



Communicative capital: a key resource for human–machine shared agency and collaborative capacity

Kory W. Mathewson^{1,2} · Adam S. R. Parker^{2,3} · Craig Sherstan⁴ · Ann L. Edwards² · Richard S. Sutton^{1,2,3,5} · Patrick M. Pilarski^{1,2,3,5}

Received: 25 February 2022 / Accepted: 12 October 2022 / Published online: 14 November 2022
© The Author(s) 2022

Abstract

In this work, we present a perspective on the role machine intelligence can play in supporting human abilities. In particular, we consider research in rehabilitation technologies such as prosthetic devices, as this domain requires tight coupling between human and machine. Taking an agent-based view of such devices, we propose that human–machine collaborations have a capacity to perform tasks which is a result of the combined agency of the human and the machine. We introduce *communicative capital* as a resource developed by a human and a machine working together in ongoing interactions. Development of this resource enables the partnership to eventually perform tasks at a capacity greater than either individual could achieve alone. We then examine the benefits and challenges of increasing the agency of prostheses by surveying literature which demonstrates that building communicative resources enables more complex, task-directed interactions. The viewpoint developed in this article extends current thinking on how best to support the functional use of increasingly complex prostheses, and establishes insight toward creating more fruitful interactions between humans and supportive, assistive, and augmentative technologies.

Keywords Communicative capital · Machine intelligence · Prostheses · Prosthetic devices · Human–machine interaction · Agency · Prediction learning · Communication · Rehabilitation technology · Robotics

1 Introduction

Technology can be used to amplify natural human abilities, provide access to new abilities, and supplement abilities changed due to injury or illness [1–6]. Various tools and technological interventions are well known to support humans in physically interacting with their world, improving perceptual abilities, and supporting decision-

making and memory [1, 7–9]. Interventions to provide people with the functions they require for daily life are a core area of interest in rehabilitation, as outlined by the International Classification of Functioning, Disability and Health (ICF) [10, 11]. For example, Geary [1] describes ways that technology is used to enhance sight, touch, hearing, taste, smell, and mental processes. Millán et al. [12], Castellini et al. [13], and Carmena [14] further present views on the use of technology to supplement and enhance motor and sensory abilities for people who have lost body parts or body functions. Of interest to this work are technological advances in assistive or augmentative technology involving tight coupling [15] between a person and a machine with the capacity to learn. This coupling affects the ability of the combined human–machine partnership to have, seek, and achieve goals.

We present the perspective that a human’s ability to have, seek, and achieve goals can be supported using

✉ Kory W. Mathewson
korymath@deepmind.com

¹ DeepMind, Montreal, Canada

² University of Alberta, Edmonton, Canada

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

⁴ Sony AI, Tokyo, Japan

⁵ DeepMind, Edmonton, Canada

machine intelligence, specifically by combining human ability with reinforcement learning agents [16]. We term this *human–machine shared agency*. This perspective suggests that a human and their machine counterpart should be viewed as partners attempting to accomplish a shared task, where the agency of each partner combines to allow for greater potential capacity to accomplish tasks.

As a main contribution, we introduce *communicative capital: a resource that is built up over time in a human–machine partnership that allows the partners to eventually perform tasks at a capacity greater than either individual could achieve alone*. The resource can consist of accumulated propositional or procedural knowledge, conventions, beliefs, models, and predictions of the other agent. Communicative capital is represented within each agent and is stored within the individual memory of both agents. Communicative capital directly affects the behavioural collaborative capacity of the human–machine partnership.

In this paper, we specifically consider the case where the resource is in the form of predictions learned over time from interaction between human and prosthetic devices. While our setting of interest is human–machine interaction, a helpful motivating example is a human–guide dog partnership that allows both independent agents—human and canine—to accomplish a greater range and complexity of shared tasks (discussed in Sect. 6.1).

