79), or a combination of physiological data, such as heart rate, breathing frequency, skin conductance, and skin temperature (80). As with metrics of physical effort, metrics of cognitive load can be used in closed loop to provide real-time adjustments to the provided assistance (80) and maintain patient engagement.

While it is challenging to quantitatively assess task performance in real time and to mathematically specify motion quality or effort with enough generality, flexible specifications of motion can improve the adjustability of robotic assistance as well as enable the development of provable safety guarantees for pHRI systems.

2.4. Task-Based Shared Control Paradigms

Once a task goal and performance metric have been established, a consideration for the autonomy is when and how much assistance to provide. Shared control paradigms distribute control between the human and robot to improve overall performance and/or safety.

Performance-based shared control schemes rely on an estimate of the user's intent, as described in Section 2.2. Intent could be a high-level goal, such as a navigation landmark to drive a wheelchair toward or an object to grasp using an assistive robot arm (81). Intent could also involve avoiding obstacles, in which case distance to obstacles could be a metric for allocating control between the human and robot (82).

Once a task goal has been identified, user control inputs can be directly modified by the autonomy in a variety of ways. User input can be blended with the assumed optimal action, obtaining an average action at every time step (83). User input can be filtered using a task-based criterion, and if sufficiently suboptimal, it can be either ignored or replaced with an assumed optimal action (84). Intervention could be conditioned on the certainty of the autonomy in its prediction of the task goal (85, 86). If the robot is operated through mode switching (explained in more detail in Section 3.5), a shared control paradigm may be designed to assist with toggling between modes based on a prediction of the most likely mode (69, 87).

While many approaches vary assistance to optimize performance, others provide assistance based solely on safety. The latter strategy has two significant advantages: (*a*) It limits interference, relinquishing control to the human operator, and (*b*) it is often easier to define—in a human–machine system, there are usually fewer safety constraints compared with all possible task goals. Even so, the optimal shared control paradigm is often unique to the person and the application (88), highlighting the need for solutions that enable adaptation. Bidirectional communication can help a person and a robot to most flexibly agree on an optimal shared control strategy and to renegotiate that agreement throughout the course of use. We discuss this idea in more detail in Section 4.

2.5. Adaptation

Adaptable control interfaces show promise in their ability to improve performance and safety, as well as to increase user satisfaction. The robot can adapt (*a*) its predictions or model of its human partner, (*b*) the mapping (or interpretation) of human-generated signals onto predicted intent, and (*c*) the shared control paradigm, adjusting its intervention strategy based on user capabilities and individual preferences. Unlike the majority of work described above, adaptation assumes that how people start using a device will not necessarily be the same across individuals and that how they use the device will change over the course of use.

Many researchers have investigated options to personalize algorithms for pHRI to accommodate biomechanical and physiological differences between individuals. A common example is the customization of gait in an exoskeleton (89, 90), usually through an initial calibration period of walking without assistance. Another example is the calibration of EMG (13, 37, 38, 43, 91) or EEG (92, 93) mappings to an individual's signal pattern or to the specific placement of the electrodes on the body. A personalized remapping has also been shown to be beneficial for more established input devices, such as a joystick, when the human operator suffers from physical limitations due to a neuromotor impairment (94).

In addition to personalization based on physiology, biomechanics, and physical capabilities, users value the ability to customize an interface based on preferences. As an example, individuals might exhibit preferences for different shared control paradigms (e.g., those that are less aggressive) irrespective of the objective performance benefits (e.g., time to task completion) (95). As a result, some work utilizes the user's control behavior to automatically vary parameters of the shared control paradigm (13, 96). While preferences vary between participants, there is a general trend favoring paradigms that retain stronger autonomy for the human operator (95). People prefer interfaces that are intuitive to use and transparent in how they work (97).

Initial adaptation to the individual is important, but it is also important to recognize that individuals' performance and capabilities change over time, due to factors such as learning, fatigue, or disease progression. In particular, in a study with a smart wheelchair, people's performance improved significantly between just two experimental sessions (95), a trend that was observed across different control interfaces and shared autonomy paradigms. Another study showed that novices and experts prefer different shared control paradigms (58). These results suggest that the shared control approach should not be static over time and that continual adaptation of the communication paradigm can benefit the human–machine system. While few research studies have evaluated the adaptation of pHRI interfaces over time, continual adaptation is an important research direction for pHRI, as discussed in Section 4.

For a summary of the key points regarding algorithmic approaches, see the sidebar titled Algorithmic Approaches: A Summary.

ALGORITHMIC APPROACHES: A SUMMARY

1. Over the past decades, algorithmic approaches that facilitate pHRI have progressed from algorithms based on physical analogies, such as admittance and impedance control, to computational methods based on higher-level reasoning.
2. Intent inference is often multisensory, taking advantage of technologies that measure biomechanical and bioelectrical activity as well as explicit human commands to inform robot actions.
3. Existing algorithmic paradigms (signal processing, optimal control, machine learning, and more) can be customized to the robotic systems and personalized to the individual, but most do not adapt beyond an initial calibration, limiting the human–robot collaborative potential.
4. While current algorithms successfully enable task-specific interactions, future progress can expand robotic applications from, e.g., semiautomated manufacturing and precision surgery to versatile at-home assistance.

3. APPLICATIONS OF PHYSICAL HUMAN–ROBOT INTERACTION

pHRI plays an important role in a number of applications, ranging from robot-assisted rehabilitation to collaborative assembly during manufacturing. Here, we describe how physical robotic assistance is being used in different settings. We comment on state-of-the-art performance and discuss the current challenges associated with human–machine communication in each application area. Example robotic platforms that have been developed for these applications are shown in **Figure 3**.

3.1. Robot-Assisted Rehabilitation

One of the major applications of pHRI focuses on robot-assisted physical therapy (98, 99). Robot-assisted rehabilitation offers the promise of therapeutic regimens that are more effective, more accessible, and more engaging than traditional therapy. To date, results have been mixed on the impact of robot-assisted training on clinical outcome measures. While some studies have shown benefits of adaptive robot-assisted training compared with conventional therapy (100, 101), others have shown no statistically significant differences in clinical outcomes (102). However, the evidence is strong that there are benefits from patients' active participation in therapy (103). These positive findings encourage research that promotes patient engagement through reliable detection of intent and real-time metrics of performance and effort.

Figure 3

Example robotic platforms: (*a*) Kinova MICO arm, (*b*) Quickie Q500 M powered wheelchair, (*c*) Ekso Bionics lower-limb exoskeleton, (*d*) Shirley Ryan AbilityLab upper-limb prosthesis, (*e*) Open-Source Leg, and (*f*) ACT3D (a modified HapticMaster end-effector robot). Panel *c* reproduced with permission from Ekso Bionics.

Research solutions attempt to maximize patient engagement using two approaches: (*a*) assist-as-needed paradigms and (*b*) measures of patient intent that guide robotic assistance (99). Assist-as-needed strategies use adaptive controllers that estimate patient effort and assess performance in real time, enabling adjustments to the level of assistance even during a single movement (104). As an example, in 2017, the HAL (Hybrid Assistive Limb) rehabilitation exoskeleton, which uses EMG signals to adjust assistance, was officially approved by the US Food and Drug Administration to enter the US medical rehabilitation market. In parallel, intent detection—e.g., using biometrics from EMG or EEG—can further involve the patient by providing them agency over generating movement trajectories (105). Intent prediction methods are particularly important for engaging more severely impaired patients in robotic training, because these patients have difficulty with independently generating voluntary movement (106).

Work has also been done on developing impairment-specific assessments of motor abilities (107–109), with some metrics showing potential for real-time use (67). With new algorithmic solutions, it should be possible to design patient-specific training protocols, dependent upon each patient's type of injury, level of impairment, and phase of recovery. Embodied communication could play a role in designing interactive paradigms for robot-assisted assessment of physical deficits. Consequently, interactive methods could enable more accurate diagnosis of motor capabilities and lead to improved therapeutic efficacy through adaptive, impairment-specific rehabilitation protocols.

For individuals with physical limitations, the ultimate goal is for robotic devices to provide both therapy and assistance (110). Some devices are already attempting to show efficacy for combined assistance and at-home rehabilitation (111). With robot-based continuous metrics of impairment, person-specific rehabilitation protocols can be more effectively incorporated into activities of daily living.

3.2. Physical Assistance with Exoskeletons

Exoskeletons—portable devices that are physically coupled with an individual—are the most commonly considered type of robot for physically assisting with human movement. Many exoskeletons have been developed to augment people's physical capabilities for both impaired and able-bodied individuals (75, 112, 113). Most rigid exoskeletons are controlled in admittance mode using force amplification, and for some applications, such as level-ground walking, this control mode works reasonably well (112). In fact, for level-ground walking, it is challenging to do better than admittance control in terms of movement freedom and flexibility. However, while task-agnostic assistance performs well for walking, it does not provide helpful support during unexpected events (e.g., tripping) or changing activities (e.g., climbing stairs or transitioning to sitting) (112). Alternatively, exoskeletons rely on a state-machine controller, where an initiation event—e.g., a button press or a specific motion—is used to communicate movement intent (114). Other task-based controllers provide assistance through a sequence of learned movement trajectories, giving the person little control over motion execution (114).

While lower-limb exoskeletons have received substantial academic and commercial attention, their practical adoption has been limited. Traditional rigid exoskeletons (75, 115) are bulky and expensive. While they offer assistance with generation of gait, even for individuals with full paralysis (114), they do not guarantee assistance with balance and hence require crutches or therapist assistance during use. More recently, lightweight designs have shown promising results for able-bodied individuals, reducing the physical effort required to walk long distances (116) or increasing the ability to carry heavy loads (117).

In parallel, there is ongoing research on the design of upper-limb exoskeletons (118). Although upper-limb devices do not have to assist with balance, their task space is larger and more varied.

The hardware requirements (e.g., weight, ergonomics, and biomechanical compatibility) are more rigorous than those of lower-limb exoskeletons, and the application areas are broad. As a result, the research is at an earlier stage, focusing largely on the hardware design of these devices (118). Most of the commercially available exoskeletons are application specific and offer limited assistance. For example, the SKELEX exoskeleton from GOBIO can nullify the weight of a tool while working on the shop floor.

In an exoskeleton, the physical coupling of a human and a machine brings both a unique opportunity for continuous communication [as in partner dancing (119)] and the challenge of arbitration between human intent and robotic assistance. Communication takes place through subtle haptic cues and concurrent negotiation of movement. Interestingly, human–human physical coupling can reliably lead to improved performance even if one of the partners is less adept at the task (120). Studies on human–robot coupling show that forceful interaction can lead to improved task performance even after the robot coupling is removed and the person performs the task independently (121). These pieces of evidence illustrate the potential of communication through physically coupled motion. However, while physical interaction is a natural part of joint activity between people, how humans control motor interaction with peers is still largely unknown (122). Human–robot interaction (e.g., via exoskeletons) can help us discover the potential of physical communication and in turn help improve algorithms for human–machine collaboration.

3.3. Physical Augmentation Through Prosthetics

Powered prostheses offer the promise of restoring lost functionality to individuals with a missing limb. With recent progress in reducing the weight of the physical devices, the prostheses' usability depends largely on the effectiveness of the algorithmic solutions. Commercially available lower-limb powered prostheses (e.g., the Empower prosthetic foot from Ottobock) use onboard sensing (e.g., a load cell or motion sensor) to drive autonomous controllers for preprogrammed activities (32). Transitions between gait phases can be triggered by sensor measurements, while transitions between locomotion modes (e.g., overground walking or stair climbing) often require input from the human user (e.g., specific body motions measured by sensors) (32). Some research solutions consider incorporating contextual information (e.g., a laser distance meter) (123) to improve decision-making by detecting terrain, and some explore incorporating EMG signals to improve prediction of user intent (32). Existing control techniques can already achieve decent performance; for example, most lower-limb amputees feel confident walking forward on level ground (124), although maintaining balance while walking on uneven terrain or on slopes remains a concern (124).

For the upper limb, the task space is more varied and open-ended, increasing the difficulty of control. The main challenge is obtaining a rich enough control signal from the human partner to infer their movement intent without interrupting the person's workflow. Intent prediction using EMG has received the most attention (13, 38, 125), with EEG being a close second. While EMG signals are usually stronger and more discernible than EEG signals, amputees might have limited muscle fibers in their residual limb to provide a robust EMG readout. To increase the strength and clarity of the EMG signals in the residual limb, two clinical procedures have shown promising results. The first is targeted motor and sensory reinnervation (126, 127). During this type of procedure, a surgeon reroutes nerves from the residual limb to a large muscle group, such as pectoral muscles, restoring the number of intuitive EMG channels and sensory inputs available to the individual. The second and more recent procedure is to implant electrodes into the muscle and add a metal anchor into the bone of an amputated limb, which then enable control of osseointegrated prosthetics via intramuscular EMG (128). These surgical advancements significantly increase the capacity of the human–prosthesis communication channel, improving usability. Even so, control of a powered upper-limb prosthesis is for now limited, and outside of laboratory settings, the prosthetic arm is usually used passively as a support arm in bimanual tasks (e.g., holding a bowl while stirring with the other hand) (129).

A relatively new research effort focuses on restoring sensing capabilities to the individual through additional feedback. Currently, there are almost no commercially available powered prostheses that transfer sensations to the user (124). Interestingly, some amputees find mechanical prostheses more intuitive to use because the more rigid devices mechanically transmit haptic feedback from the prosthesis tip to the residual limb. Osseointegrated prostheses share this benefit (128). In socket-based powered prostheses, researchers are considering alternatives for providing sensory feedback (13, 130, 131). As an example, pressure stimulation in a prosthetic socket is being used to communicate forces perceived by the hand (131). Systems leveraging targeted sensory reinnervation (127) can increase the amputee's capacity for sensory feedback as well as make the sensory feedback more intuitive.

The human–prosthesis interaction can also benefit from communication to the user about the current state or movement intent of the prosthesis (70). As an example, haptic feedback about the opening or closing of a prosthetic hand has been shown to lead to a lesser need for visual attention and improved task performance (132). Rich sensing and reciprocal communication of both the device's perception and its intent are expected to improve the effectiveness of the human–prosthesis interaction. While powered prostheses are meant to be natural extensions of the human body, there is a desire to seamlessly integrate the artificial limb into the person's workflow without the need for explicit, cognitively taxing communication of intent. In this viewpoint, powered prostheses remain a challenging domain for inferring intent, especially in contextually rich settings during activities of daily living. Emergent multimodal dialogue for pHRI is an important area of research that could improve performance, as we discuss in Section 4.

3.4. Collaboration in Close Proximity with an Autonomous Robot

Some assistive robots are stand-alone devices that operate autonomously in close proximity to a human partner. Physical interaction during this type of collaboration is arguably one of the more difficult pHRI scenarios, because the human and robot must be able to work together while (*a*) mostly avoiding each other and (*b*) intermittently seeking out safe physical contact. In the first scenario, physical contact is an unintended by-product of working in close proximity. A large body of research focuses on intent prediction for avoiding accidental interaction while working in close proximity with a human partner (133, 134) or for avoiding collisions when navigating in crowded environments (135, 136). In the second scenario, a common focus area is interacting with a robot by jointly manipulating objects in the workspace. This can involve collaboratively moving a table (137, 138) or handing objects between the human and the robot (139). These studies have made progress on relevant pHRI challenges, yet they often explicitly separate the periods of physical interaction and autonomous operation. Incorporating intermittent pHRI as a natural component of collaborative interaction will be key to the emergence of complex, rich, and flexible collaboration.

A commonly studied form of prescribed human–robot contact involves kinesthetic demonstrations, where the human physically manipulates the robot to demonstrate desired movements. Kinesthetic demonstrations provide an intuitive avenue for people to teach robots new skills or customize assistance without having to program reward functions or low-level behaviors (140). As an example, learning-from-demonstration paradigms have been successfully used for teaching a robot arm how to perform manipulation tasks (141), such as drinking or pouring liquids and moving objects. Many studies have shown that a learned policy can generalize to new environmental conditions (142, 143), such as new locations of objects, similar objects, and clutter. However, fewer works have demonstrated a generalization of learned skills between different tasks. Without the ability to reuse motor skills for novel goal-oriented activities, the learning solutions are less applicable to a versatile assistive robot. The question remains of how to algorithmically structure the robot's interaction with a person to enable the robot to effectively extract task-relevant information from demonstrations (140).

New research is exploring interactive learning, with the goal of improving learning of salient task elements by actively soliciting clarifying input from the human teacher. Some approaches ask users to rank demonstrations (144, 145) or provide corrective input on the learned task executions (146, 147). Other studies incorporate demonstrations of what *not* to do, enabling the robot to extract information from failed attempts at a task (148). Additional feedback often increases task performance, but it does so at the cost of an individual's time and effort. Even though in their current implementation the interactive paradigms might be inefficient, they are a step toward bidirectional communication that will enable more versatile human–robot collaboration.

3.5. Teleoperation of a Semiautonomous Robot

Robotic teleoperation has enabled robotic surgery, space exploration, assistance for physically impaired individuals, and more. Here, we focus on teleoperation as a way to control a colocated collaborative device (e.g., a powered wheelchair or wheelchair-mounted robot arm) rather than a remote robot (e.g., a Mars rover or the da Vinci surgical system).