2 Robotic upper-limb prostheses


Robotic prostheses and other examples from the field of rehabilitation technology help us focus our thinking on direct human–machine interactions that can be well supported by machine intelligence. The rehabilitation technology setting is appealing in that it involves a direct, immediate, tightly coupled collaboration between a human and their technology to achieve a goal [15, 17]. Examples of assistive rehabilitation devices include semi-autonomous wheelchairs [12, 18], robotic manipulators and locomotors [13, 19], exoskeletons [20], smart living environments [21], and socially assistive robotic coaches [22]. The representative example of assistive rehabilitation technology we focus on in the present work is *robotic upper-limb prostheses*: assistive electromechanical devices attached to the body of individuals with amputations [23] (Fig. 1). Despite the evolution of prosthetic devices from iron hands to more dexterous mechanical manipulators, and improvements in quality of life for some users, state-of-the-art devices have yet to create a satisfactory solution for many individuals [13, 24–27].

In the prosthetic setting, movement control contributions from both human and machine must combine effectively in order for the device to benefit the human user. In

this setting challenges result from the limited number of degrees of human control and the lack of feedback from the device [13, 30]. The coupling of human and device is further complicated by the dynamic, non-stationary nature of human environments [31]. This coupling has been improved by muscular, neural, and osseointegration allowing for a more direct, high-bandwidth connection between human and machine [13, 19, 27, 32, 33]. To provide a bidirectional flow of information between prostheses and their users, cameras have been used to augment perception [34], microphones and speakers have been used to facilitate natural language interactions [35], and both surgical practices and prosthetic feedback approaches have evolved [30, 36]. Prosthetic devices of the future will receive an unprecedented density of data about human users and their environment, and they should be well equipped to translate such data into actions which support the goals of the users.

Despite the potential of advanced prostheses to support human abilities, current neuroprosthetic literature describes that one remaining limitation on the interaction between human and machine is the number of independent signals flowing between human and machine partners [13]. This constrains control strategy design of upper-limb prostheses to a small number of degrees of freedom, actuated by classification or regression algorithms for real-time control. Giving the upper-limb prostheses some autonomy in their control mechanism has been shown to allow for simultaneous control of multiple degrees of freedom while still using the same number of independent human generated control signals [13]. For example, pattern recognition-based controllers have provided an improvement over conventional controllers in standardized tasks in randomized clinical trials in part because of their ability to learn to interpret and act upon diverse collections of signals provided by a human user [37, 38]. Importantly, these systems therefore require upfront investment on the part of both the device and the user in the form of initial training and subsequent adjustments in order to see the autonomy-related improvements they offer. Increasing the autonomy of a prosthetic device has been shown in many specific cases to significantly increase the capacity of the human–prosthesis partnership to efficiently and effectively accomplish functional tasks [13]. Perhaps surprisingly then, given the diverse data streams and automation capabilities noted above, the specific consequences of prostheses themselves being considered to have and share in agency during human prosthesis interaction has remained relatively under-explored. We now examine the relationship between agency and capabilities in human–prosthesis partnerships.

Fig. 1 Prostheses examples: **a** Robotic upper-limb prosthetic, **b** Human using a research prosthesis [28], **c** Human using a supernumerary limb [29]



3 Prostheses as agents

In this section, we consider the implications of treating a prosthetic device as an *agent*—an autonomous goal-seeking system. This is not a common perspective—it suggests both sides of a tightly coupled human–machine interface should be thought of as agents with goals. Drawing insight from relationships found in human–human joint action and interaction [39–42], treating a human–prosthesis interaction in this way is in fact not as unfamiliar as it might first seem; with an agent-centric view, each agent would be expected, within its capability, to grow to understand the capabilities of the other and predict how to act accordingly. That is, each agent would naturally and, to the best of its ability, explicitly model the agency of the other to increase the capacity of the partnership in a continual and incrementally increasing fashion. This form of model building and adaptation is present in rather constrained ways in existing state-of-the-art upper-limb prostheses, and something the community hopes to enhance within future prosthetic systems [13].

We first delineate degrees of agency and the resulting capabilities that each side of the prosthetic human–machine partnership may obtain. Here, the human and the machine are considered analogous to co-actors in a joint action task [39–42] or the leader and follower in a two-agent partnership [43]; this collective shared agency is cooperation between a natural and an artificial system [44]. We define agency as the degree to which an autonomous system has the ability to have, seek, and achieve goals. This definition is inspired by the Belmont Report [45], wherein a system assumes agency if it is “capable of deliberation about personal goals and of acting under the direction of such deliberation”. Hallmarks of agency include the ability to take actions, have sensation, persist over time, and improve with respect to a goal. These hallmarks give rise to an agent’s ability to predict, control, and model its environment and other agents. By taking prior perspectives on

agency into consideration [46], along with the nuances of the prosthetic setting of interest, we focus on five attributes of agency that may be present in the human or machine agent.