A semiautonomous assistive device, such as a smart powered wheelchair, can be successfully controlled by a handheld joystick (56). In this case, control is near trivial because a 2D joystick intuitively maps to a 2D control space. Direct control of the robot is more difficult when the controller is lower-dimensional (e.g., a sip-n-puff, as described in Section 2.2.4) or the robot is more complex (e.g., a robot arm with multiple degrees of freedom), because low-dimensional input must be used to control high-dimensional motion. Mode switching—toggling between independent and typically orthogonal control directions, which usually span translation and rotation in the *xy*- or *xyz*-coordinate space—has been proposed as a possible solution (149). Autonomous switching of modes can speed up task performance (87). However, control via single degrees of freedom in a Cartesian coordinate space is an unintuitive approach to specifying robot movement, because people do not think about motion in terms of 1D adjustments. Instead, they tend to plan trajectories in terms of functional movements or motion primitives (150).

In recent work, researchers have proposed using latent variables to map control inputs onto task-specific movement trajectories (86, 151). These approaches offer a way to map control inputs onto more natural and functionally relevant movements of the robot. In their current implementation, these approaches require a task library and a high-level controller to determine the applicable latent space for the task, but this is a promising research direction as it enables learning a communication manifold from human–robot interaction. Algorithmically structuring the interaction, in a way that will effectively enable these mappings to emerge, is an important future research direction.

4. FUTURE DIRECTIONS FOR EMBODIED COMMUNICATION IN PHYSICAL HUMAN–ROBOT INTERACTION

Future work on pHRI will benefit from greater emphasis on embodied communication to exploit the richness of the available communication channels both for human–robot communication and, increasingly, reciprocal robot–human communication. Aspirations for future pHRI systems

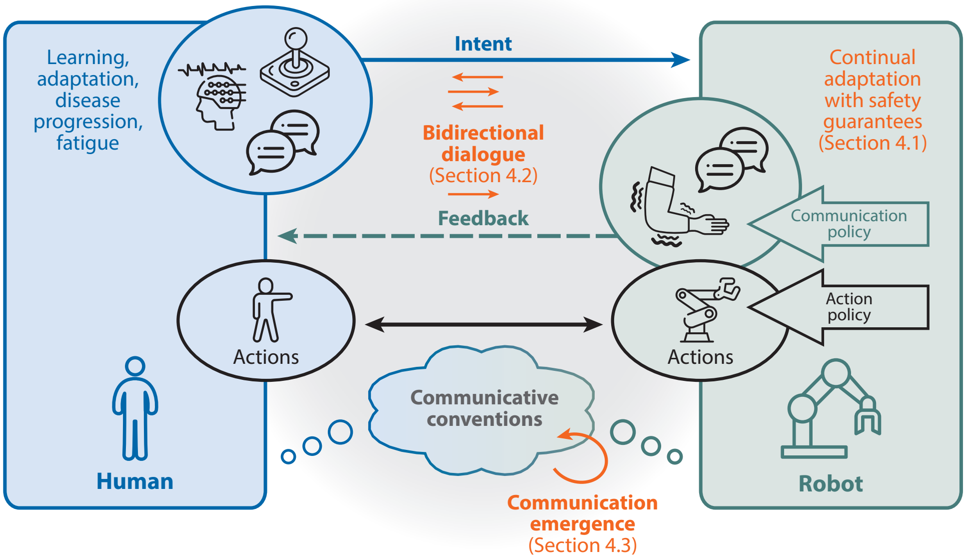

Figure 4

Future directions (*orange*) for embodied communication during effective human–robot collaboration. In existing human–robot partnerships, the robot infers a person's intent and interacts with them through assistive actions, sometimes providing limited feedback through additional signaling. Even though the person changes over time—due to learning, fatigue, etc.—the communicative conventions remain the same throughout use. Developing new algorithmic tools to facilitate human–machine dialogue and enable the emergence of communicative conventions without compromising safety is one of the next frontiers for the field of physical human–robot interaction (pHRI).

include safe operation enforced by the autonomy while allowing continual adaptation, multimodal dialogue, and support for emergent, unanticipated behavior and communication during collaboration (as shown in **Figure 4**). These goals anticipate that the person and the environment change over time and that the human–robot pair will need to coadapt to each other. They also have something in common technically—they lack standardized algorithmic tools, which are currently in development in terms of theory and implementation, as we discuss below. As a result, datasets are needed for validating pHRI-relevant algorithmic tools, benchmarking algorithms prior to use with people, and comparing and evaluating current pHRI capabilities.

4.1. Continual Adaptation with Safety Guarantees

pHRI involves potential safety hazards because of the direct mechanical contact and energy exchange. As a result, safety has been a core focus of facilitating pHRI (152). Even so, it is not clear how to guarantee safety without limiting parameters of robot motion (e.g., torque or velocity) and bounding the robot's performance (152, 153). Due to the lack of reliable alternatives, ISO guidelines released in 2016 suggest just that—they recommend quantitative biomechanical limits, such as allowable peak forces or pressures for various parts of the body, as a requirement for collaborative robots (154).

Recent work takes advantage of predictive modeling to ensure safety. Folkestad et al. (155) used control barrier functions and Koopman operators to impose data-driven safety guarantees.

Other approaches involve modeling and closely monitoring people's behavior in anticipation of unsafe robot motion. As an example, researchers have shown the utility of hidden Markov models for estimating, in real time, human affective states during pHRI (156) and proposed using this approach as a feedback mechanism to prevent unsafe interactions. Separately, Brown et al. (157) formalized and theoretically analyzed safety in the context of efficient value alignment, with the goal of constructing a kind of driver's test that a human can give to a robot for assessment prior to use. These efforts are a step toward providing standardized safety guarantees (under reasonable assumptions) for data-driven robotic systems.

In Section 2.5, we discussed the value of adaptation and data-driven methods for algorithmic regulation of human–robot interaction. Notably, most work has focused on learning within a confined period of initial interaction, and few studies have considered continual or long-term adaptation. One reason for this is that although it is challenging to assess the safety and performance of data-driven methods, it is even more challenging to do so for systems that are continually adapting. Developing methods for data-driven verification of safety while enabling long-term learning is an important direction for future work. As discussed by Kress-Gazit et al. (158), we should aim to provide safety guarantees that at minimum maintain states within safe sets and satisfy temporal logic guarantees while also providing specifications that are appropriate for nonstationary models and reason about human variability and adaptation.

4.2. Multimodal Dialogue

Just as we seek to infer informative measurement signals from the human for the robot, the robot should seek to provide informative signals to the person (181). Many human–robot systems rely on the robot movement as implicit feedback for the human partner. Some work augments the interaction with exaggerated robot movements (10, 159, 160) or anthropomorphic signaling (161). Purposefully communicative motion has been shown to benefit the interaction, because it makes robot behavior more predictable, allowing the person to anticipate robot movements and adapt their own behavior accordingly (162). If robot motion is predictable and allows humans to adapt—e.g., in the way that humans swiftly adapt to robot motion in crowd navigation (135)—then safety constraints on the robot might be relaxed and enable more effective collaboration.

Other work has taken advantage of communication through modes other than motion to provide explicit feedback about robot intent. For instance, recent studies have used vibrotactile feedback in the form of a vibrating sleeve (132), pressure spots in a prosthetic socket (131), or friction modulation on the surface of a touch screen (163). Other studies have explored using vibrotactile feedback to communicate emotion, imitating social touch with patterns of haptic sensations on the arm (164, 165), or to communicate learned predictions of upcoming hazards (166, 167). Without disturbing the task flow, vibrotactile stimulation and haptic cues can significantly increase the information communicated to the person by the machine, improving the pair's collaborative potential.

As discussed throughout the review, one-way signaling about intent—from the person to the robot—has enabled task-specific robotic assistance. We described how most work on communication in pHRI treats the robot as a passive observer of human behavior and, at best, enables the robot to provide feedback about its current state and/or intent. We discussed evidence showing that reciprocal signaling from the robot to the person overwhelmingly improves the interaction (68, 70, 132, 168). Though bidirectional, thus far this type of communication has enabled negotiating task solutions or refining mutual understanding and alignment over time. Bidirectional dialogue—where communication is dependent on both the mental state of the communication partner and the communication history—could enable more versatile collaboration.

4.3. Emergent Interfaces

The model of signaling described above assumes that the communicating agent (human or robot) has a correct mental model of the communicative conventions, and this assumed mental model is typically static. If we are to move toward active nonverbal dialogue between the human and the machine, novel nonverbal languages will form a foundation for communication. As described by Dominijanni et al. (169), we can capitalize on our kinematic, muscular, and neural null space—excess degrees of freedom—to enable communication with assistive devices without causing excessive increases in cognitive load. The question remains how to facilitate efficient neural resource allocation for a novel human–robot system (169).

As a baseline approach, the languages can be predefined and encoded in the robot—the programmer determines a set of vectors in the neural null space, and the human operator is instructed in their interpretation and trained to generate them. However, how can functionally optimal or near-optimal symbols be specified? And how can the cognitive burden on the human partner—who is expected to learn the newly introduced vocabulary of nonintuitive symbols—be reduced? If we are mapping joystick commands to the action space of a powered wheelchair, there is an intuitive way to specify the mapping of the continuous 2D control space onto the analogous 2D action space of the robot, making the language easy to learn. However, if we are using shoulder shrugs to control movements of a robot arm (170) or sips and puffs to maneuver a wheelchair (58), the mapping is no longer as intuitive, either to specify or to learn, and may not even be within the motor control capacity of a given individual.

Interestingly, experiments have shown that humans are good at adapting to and developing novel communication protocols (171). For instance, in an experimental setup with an unfamiliar task (172), people learned how to interpret vibrotactile stimulation without being given a description of what the stimulation is intended to evoke. In the experiment, the vibrotactile stimulation was synthesized by an optimal controller (rather than embodying a state measurement), and the context was sufficient for participants to infer the meaning of the vibrotactile signals—subjects successfully learned to use the stimulation as a cue for motor response to improve performance. Such anecdotal evidence illustrates potential for fluent coadaptation of the human–robot pair, on fast timescales and in an individualized manner.

There is a body of work that studies how to enable and facilitate communication emergence in autonomous agents (173), with recent studies exploring the role of deep reinforcement learning (174, 175). These algorithmic approaches to facilitating collaborative development of nonverbal languages have the potential to form the foundation for humans and robots to cooperate flexibly (176) and to continuously negotiate their partnership. Enabling the human–robot pair to jointly develop a communication protocol could reduce the cognitive burden on the human partner and enable more effective communication protocols that are uniquely relevant to the human–machine system and to the corresponding task space. However, unlike the anecdote mentioned above, the deep learning algorithms are data intensive, requiring thousands or millions of interactions in order to achieve a nonverbal language. The opportunity is that these approaches assume naive agents that know little about their environment or about each other. Incorporating these approaches into collaborative language creation between a human and a robot is the next step in bringing these simulation-based studies closer to the fast, individualized physical communication needed in human–machine interaction.

4.4. Datasets and Benchmarking

Acquiring human data is costly in terms of the time and effort required from both the researchers and the study participants, particularly in experiments with vulnerable populations. As a result, research studies have focused on evaluating novel algorithmic solutions, and few have carried out direct comparisons with existing approaches. To accelerate progress, sharing datasets of benchmark tasks would be beneficial. Although statements that data are available upon request are common, they have been shown to be inefficient (177). This suggests that the pHRI community needs open datasets and standards that define them, in terms of both standardized sharing practices and standardized data formats.

Some researchers have already initiated this practice. A recent paper featured a dataset of kinematic and EMG signals collected during reach-to-grasp movements with online adjustments in response to visual perturbations (178). Datasets like this one, which record human kinematics in response to environmental stimuli, can be used for training robotic controllers that aim to mimic human behaviors. Similarly, a multimodal dataset of assistive human–robot collaboration has been released, including eye tracking, EEG, EMG, camera images, joystick signals, and more (179). Although these data cannot substitute for an experimental evaluation of a novel algorithm because of the interactive aspect of online control, they can enable a baseline comparison of task performance. Sharing data is of particular importance in fields such as pHRI, where active participation of human volunteers is required to collect data.

While the benefits of sharing data are clear, wide adoption and reliable benefits will require us to create and implement community standards. Such standards may require a standardization of tasks, which could be similar to clinical assessments of impairment (180)—sets of well-defined minitasks that are scored independently and summed to provide an estimate of motor deficit. The standards will also need to focus on ease of use—e.g., sharing raw anonymized data along with Jupyter Notebooks for server-side visualization, or making code available for reproducing statistics used in the original study. Sharing data can improve benchmarking, facilitate collaborations, and accelerate progress in designing machines that can effectively interact and communicate with people.

5. CONCLUSIONS

In current applications of human–robot collaboration, pHRI is largely avoided or completely prescribed. The robot is often treated as a passive observer of human intent; in some scenarios, it can provide reciprocal feedback. While this type of embodied communication has enabled successful collaboration within the constraints of a specific task, effective collaboration in unanticipated settings will require human–machine dialogue—comprising a multimodal exchange of information based on a mental model of the collaborative partner and interaction history—and safely incorporating intermittent physical contact. As such, developing new algorithmic tools to facilitate multimodal dialogue and enable the emergence of communicative conventions is one of the next frontiers for pHRI. If we achieve flexible dialogue between the human–robot pair, human–robot collaboration will become possible in currently impractical applications, such as coexistence in the home or continuous rehabilitation during robot-assisted execution of daily activities.

FUTURE DIRECTIONS

1. Continual adaptation with safety guarantees: With the growing use of data-driven methods that learn over time, there is an increased need for provable safety bounds that do not unnecessarily constrain the human–robot interaction or limit overall performance.
2. Multimodal dialogue: While many existing solutions rely on the robot to infer intent from communicative signaling generated by its human partner, and some enable

reciprocal signaling from the robot, continual bidirectional communication could enable more versatile collaboration.

3. Emergent interfaces: While it is difficult to prespecify a comprehensive multimodal vocabulary of symbols, a human–robot pair could gradually expand its communicative capacity by establishing relevant conventions through interaction.
4. Datasets and benchmarking: With a growing number of algorithmic approaches and the costliness of comparison-based evaluations, robust benchmarks and reliable data sharing could speed up innovation while reducing experimental overhead and barriers to entry.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grant CNS-1837515. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. P.M.P. was supported in part by a Canada CIFAR AI Chair (Amii).

LITERATURE CITED

1. Haddadin S, Croft E. 2016. *Physical Human–Robot Interaction*. Cham, Switz.: Springer
2. Mathewson KW, Parker ASR, Sherstan C, Edwards AL, Sutton RS, Pilarski PM. 2022. Communicative capital: a key resource for human-machine shared agency and collaborative capacity. *Neural Comput. Appl.* In press. **https://doi.org/10.1007/s00521-022-07948-1**
3. Scott-Phillips T. 2014. *Speaking Our Minds: Why Human Communication Is Different, and How Language Evolved to Make It Special*. London: Bloomsbury
4. Sebanz N, Bekkering H, Knoblich G. 2006. Joint action: bodies and minds moving together. *Trends Cogn. Sci.* 10:70–76
5. Knoblich G, Butterfill S, Sebanz N. 2011. Psychological research on joint action: theory and data. *Psychol. Learn. Motiv.* 54:59–101
6. Pezzulo G, Dindo H. 2011. What should I do next? Using shared representations to solve interaction problems. *Exp. Brain Res.* 211:613–30
7. Bicho E, Erlhagen W, Louro L, Costa e Silva E. 2011. Neuro-cognitive mechanisms of decision making in joint action: a human–robot interaction study. *Hum. Mov. Sci.* 30:846–68
8. Pezzulo G, Donnarumma F, Dindo H. 2013. Human sensorimotor communication: a theory of signaling in online social interactions. *PLOS ONE* 8:e79876
9. Wolpert DM, Doya K, Kawato M. 2003. A unifying computational framework for motor control and social interaction. *Philos. Trans. R. Soc. Lond. B* 358:593–602
10. Moon A, Parker CA, Croft EA, Van der Loos HM. 2013. Design and impact of hesitation gestures during human-robot resource conflicts. *J. Hum.-Robot Interact.* 2:18–40
11. Hogan N. 1985. Impedance control: an approach to manipulation: part I—theory. *J. Dyn. Syst. Meas. Control* 107:1–7
12. Gopinath D, Jain S, Argall BD. 2016. Human-in-the-loop optimization of shared autonomy in assistive robotics. *Robot. Autom. Lett.* 2:247–54