3.1 Be a mechanism

The agent acts in a predetermined way in response to stimulus. For example, a myoelectric controller that processes electromyographic (EMG) signals via a fixed linear proportional mapping to create control commands for prosthetic actuators [47].

3.2 Adapt over time

In addition to being a mechanism, the agent has the capacity to adapt in response to the signals perceived. Through adaptation, the agent may acquire knowledge about its situation (e.g. by modelling and adapting to perceived signals). Adaptation can occur during training, as in the supervised learning of a pattern recognition classifier, or during ongoing experience [13, 48].

3.3 Pursue a goal

The agent has defined goals and an intent to optimize some measure of its own situation. One example of the pursuit of a goal is the maximization of a scalar reward signal, as in computational and biological reinforcement learning [16].

3.4 Model the other agent as adapting

The agent views the other agent as adapting during ongoing interaction. This can alter the way one agent presents signals to the other. For example, a human user trains a pattern recognizing prosthetic with knowledge that the device is adapting to their signals.

3.5 Model the other agent as pursuing a goal

The agent views the other agent as not only changing in response to received signals, but also as pursuing its own objectives. This preliminary theory of mind further alters the way that the one agent presents signals to the other agent.

We present this list of attributes with the caveat that it is likely not exhaustive. We can imagine that there may be higher order attributes of agency which mirror the recursive theory of mind. Additional attributes may parallel high order intentionality and reasoning, as in research in animal ethology, machine theory of mind, and cultural intelligence [49–52]. This line of thinking is discussed further in Sect. 5.3.

We now outline a schema (Fig. 2) for considering degrees of agency and relate agency to the combined capacity of a human–machine partnership. Capacity and agency in this schema are agnostic to the units of measurement and the exact attributes of agency, so as to be compatible with, and still helpful across, multiple definitions of agency. *Capacity* is a measure of task performance accomplished by the human–machine partnership as quantified by some metric. *Maximum capacity* is the optimal performance that could be achieved by the partnership, illustrated and labelled ‘max partnership capacity’ in Fig. 2a. This maximum capacity can be realized or unrealized. *Realized capacity* is the actual achieved capacity of the partnership, shown as a solid red bar in Fig. 2a.

3.6 Agency

Agency is the summation of contributions from individual degrees of agency, either discrete or continuous in nature.

Multiple degrees combine to increase agency of the agent and shared agency of the partnership (Fig. 2).

3.7 Capacity function

Agency is related to capacity by a *capacity function*. By finding the point on a capacity function corresponding to a given level of agency, we can visualize the maximum capacity of partnership. A system that is a mechanism has less agency and less capacity than a system that is a mechanism, adapts over time, and pursues a goal. A partnership may result in greater capacity than the sum of the two individual systems if both partners model each other and how to effectively utilize the capabilities of both agents. A partnership can also result in a capacity less than the sum of the two individual systems if, for example, the partners interfere with each other.

As an illustrative example, Fig. 2 uses this agency-capacity schema to compare a human-mechanism partnership (without shaded rectangles) to a partnership where the machine is able to adapt (with shaded rectangles). Note how the maximum capacity of the partnership is greater than either could achieve on their own. That capacity may be initially unrealized and change over time, or it might only be realized if both agents can model the other as pursuing a goal.

The way that the goals of the human and machine align is a problem related to team formation in human-human and human-animal partnerships [53]. Such alignment can occur during normal sensorimotor interactions between agents [41, 42, 54, 55]. To examine the process by which such alignment might occur during human–machine interaction, we now introduce the idea of communicative capital. Communicative capital is a resource built up through

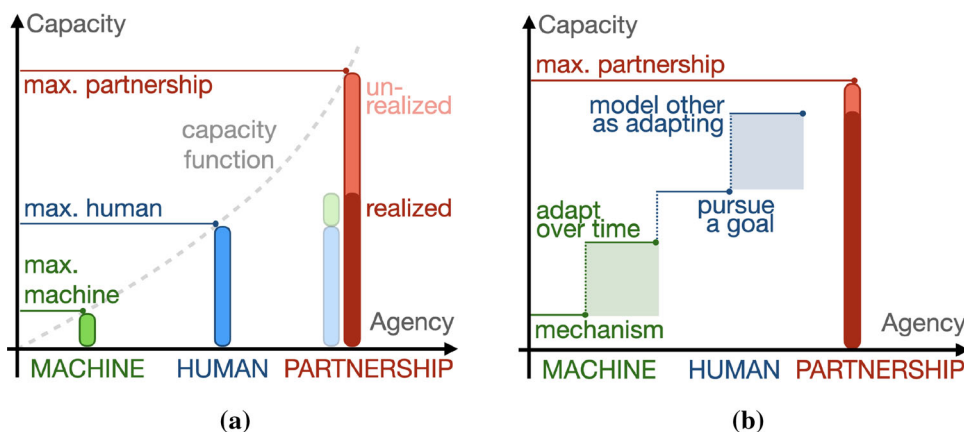


Fig. 2 The capacity function (dashed grey line) is the relationship between *capacity* and *agency*. **a** The capacity of the partnership (red) is a function of the contributions from the machine agent (green) and the human agent (blue). **b** Illustrative example of how attributes of human and machine agency can relate to maximum partnership

capacity. The light green shaded rectangle represents the capacity increase when a machine agent adapts over time versus when it acts only as a mechanism. The light blue shaded rectangle represents the increase in capacity when a human pursues a goal versus when it also models a machine partner as adapting (Color figure online)

ongoing interactions between a human and their machine counterpart that correspond to how well both agents understand each other and the partnership [56].

4 Communicative capital

As depicted in Figs. 2 and 3, the agency of the human and the machine contribute to the capacity of the partnership. *Communicative capital* is a resource built through interaction between both sides of the partnership. It enables a partnership to eventually perform a task at a capacity greater than either individual could achieve alone. Accumulating communicative capital requires investment to establish and maintain (see the ‘cost of signalling’ described by Pezzulo and Dindo [41]). The cost of investing in communicative capital may be incurred passively during the interactions of a partnership, or, in many cases, through dedicated effort tangentially related to the ultimate goals of the partnership. For example, users of prosthetic devices learn about the use of their prosthesis before they take it home for use in activities of daily living. In advanced devices that use pattern recognition, teaching both sides of a partnership to engage in a system of meaning-by-convention [57] (e.g., a series of commands to a prosthesis phrased in terms of patterns of myoelectric signals) may require significant additional time and energy but lead to increased future efficiency.

Building communicative capital can also be viewed as a process of compression and decompression, or via the lens of Scott-Phillips et al. [58, 59], one related to ostension and inference. One agent takes an action and thereby encodes information into a signal. The other agent must decode the signal as it arrives, and thereby recover the associated information. To begin to form communicative capital, at least one of the two agents must be able to adapt. Further,

we expect the greatest opportunities to build communicative capital will exist when both the human and the machine exhibit the highest possible degrees of agency. We now discuss how communicative capital can be built and used to progressively realize more capacity in prosthetic human–machine partnerships.

5 Building capital through interaction

So far we have considered settings where a communication channel exists between the human and the machine. While this channel can be either unidirectional or bidirectional, two-way communication is often beneficial for interactions between multiple goal-seeking agents. If the agent’s goals are not furthered by the information received, then it may ignore the received information. If one agent’s goals are not furthered by what the other agent does with received information, it will choose to not send such information in the future. The agent can send many possible things, and can therefore choose how to balance the cost of sending information with the expected outcomes for itself and the partnership [41]. It follows that both agents should vary their communication to send information that results in both improving with respect to their goals. The variation of communication could be independent, or guided by other parties—e.g., the work of clinical staff to train a patient for prosthesis use, or an instructor helping someone collaborate with a guide dog [60].