13. Castellini C, Artemiadis P, Wininger M, Ajoudani A, Alimusaj M, et al. 2014. Proceedings of the first workshop on peripheral machine interfaces: going beyond traditional surface electromyography. *Front. Neurorobot.* 8:22
14. Whitney DE. 1977. Force feedback control of manipulator fine motions. *J. Dyn. Syst. Meas. Control* 99:91–97
15. Marchal-Crespo L, Reinkensmeyer DJ. 2009. Review of control strategies for robotic movement training after neurologic injury. *J. NeuroEng. Rehabil.* 6:20
16. Newman WS. 1992. Stability and performance limits of interaction controllers. *J. Dyn. Syst. Meas. Control* 114:563–70
17. Keemink AQ, van der Kooij H, Stienen AH. 2018. Admittance control for physical human–robot interaction. *Int. J. Robot. Res.* 37:1421–44
18. He W, Xue C, Yu X, Li Z, Yang C. 2020. Admittance-based controller design for physical human–robot interaction in the constrained task space. *IEEE Trans. Autom. Sci. Eng.* 17:1937–49
19. Chen X, Wang N, Cheng H, Yang C. 2020. Neural learning enhanced variable admittance control for human–robot collaboration. *IEEE Access* 8:25727–37
20. Abdallah M, Chen A, Campeau-Lecours A, Gosselin C. 2022. How to reduce the impedance for pHRI: admittance control or underactuation? *Mechatronics* 84:102768
21. van der Linde RQ, Lammertse P. 2003. HapticMaster – a generic force controlled robot for human interaction. *Ind. Robot* 30:515–24
22. Patton JL, Stoykov ME, Kovic M, Mussa-Ivaldi FA. 2006. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Exp. Brain Res.* 168:368–83
23. Fan T, Alwala KV, Xiang D, Xu W, Murphey T, Mukadam M. 2021. Revitalizing optimization for 3D human pose and shape estimation: a sparse constrained formulation. In *2021 IEEE/CVF International Conference on Computer Vision*, pp. 11437–46. Piscataway, NJ: IEEE
24. Williams HE, Chapman CS, Pilarski PM, Vette AH, Hebert JS. 2019. Gaze and Movement Assessment (GaMA): inter-site validation of a visuomotor upper limb functional protocol. *PLOS ONE* 14:e0219333
25. Rudenko A, Palmieri L, Herman M, Kitani KM, Gavrila DM, Arras KO. 2020. Human motion trajectory prediction: a survey. *Int. J. Robot. Res.* 39:895–935
26. Letham B, Rudin C, Madigan D. 2013. Sequential event prediction. *Mach. Learn.* 93:357–80
27. Kalinowska A, Berrueta TA, Zoss A, Murphey T. 2019. Data-driven gait segmentation for walking assistance in a lower-limb assistive device. In *2019 International Conference on Robotics and Automation*, pp. 1390–96. Piscataway, NJ: IEEE
28. Pérez-D'Arpino C, Shah JA. 2015. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *2015 IEEE International Conference on Robotics and Automation*, pp. 6175–82. Piscataway, NJ: IEEE
29. Lasota PA, Shah JA. 2017. A multiple-predictor approach to human motion prediction. In *2017 IEEE International Conference on Robotics and Automation*, pp. 2300–7. Piscataway, NJ: IEEE
30. Pilarski PM, Dawson MR, Degris T, Carey JP, Chan KM, et al. 2013. Adaptive artificial limbs: a real-time approach to prediction and anticipation. *Robot. Autom. Mag.* 20(1):53–64
31. Liu H, Wang L. 2018. Gesture recognition for human-robot collaboration: a review. *Int. J. Ind. Ergon.* 68:355–67
32. Fleming A, Stafford N, Huang S, Hu X, Ferris DP, Huang HH. 2021. Myoelectric control of robotic lower limb prostheses: a review of electromyography interfaces, control paradigms, challenges and future directions. *J. Neural Eng.* 18:041004
33. Veeriah V, Pilarski PM, Sutton RS. 2016. *Face valuing: training user interfaces with facial expressions and reinforcement learning*. Paper presented at the Interactive Machine Learning Workshop, 25th International Joint Conference on Artificial Intelligence, New York, NY, July 9
34. Palinko O, Rea F, Sandini G, Sciutti A. 2016. Robot reading human gaze: why eye tracking is better than head tracking for human-robot collaboration. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5048–54. Piscataway, NJ: IEEE
35. Admoni H, Scassellati B. 2017. Social eye gaze in human-robot interaction: a review. *J. Hum.-Robot Interact.* 6:25–63

36. Rastgoo R, Kiani K, Escalera S. 2021. Sign language recognition: a deep survey. *Expert Syst. Appl.* 164:113794
37. Shehata AW, Williams HE, Hebert JS, Pilarski PM. 2021. Machine learning for the control of prosthetic arms: using electromyographic signals for improved performance. *IEEE Signal Process. Mag.* 38(4):46–53
38. Micera S, Carpaneto J, Raspopovic S. 2010. Control of hand prostheses using peripheral information. *Rev. Biomed. Eng.* 3:48–68
39. Lenzi T, De Rossi SMM, Vitiello N, Carrozza MC. 2011. Proportional EMG control for upper-limb powered exoskeletons. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 628–31. Piscataway, NJ: IEEE
40. Bi L, Feleke AG, Guan C. 2019. A review on EMG-based motor intention prediction of continuous human upper limb motion for human-robot collaboration. *Biomed. Signal Process. Control* 51:113–27
41. Hargrove LJ, Simon AM, Young AJ, Lipschutz RD, Finucane SB, et al. 2013. Robotic leg control with EMG decoding in an amputee with nerve transfers. *N. Engl. J. Med.* 369:1237–42
42. Chowdhury RH, Reaz MB, Ali MABM, Bakar AA, Chellappan K, Chang TG. 2013. Surface electromyography signal processing and classification techniques. *Sensors* 13:12431–66
43. Hargrove LJ, Miller LA, Turner K, Kuiken TA. 2017. Myoelectric pattern recognition outperforms direct control for transhumeral amputees with targeted muscle reinnervation: a randomized clinical trial. *Sci. Rep.* 7:13840
44. Resnik L, Huang HH, Winslow A, Crouch DL, Zhang F, Wolk N. 2018. Evaluation of EMG pattern recognition for upper limb prosthesis control: a case study in comparison with direct myoelectric control. *J. NeuroEng. Rehabil.* 15:23
45. Kristoffersen MB, Franzke AW, Van Der Sluis CK, Murgia A, Bongers RM. 2020. Serious gaming to generate separated and consistent EMG patterns in pattern-recognition prosthesis control. *Biomed. Signal Process. Control* 62:102140
46. Lalitharatne TD, Teramoto K, Hayashi Y, Kiguchi K. 2013. Towards hybrid EEG-EMG-based control approaches to be used in bio-robotics applications: current status, challenges and future directions. *Paladyn J. Behav. Robot.* 4:147–54
47. DelPreto J, Salazar-Gomez AF, Gil S, Hasani R, Guenther FH, Rus D. 2020. Plug-and-play supervisory control using muscle and brain signals for real-time gesture and error detection. *Auton. Robots* 44:1303–22
48. Cipriani C, Zaccone F, Micera S, Carrozza MC. 2008. On the shared control of an EMG-controlled prosthetic hand: analysis of user-prosthesis interaction. *IEEE Trans. Robot.* 24:170–84
49. Al-Quraishi MS, Elamvazuthi I, Daud SA, Parasuraman S, Borboni A. 2018. EEG-based control for upper and lower limb exoskeletons and prostheses: a systematic review. *Sensors* 18:3342
50. McMullen DP, Hotson G, Katyal KD, Wester BA, Fifer MS, et al. 2013. 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.* 22:784–96
51. Herff C, Krusienski DJ, Kubben P. 2020. The potential of stereotactic-EEG for brain-computer interfaces: current progress and future directions. *Front. Neurosci.* 14:123
52. Wolpaw JR, McFarland DJ. 2004. Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *PNAS* 101:17849–54
53. Jeong JH, Shim KH, Kim DJ, Lee SW. 2020. Brain-controlled robotic arm system based on multi-directional CNN-BiLSTM network using EEG signals. *IEEE Trans. Neural Syst. Rehabil. Eng.* 28:1226–38
54. Millán JDR, Renkens F, Mourino J, Gerstner W. 2004. Noninvasive brain-actuated control of a mobile robot by human EEG. *IEEE Trans. Biomed. Eng.* 51:1026–33
55. Batzianoulis I, Iwane F, Wei S, Correia CGPR, Chavarriaga R, et al. 2021. Customizing skills for assistive robotic manipulators, an inverse reinforcement learning approach with error-related potentials. *Commun. Biol.* 4:1406
56. Argall BD, Billard AG. 2010. A survey of tactile human–robot interactions. *Robot. Auton. Syst.* 58:1159–76
57. Gopinath D, Javaremi MN, Argall B. 2021. Customized handling of unintended interface operation in assistive robots. In *2021 IEEE International Conference on Robotics and Automation*, pp. 10406–12. Piscataway, NJ: IEEE