In effect, the processes of building communicative capital toward the attainment of goals is about the specification and identification of things each agent cares about, as in “when I do *this* it means *this*”. There can then be a natural progression in the interaction as the two sides get to know each other better. For example, in the progression shown in Fig. 3, improved predictions represent one form

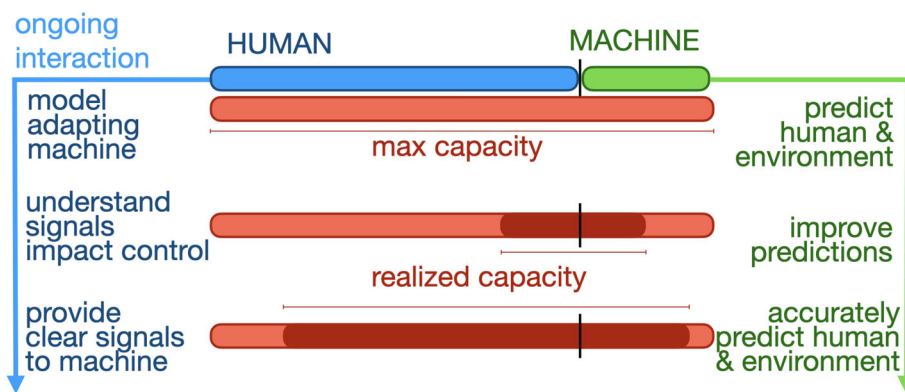


Fig. 3 Communicative capital is acquired by the partnership over the course of ongoing interaction. Prior to the partnership interacting, the human (blue) and machine (green) have acquired no communicative capital and thus have no realized capacity. Then, over the course of

ongoing interaction (from top to bottom) through modelling, improved predictions, and understanding of the signals of one another, the partnership acquires communicative capital which leads to increased realized capacity (dark red) (Color figure online)

of communicative capital. Beneficial collaboration often requires that at least one agent model and predict information about the other. This modelling of the other enables the partnership to achieve tasks with less effort and less explicit communication. This viewpoint is compatible with perspectives on human-human motor coordination [43] and with prosthetic control approaches like pattern recognition [61, 62] as discussed below.

In the following sections, we use the idea of communicative capital and the agency-capacity schema defined in Sect. 3 to examine experimental work where prosthetic control has been improved by ongoing interactions between the device and the user. First, we explore human interactions with *adaptive mechanisms* like pattern recognition systems in commercially available prostheses, and then we detail interactions with *goal-seeking prosthetic agents*.

5.1 Adaptation: prediction enhanced control

First, we consider communicative capital in adaptive control paradigms—specifically, machine learning based prosthetic controllers. There are multiple examples where the human views the machine as adapting and where the machine models and predicts information about the human to better fulfill the human's intentions [13, 48, 63, 64].

In commercial prostheses with *pattern recognition*, the human engages in a training phase to inform the device about the preferred motions to perform in response to complex patterns of myoelectric activity recorded from the human's body [13, 62]. The use of pattern recognition can provide users with more intuitive control of their prosthesis [13]. The human becomes more skilled at providing clear training commands, in part because of their knowledge that the machine is learning and adapting from the ongoing interaction. The result is improved capacity due to an increase in communicative capital: the number of human controllable functions can now exceed the number of available degrees of control available in conventional myoelectric control which depends on antagonistic muscle pairs for each degree of freedom [65].

A second example is *adaptive and autonomous switching* [63, 64, 66]. In this setting, a machine learns to make ongoing predictions about how and when a human will decide to switch between controlling one functional joint of a prosthetic device (e.g. the wrist, elbow, or shoulder) and another (Fig. 4). In manual switching, the human uses a separate biophysical control interface to send a 'change currently controlled joint to the next in a fixed list' signal to the device. In adaptive switching, the device adapts to the human by suggesting which joint it predicts the user might want to control next. The human's ability to quickly perform tasks is improved by these suggestions. The device