58. Javaremi MN, Young M, Argall BD. 2019. Interface operation and implications for shared-control assistive robots. In *2019 IEEE 16th International Conference on Rehabilitation Robotics*, pp. 232–39. Piscataway, NJ: IEEE
59. Devlin J, Chang MW, Lee K, Toutanova K. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1: *Long and Short Papers*, pp. 4171–86. Minneapolis, MN: Assoc. Comput. Linguist.
60. Alayrac JB, Donahue J, Luc P, Miech A, Barr I, et al. 2022. Flamingo: a visual language model for few-shot learning. arXiv:2204.14198 [cs.CV]
61. Tellex S, Gopalan N, Kress-Gazit H, Matuszek C. 2020. Robots that use language. *Annu. Rev. Control Robot. Auton. Syst.* 3:25–55
62. Collins S, Ruina A, Tedrake R, Wisse M. 2005. Efficient bipedal robots based on passive-dynamic walkers. *Science* 307:1082–85
63. Todorov E. 2004. Optimality principles in sensorimotor control. *Nat. Neurosci.* 7:907–15
64. Fitzsimons K, Murphey TD. 2022. Ergodic shared control: closing the loop on pHRI based on information encoded in motion. *ACM Trans. Hum.-Robot Interact.* 11:37
65. Gielniak MJ, Thomaz AL. 2011. Spatiotemporal correspondence as a metric for human-like robot motion. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 77–84. Piscataway, NJ: IEEE
66. Fitzsimons K, Acosta AM, Dewald JP, Murphey TD. 2019. Ergodicity reveals assistance and learning from physical human-robot interaction. *Sci. Robot.* 4:eaav6079
67. Kalinowska A, Rudy K, Schlafly M, Fitzsimons K, Dewald JP, Murphey TD. 2020. Shoulder abduction loading affects motor coordination in individuals with chronic stroke, informing targeted rehabilitation. In *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics*, pp. 1010–17. Piscataway, NJ: IEEE
68. Kalinowska A, Schlafly M, Rudy K, Dewald J, Murphey TD. 2022. Measuring interaction bandwidth during physical human-robot collaboration. *Robot. Autom. Lett.* 7:12467–74
69. Edwards AL, Dawson MR, Hebert JS, Sherstan C, Sutton RS, et al. 2016. Application of real-time machine learning to myoelectric prosthesis control: a case series in adaptive switching. *Prosthet. Orthot. Int.* 40:573–81
70. Edwards AL, Hebert JS, Pilarski PM. 2016. Machine learning and unlearning to autonomously switch between the functions of a myoelectric arm. In *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics*, pp. 514–21. Piscataway, NJ: IEEE
71. Günther J, Kearney A, Dawson MR, Sherstan C, Pilarski PM. 2018. Predictions, surprise, and predictions of surprise in general value function architectures. In *Proceedings of the AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*, ed. KH Wray, JA Shah, P Stone, SJ Witwicki, S Zilberstein, pp. 22–29. Palo Alto, CA: AAAI Press
72. Bethel CL, Salomon K, Murphy RR, Burke JL. 2007. Survey of psychophysiology measurements applied to human-robot interaction. In *RO-MAN 2007: The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 732–37. Piscataway, NJ: IEEE
73. Peternel L, Tsagarakis N, Caldwell D, Ajoudani A. 2018. Robot adaptation to human physical fatigue in human–robot co-manipulation. *Auton. Robots* 42:1011–21
74. O'Neill C, Proietti T, Nuckols K, Clarke ME, Hohimer CJ, et al. 2020. Inflatable soft wearable robot for reducing therapist fatigue during upper extremity rehabilitation in severe stroke. *Robot. Autom. Lett.* 5:3899–906
75. Zhang J, Fiers P, Witte KA, Jackson RW, Poggensee KL, et al. 2017. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science* 356:1280–84
76. Ding Y, Kim M, Kuindersma S, Walsh CJ. 2018. Human-in-the-loop optimization of hip assistance with a soft exosuit during walking. *Sci. Robot.* 3:eaar5438
77. Li K, Zhang J, Wang L, Zhang M, Li J, Bao S. 2020. A review of the key technologies for sEMG-based human-robot interaction systems. *Biomed. Signal Process. Control* 62:102074

78. Swerdloff MM, Hargrove LJ. 2020. Quantifying cognitive load using EEG during ambulation and postural tasks. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 2849–52. Piscataway, NJ: IEEE
79. Luo R, Wang Y, Weng Y, Paul V, Brudnak MJ, et al. 2019. Toward real-time assessment of workload: a Bayesian inference approach. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* 63:196–200
80. Koenig A, Novak D, Omlin X, Pulfer M, Perreault E, et al. 2011. Real-time closed-loop control of cognitive load in neurological patients during robot-assisted gait training. *IEEE Trans. Neural Syst. Rehabil. Eng.* 19:453–64
81. Losey DP, McDonald CG, Battaglia E, O'Malley MK. 2018. A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. *Appl. Mech. Rev.* 70:010804
82. Li Q, Chen W, Wang J. 2011. Dynamic shared control for human-wheelchair cooperation. In *2011 IEEE International Conference on Robotics and Automation*, pp. 4278–83. Piscataway, NJ: IEEE
83. Dragan AD, Srinivasa SS. 2013. A policy-blending formalism for shared control. *Int. J. Robot. Res.* 32:790–805
84. Kalinowska A, Fitzsimons K, Dewald J, Murphey TD. 2018. Online user assessment for minimal intervention during task-based robotic assistance. In *Robotics: Science and Systems XIV*, ed. H Kress-Gazit, S Srinivasa, T Howard, N Atanasov, pap. 46. N.p.: Robot. Sci. Syst. Found.
85. Ezeh C, Trautman P, Devigne L, Bureau V, Babel M, Carlson T. 2017. Probabilistic versus linear blending approaches to shared control for wheelchair driving. In *2017 International Conference on Rehabilitation Robotics*, pp. 835–40. Piscataway, NJ: IEEE
86. Jeon HJ, Losey DP, Sadigh D. 2020. Shared autonomy with learned latent actions. In *Robotics: Science and Systems XVI*, ed. M Toussaint, A Bicchi, T Hermans, pap. 11. N.p.: Robot. Sci. Syst. Found.
87. Gopinath DE, Argall BD. 2020. Active intent disambiguation for shared control robots. *IEEE Trans. Neural Syst. Rehabil. Eng.* 28:1497–506
88. Abbink DA, Carlson T, Mulder M, De Winter JC, Aminravan F, et al. 2018. A topology of shared control systems–finding common ground in diversity. *IEEE Trans. Hum.-Mach. Syst.* 48:509–25
89. Emken JL, Harkema SJ, Beres-Jones JA, Ferreira CK, Reinkensmeyer DJ. 2007. Feasibility of manual teach-and-replay and continuous impedance shaping for robotic locomotor training following spinal cord injury. *IEEE Trans. Biomed. Eng.* 55:322–34
90. Bortole M, Del Ama A, Rocon E, Moreno JC, Brunetti F, Pons JL. 2013. A robotic exoskeleton for overground gait rehabilitation. In *2013 IEEE International Conference on Robotics and Automation*, pp. 3356–61. Piscataway, NJ: IEEE
91. Tchimino J, Markovic M, Dideriksen JL, Dosen S. 2021. The effect of calibration parameters on the control of a myoelectric hand prosthesis using EMG feedback. *J. Neural Eng.* 18:046091
92. Ang KK, Guan C. 2016. EEG-based strategies to detect motor imagery for control and rehabilitation. *IEEE Trans. Neural Syst. Rehabil. Eng.* 25:392–401
93. Mondini V, Kobler RJ, Sburlea AI, Müller-Putz GR. 2020. Continuous low-frequency EEG decoding of arm movement for closed-loop, natural control of a robotic arm. *J. Neural Eng.* 17:046031
94. Thompson A, Loke LY, Argall B. 2022. Control interface remapping for bias-aware assistive teleoperation. In *2022 International Conference on Rehabilitation Robotics*. Piscataway, NJ: IEEE. **https://doi.org/10.1109/ICORR55369.2022.9896567**
95. Erdogan A, Argall BD. 2017. The effect of robotic wheelchair control paradigm and interface on user performance, effort and preference: an experimental assessment. *Robot. Auton. Syst.* 94:282–97
96. Musić S, Hirche S. 2017. Control sharing in human-robot team interaction. *Annu. Rev. Control* 44:342–54
97. Argall BD. 2015. Turning assistive machines into assistive robots. In *Quantum Sensing and Nanophotonic Devices XII*, ed. M Razeghi, E Tournié, GJ Brown, pp. 413–24. Proc. SPIE 9370. Bellingham, WA: SPIE
98. Reinkensmeyer DJ, Boninger ML. 2012. Technologies and combination therapies for enhancing movement training for people with a disability. *J. NeuroEng. Rehabil.* 9:17
99. Blank AA, French JA, Pehlivan AU, O'Malley MK. 2014. Current trends in robot-assisted upper-limb stroke rehabilitation: promoting patient engagement in therapy. *Curr. Phys. Med. Rehabil. Rep.* 2:184–95
100. Lo AC, Guarino PD, Richards LG, Haselkorn JK, Wittenberg GF, et al. 2010. Robot-assisted therapy for long-term upper-limb impairment after stroke. *N. Engl. J. Med.* 362:1772–83