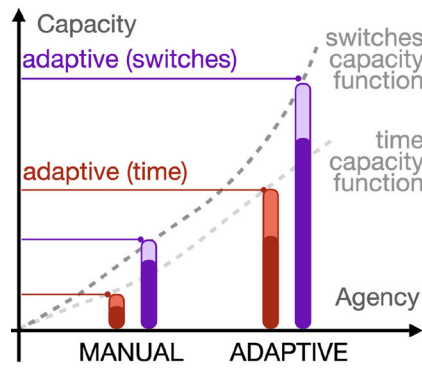
improves its suggestions based on ongoing observations about the human's actions and preferences. The adaptive nature of the machine, and the increased agency of expert humans to model the machine, lead to increased capacity to successfully complete the task efficiently in terms of reducing both total task time and total switches needed by a human user to complete a task (Fig. 4a,b). In autonomous switching, the device automatically switches which joint is currently controlled. This is done by making and using predictions to automatically switch between the functional control of different prosthetic device joints (see Fig. 5) [63, 66]. Predictions are an acquirable form of communicative capital built up by a machine learning agent during its interactions with a human and the environment.

Observations from both adaptive and autonomous switching suggest that the human begins to model the device as an agent that makes predictions [66]. As human subjects became more familiar, both with their execution of a task and with the role of machine learning as it adapted to a task, they reported greater trust in the autonomy of the device. In these experiments, certain regions of task spaces were observed where the learning system performed with close to 100% prediction accuracy. In these regions, subjects' behaviour suggested they needed to monitor the prosthetic arm less (e.g., the reduced number of manual switches in Fig. 5).

In the autonomous switching experiments of Edwards et al. [66], users began to predict autonomous switches, often moving the next functional prosthetic actuator prior to hearing a cue alerting them to the machine's automatic switching behaviour. Increased capacity in terms of reduced manual switching, and the communicative capital that supports it, is evident in users who have extensive prior experience operating adaptive prosthetic devices (see Fig. 5c,d). Users who had a greater understanding of the prosthetic learning system tended to perform actions that benefited learning, allowing the prosthetic arm to build up expectations about their behaviour more swiftly.

Another related example is the work of Sherstan et al. [67]. In this work, a human and a machine learning system share agency in controlling the movement of a robotic arm. The user is only able to control a single joint of the arm at a time and must switch between joints as needed in order to complete a task. The machine agent observes the human's behaviour and learns to predict the expected joint angles of the robot arm. These predictions are then used to move the arm in collaboration with the human's own actions [42].

As a final example of adaptive assistive technology related to the upper-limb prosthetic setting, Xu et al. [68] describe a walking-aid robot designed to autonomously adapt to different users. The robot uses reinforcement learning to adjust the relative control of the human in real-time for smoother, faster movement. Smoothness of




(a)


(b)

Fig. 4 **a** An illustrative example of how adaptive switching enables a prosthetic device to model the way a human uses the functions of a prosthesis and thereby increase agency when compared to the manual switching condition. The increase in shared agency from the manual to adaptive mode of interaction corresponds to increased capacity in


terms of (red) time to complete a task, and (purple) the number of switches required to complete a task. This plot shows data approximated from Edwards et al. [66] for illustrative purposes, and **(b)** their participant using the device [63, 64, 66] (Color figure online)




(a)



(b)



(c)



(d)

Fig. 5 Measured capacity for autonomous and adaptive switching for **(a)** expert and **(b)** non-expert humans (plots adapted from Edwards et al. [66]), summarized in capacity functions relating to **(c)** the

number of manual interactions and **(d)** the total number of switches required to complete a control task, expert humans realized more capacity than non-experts

motion, system safety, and intuitive control can all be viewed as different capacity functions that are improved by the adaptive nature of the machine.

5.2 Goals: reward-based control

Goal-seeking behaviour on the part of both the human and the machine—behaviour driven by processes of reinforcement learning—enables a more detailed progression of interactions than is possible with an adaptive, but not goal-

seeking, machine. What follows is one hypothetical progression of the training of an assistive machine, where both the human and the machine are goal-seeking agents, and where the human starts to model the device as a goal-seeking agent. This modelling and adaptation can be observed behaviourally as in the previous section.

1. At the outset, the human can only provide positive feedback (i.e. reward) signals indicating their approval; no other signals have any agreed upon meaning.