101. Klamroth-Marganska V, Blanco J, Campen K, Curt A, Dietz V, et al. 2014. Three-dimensional, task-specific robot therapy of the arm after stroke: a multicentre, parallel-group randomised trial. *Lancet Neurol.* 13:159–66
102. Iosa M, Morone G, Cherubini A, Paolucci S. 2016. The three laws of neurorobotics: a review on what neurorehabilitation robots should do for patients and clinicians. *J. Med. Biol. Eng.* 36:1–11
103. Danzl MM, Etter NM, Andreatta RD, Kitzman PH. 2012. Facilitating neurorehabilitation through principles of engagement. *J. Allied Health* 41:35–41
104. Wolbrecht ET, Chan V, Reinkensmeyer DJ, Bobrow JE. 2008. Optimizing compliant, model-based robotic assistance to promote neurorehabilitation. *IEEE Trans. Neural Syst. Rehabil. Eng.* 16:286–97
105. Cheng N, Phua KS, Lai HS, Tam PK, Tang KY, et al. 2020. Brain-computer interface-based soft robotic glove rehabilitation for stroke. *IEEE Trans. Biomed. Eng.* 67:3339–51
106. Balasubramanian S, Garcia-Cossio E, Birbaumer N, Burdet E, Ramos-Murguialday A. 2018. Is EMG a viable alternative to BCI for detecting movement intention in severe stroke? *IEEE Trans. Biomed. Eng.* 65:2790–97
107. Dukelow SP, Herter TM, Moore KD, Demers MJ, Glasgow JI, et al. 2010. Quantitative assessment of limb position sense following stroke. *Neurorehabil. Neural Repair* 24:178–87
108. Tyryshkin K, Coderre AM, Glasgow JI, Herter TM, Bagg SD, et al. 2014. A robotic object hitting task to quantify sensorimotor impairments in participants with stroke. *J. NeuroEng. Rehabil.* 11:47
109. McPherson JG, Ellis MD, Harden RN, Carmona C, Drogos JM, et al. 2018. Neuromodulatory inputs to motoneurons contribute to the loss of independent joint control in chronic moderate to severe hemiparetic stroke. *Front. Neurol.* 9:470
110. Argall BD. 2018. Autonomy in rehabilitation robotics: an intersection. *Annu. Rev. Control Robot. Auton. Syst.* 1:441–63
111. Polygerinos P, Wang Z, Galloway KC, Wood RJ, Walsh CJ. 2015. Soft robotic glove for combined assistance and at-home rehabilitation. *Robot. Auton. Syst.* 73:134–43
112. Anam K, Al-Jumaily AA. 2012. Active exoskeleton control systems: state of the art. *Procedia Eng.* 41:988–94
113. Xiloyannis M, Alicea R, Georgarakis AM, Haufe FL, Wolf P, et al. 2021. Soft robotic suits: state of the art, core technologies, and open challenges. *IEEE Trans. Robot.* 38:1342–62
114. Sun Y, Tang Y, Zheng J, Dong D, Chen X, Bai L. 2022. From sensing to control of lower limb exoskeleton: a systematic review. *Annu. Rev. Control* 53:83–96
115. Dollar AM, Herr H. 2008. Lower extremity exoskeletons and active orthoses: challenges and state-of-the-art. *IEEE Trans. Robot.* 24:144–58
116. De Looze MP, Bosch T, Krause F, Stadler KS, O'Sullivan LW. 2016. Exoskeletons for industrial application and their potential effects on physical work load. *Ergonomics* 59:671–81
117. Walsh CJ, Endo K, Herr H. 2007. A quasi-passive leg exoskeleton for load-carrying augmentation. *Int. J. Humanoid Robot.* 4:487–506
118. Gull MA, Bai S, Bak T. 2020. A review on design of upper limb exoskeletons. *Robotics* 9:16
119. Sawers A, Bhattacharjee T, McKay JL, Hackney ME, Kemp CC, Ting LH. 2017. Small forces that differ with prior motor experience can communicate movement goals during human-human physical interaction. *J. NeuroeEng. Rehabil.* 31:8
120. Takagi A, Ganesh G, Yoshioka T, Kawato M, Burdet E. 2017. Physically interacting individuals estimate the partner's goal to enhance their movements. *Nat. Hum. Behav.* 1:0054
121. Fitzsimons K, Kalinowska A, Dewald JP, Murphey TD. 2020. Task-based hybrid shared control for training through forceful interaction. *Int. J. Robot. Res.* 39:1138–54
122. Jarrasse N, Sanguineti V, Burdet E. 2014. Slaves no longer: review on role assignment for human–robot joint motor action. *Adapt. Behav.* 22:70–82
123. Liu M, Wang D, Huang H. 2015. Development of an environment-aware locomotion mode recognition system for powered lower limb prostheses. *IEEE Trans. Neural Syst. Rehabil. Eng.* 24:434–43
124. Farina D, Vujaklija I, Brånemark R, Bull AM, Dietl H, et al. 2021. Toward higher-performance bionic limbs for wider clinical use. *Nat. Biomed. Eng.* In press. **https://doi.org/10.1038/s41551-021-00732-x**
125. Igual C, Pardo LA Jr., Hahne JM, Igual J. 2019. Myoelectric control for upper limb prostheses. *Electronics* 8:1244

126. Kuiken TA, Li G, Lock BA, Lipschutz RD, Miller LA, et al. 2009. Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms. *JAMA* 301:619–28
127. Hebert JS, Olson JL, Morhart MJ, Dawson MR, Marasco PD, et al. 2013. Novel targeted sensory reinnervation technique to restore functional hand sensation after transhumeral amputation. *IEEE Trans. Neural Syst. Rehabil. Eng.* 22:765–73
128. Ortiz-Catalan M, Mastinu E, Sassu P, Aszmann O, Brånemark R. 2020. Self-contained neuromusculoskeletal arm prostheses. *N. Engl. J. Med.* 382:1732–38
129. Simon A, Turner K, Miller L, Potter B, Beachler M, et al. 2022. User performance with a transradial multi-articulating hand prosthesis during pattern recognition and direct control home use. TechRxiv 19859281. **https://doi.org/10.36227/techrxiv.19859281**
130. Schofield JS, Evans KR, Carey JP, Hebert JS. 2014. Applications of sensory feedback in motorized upper extremity prosthesis: a review. *Expert Rev. Med. Devices* 11:499–511
131. Barontini F, Catalano MG, Grioli G, Bianchi M, Bicchi A. 2021. Wearable integrated soft haptics in a prosthetic socket. *Robot. Autom. Lett.* 6:1785–92
132. Battaglia E, Clark JP, Bianchi M, Catalano MG, Bicchi A, O'Malley MK. 2019. Skin stretch haptic feedback to convey closure information in anthropomorphic, under-actuated upper limb soft prostheses. *IEEE Trans. Haptics* 12:508–20
133. Lasota PA, Rossano GF, Shah JA. 2014. Toward safe close-proximity human-robot interaction with standard industrial robots. In *2014 IEEE International Conference on Automation Science and Engineering*, pp. 339–44. Piscataway, NJ: IEEE
134. Liu H, Wang L. 2021. Collision-free human-robot collaboration based on context awareness. *Robot. Comput.-Integr. Manuf.* 67:101997
135. Trautman P, Ma J, Murray RM, Krause A. 2015. Robot navigation in dense human crowds: statistical models and experimental studies of human–robot cooperation. *Int. J. Robot. Res.* 34:335–56
136. Chen C, Liu Y, Kreiss S, Alahi A. 2019. Crowd-robot interaction: crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 International Conference on Robotics and Automation*, pp. 6015–22. Piscataway, NJ: IEEE
137. Mörtl A, Lawitzky M, Kucukyilmaz A, Sezgin M, Basdogan C, Hirche S. 2012. The role of roles: physical cooperation between humans and robots. *Int. J. Robot. Res.* 31:1656–74
138. Ghadirzadeh A, Bütepage J, Maki A, Kragic D, Björkman M. 2016. A sensorimotor reinforcement learning framework for physical human-robot interaction. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2682–88. Piscataway, NJ: IEEE
139. Edsinger A, Kemp CC. 2007. Human-robot interaction for cooperative manipulation: handing objects to one another. In *RO-MAN 2007: The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 1167–72. Piscataway, NJ: IEEE
140. Ravichandar H, Polydoros AS, Chernova S, Billard A. 2020. Recent advances in robot learning from demonstration. *Annu. Rev. Control Robot. Auton. Syst.* 3:297–330
141. Pastor P, Hoffmann H, Asfour T, Schaal S. 2009. Learning and generalization of motor skills by learning from demonstration. In *2009 IEEE International Conference on Robotics and Automation*, pp. 763–68. Piscataway, NJ: IEEE
142. Fitzgerald T, Goel A, Thomaz A. 2018. Human-guided object mapping for task transfer. *ACM Trans. Hum.-Robot Interact.* 7:17
143. Lauretti C, Cordella F, Ciancio AL, Trigili E, Catalan JM, et al. 2018. Learning by demonstration for motion planning of upper-limb exoskeletons. *Front. Neurorobot.* 12:5
144. Brown DS, Goo W, Niekum S. 2020. Better-than-demonstrator imitation learning via automatically-ranked demonstrations. In *Proceedings of the Conference on Robot Learning*, ed. LP Kaelbling, D Kragic, K Sugiura, pp. 330–59. Proc. Mach. Learn. Res. 100. N.p.: PMLR
145. Myers V, Biyik E, Anari N, Sadigh D. 2022. Learning multimodal rewards from rankings. In *Proceedings of the 5th Conference on Robot Learning*, ed. A Faust, D Hsu, G Neumann, pp. 342–52. Proc. Mach. Learn. Res. 164. N.p.: PMLR
146. Niekum S, Chitta S, Barto AG, Marthi B, Osentoski S. 2013. Incremental semantically grounded learning from demonstration. In *Robotics: Science and Systems IX*, ed. P Newman, D Fox, D Hsu, pap. 48. N.p.: Robot. Sci. Syst. Found.

147. Bobu A, Bajcsy A, Fisac JF, Deglurkar S, Dragan AD. 2020. Quantifying hypothesis space misspecification in learning from human–robot demonstrations and physical corrections. *IEEE Trans. Robot.* 36:835–54
148. Kalinowska A, Prabhakar A, Fitzsimons K, Murphey T. 2021. Ergodic imitation: learning from what to do and what not to do. In *2021 IEEE International Conference on Robotics and Automation*, pp. 3648–54. Piscataway, NJ: IEEE
149. Herlant LV, Holladay RM, Srinivasa SS. 2016. Assistive teleoperation of robot arms via automatic time-optimal mode switching. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 35–42. Piscataway, NJ: IEEE
150. Hogan N, Sternad D. 2012. Dynamic primitives of motor behavior. *Biol. Cybernet.* 106:727–39
151. Losey DP, Srinivasan K, Mandlekar A, Garg A, Sadigh D. 2020. Controlling assistive robots with learned latent actions. In *2020 IEEE International Conference on Robotics and Automation*, pp. 378–84. Piscataway, NJ: IEEE
152. Lasota PA, Fong T, Shah JA. 2017. A survey of methods for safe human-robot interaction. *Found. Trends Robot.* 5:261–349
153. Heinzmann J, Zelinsky A. 2003. Quantitative safety guarantees for physical human-robot interaction. *Int. J. Robot. Res.* 22:479–504
154. ISO. 2106. *ISO/TS 15066:2016: Robots and robotic devices – collaborative robots*. Stand., ISO, Geneva, Switz. **https://www.iso.org/standard/62996.html**
155. Folkestad C, Chen Y, Ames AD, Burdick JW. 2020. Data-driven safety-critical control: synthesizing control barrier functions with Koopman operators. *Control Syst. Lett.* 5:2012–17
156. Kulic D, Croft EA. 2007. Affective state estimation for human–robot interaction. *IEEE Trans. Robot.* 23:991–1000
157. Brown DS, Schneider J, Dragan A, Niekum S. 2021. Value alignment verification. In *Proceedings of the 38th International Conference on Machine Learning*, ed. M Meila, T Zhang, pp. 1105–15. Proc. Mach. Learn. Res. 139. N.p.: PMLR
158. Kress-Gazit H, Eder K, Hoffman G, Admoni H, Argall B, et al. 2021. Formalizing and guaranteeing human-robot interaction. *Commun. ACM* 64(9):78–84
159. Gielniak MJ, Thomaz AL. 2012. Enhancing interaction through exaggerated motion synthesis. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 375–82. Piscataway, NJ: IEEE
160. Dragan AD, Bauman S, Forlizzi J, Srinivasa SS. 2015. Effects of robot motion on human-robot collaboration. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 51–58. Piscataway, NJ: IEEE
161. Hoffman G, Ju W. 2014. Designing robots with movement in mind. *J. Hum.-Robot Interact.* 3:91–122
162. Rojas RA, Ruiz Garcia MA, Wehrle E, Vidoni R. 2019. A variational approach to minimum-jerk trajectories for psychological safety in collaborative assembly stations. *Robot. Autom. Lett.* 4:823–29
163. Winfield L, Glassmire J, Colgate JE, Peshkin M. 2007. T-pad: tactile pattern display through variable friction reduction. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 421–26. Piscataway, NJ: IEEE
164. Culbertson H, Nunez CM, Israr A, Lau F, Abnousi F, Okamura AM. 2018. A social haptic device to create continuous lateral motion using sequential normal indentation. In *2018 IEEE Haptics Symposium*, pp. 32–39. Piscataway, NJ: IEEE
165. Salvato M, Williams SR, Nunez CM, Zhu X, Israr A, et al. 2021. Data-driven sparse skin stimulation can convey social touch information to humans. *IEEE Trans. Haptics* 15:392–404
166. Parker AS, Edwards AL, Pilarski PM. 2019. Exploring the impact of machine-learned predictions on feedback from an artificial limb. In *2019 IEEE 16th International Conference on Rehabilitation Robotics*, pp. 1239–46. Piscataway, NJ: IEEE
167. Brenneis DJ, Parker AS, Johanson MB, Butcher A, Davoodi E, et al. 2022. *Assessing human interaction in virtual reality with continually learning prediction agents based on reinforcement learning algorithms: a pilot study*. Paper presented at the Adaptive and Learning Agents Workshop, International Conference on Autonomous Agents and Multiagent Systems, Auckland, N.Z., May 9–10

168. Farkhatdinov I, Ryu JH, An J. 2010. A preliminary experimental study on haptic teleoperation of mobile robot with variable force feedback gain. In *2010 IEEE Haptics Symposium*, pp. 251–56. Piscataway, NJ: IEEE
169. Dominijanni G, Shokur S, Salvietti G, Buehler S, Palmerini E, et al. 2021. The neural resource allocation problem when enhancing human bodies with extra robotic limbs. *Nat. Mach. Intell.* 3:850–60
170. Lee JM, Gebrekristos T, De Santis D, Javaremi MN, Gopinath D, et al. 2022. Learning to control complex rehabilitation robot using high-dimensional interfaces. bioRxiv 2022.03.07.483341. **https://doi.org/10.1101/2022.03.07.483341**
171. Kirby S, Hurford JR. 2002. The emergence of linguistic structure: an overview of the iterated learning model. In *Simulating the Evolution of Language*, ed. A Cangelosi, D Parisi, pp. 121–47. London: Springer
172. Tzorakoleftherakis E, Murphey TD, Scheidt RA. 2016. Augmenting sensorimotor control using goal-aware vibrotactile stimulation during reaching and manipulation behaviors. *Exp. Brain Res.* 234:2403–14
173. Wagner K, Reggia JA, Uriagereka J, Wilkinson GS. 2003. Progress in the simulation of emergent communication and language. *Adapt. Behav.* 11:37–69
174. Kalinowska A, Davoodi E, Mathewson KW, Murphey T, Pilarski PM. 2022. Towards situated communication in multi-step interactions: time is a key pressure in communication emergence. *Proc. Annu. Meet. Cogn. Sci. Soc.* 44:615–20
175. Lazaridou A, Baroni M. 2020. Emergent multi-agent communication in the deep learning era. arXiv:2006.02419 [cs.CL]
176. Crandall JW, Oudah M, Ishowo-Oloko F, Abdallah S, Bonnefon JF, et al. 2018. Cooperating with machines. *Nat. Commun.* 9:233
177. Tedersoo L, Küngas R, Oras E, Köster K, Eenmaa H, et al. 2021. Data sharing practices and data availability upon request differ across scientific disciplines. *Sci. Data* 8:192
178. Furmanek MP, Mangalam M, Yarossi M, Lockwood K, Tunik E. 2022. A kinematic and EMG dataset of online adjustment of reach-to-grasp movements to visual perturbations. *Sci. Data* 9:23
179. Newman BA, Aronson RM, Srinivasa SS, Kitani K, Admoni H. 2022. Harmonic: a multimodal dataset of assistive human–robot collaboration. *Int. J. Robot. Res.* 41:3–11
180. Rabadi MH, Rabadi FM. 2006. Comparison of the action research arm test and the Fugl-Meyer assessment as measures of upper-extremity motor weakness after stroke. *Arch. Phys. Med. Rehabil.* 87:962–66
181. Yuan L, Gao X, Zheng Z, Edmonds M, Wu YN, et al. 2022. In situ bidirectional human-robot value alignment. *Sci. Robot.* 7:eabm4183