2. Using these rewards, the machine can learn a function that maps signals from the human, or other environmental cues, to a valuation that is grounded in cumulative reward (a *value function*, as detailed by Sutton and Barto [16], and used in face valuing by Veeriah et al. [69]).
3. Using this value function, the human teaches the machine a convention that may be used to interact at a low level—e.g., simple commands, body language, cues like pointing, and the basics of shifting between different functions of a system. The human begins to model how their behaviour affects the learning and adaptation of the machine.
4. Using these developed conventions, higher-level abstractions can be established between the human and the machine. These built-up conventions are one component of communicative capital which enable the realization of additional partnership capacity.

With this progression in mind, there are a variety of compatible ways to incorporate human knowledge into a learning system [70–73]. Starting with the idea of training based on primary reward, as in the progression described above, Knox and Stone [74] introduced the *Interactive Shaping Problem*, wherein an agent is acting in an environment and a human is observing the agent's performance and providing feedback to the agent such that the agent must learn the best possible way to act based on that feedback. The interactive shaping problem is related to communicative capital, as it is a readily observable case of information sharing between two goal-seeking systems with a limited channel of communication.

Goal-seeking behaviour in a machine, and developing communicative capital through the human's modelling of the machine as goal-seeking agent, increases the maximum capacity of a partnership. A human's interactions with a machine are supported by a channel of communication with defined semantics (e.g., the reward channel in reinforcement learning [16]) that allows the human to shape the machine's behaviour in ways that are not possible for an adaptive, non-goal-seeking machine. This communication channel is integral to realizing the goal-seeking agent's capacity to deal with non-stationary tasks, changing problem domains, and novel environments, in a way that aligns with the human's goals. Providing the means by which to shape behaviour can also reduce the amount of pretraining for the system, as interactions are now accompanied by online, real-time human feedback. Reward allows the human to shape the machine learning agent to perform the task in a personalized, and situation-specific way—an adaptive goal-seeking agent has the ability to incorporate engineered knowledge, but also move beyond it.

Previous work has demonstrated how both predefined and human-delivered reward could be provided to a goal-seeking agent to gradually improve the control capabilities of a myoelectric control interface [48, 75]. By using a goal-seeking reinforcement learning agent to control the joints of a prosthesis, informed by predictions about future movement, the human–machine partnership was found to be able to progressively refine the simultaneous multi-joint myoelectric control of a robotic arm. In these studies, human approval and disapproval was delivered to the machine with full knowledge of the machine's learning capabilities. These initial results have been extended to more complex settings which informs how mutual, goal-seeking behaviour supports myoelectric control [76]. These results demonstrate the value of developing communicative capital through the explicit incorporation of human feedback signals. In this representative work communicative capital led to an increased partnership capacity.

5.3 Models, shared agency, and feedback

Beliefs about the nature of internal and external signals are a kind of knowledge that we broadly denote as *models*. Models are required for the higher level attributes of agency; it is useful for a machine to represent, or construct a model, of its partner and the world, in order to achieve more effective interaction. Agent models, as they apply to a human-prosthetic partnerships, may take many forms. They may include, for instance, a collection of learned, temporally extended predictions about the dynamics of the world and the behaviour of the human [16, 77, 78].

As described by Pezzulo and Dindo [54] shared representations may be a critical part of communication during human–machine interaction, and central to the formation of more effective models in terms of beliefs, actions, and intentions. This moves us towards developing a theory of mind—an agent predicting the internal beliefs, motivations, and thoughts of another especially as applied to observable sensorimotor interactions [41, 43, 54]. Recursive theory of mind might imply higher levels of agency, as presented in Sect. 4, and parallel higher order intentionality [49–52]. Future work may explore this other-modelling and how it can be leveraged to build shared knowledge.

As one example of how models can impact a human–machine partnership, Bicho et al. [79] describe a shared construction task in which a robot and a human must work together to assemble a toy. Completion of the assembly task required actions from both agents. The robot infers the goal of the human from contextual clues and acts accordingly, communicating its intention at each point during the task using a speech synthesizer. This allows the human to further model the internal processes of the machine. Another example of a joint task in which a robot infers the

goal of the human comes from Liu and Hedrick [80]. In their work, participants and virtual robots collaborate to accomplish a task, and the robot infers the human's goal based on motion. This research suggests that goal inference (i.e., the modelling goals) decreased the time required to finish tasks and improved other measures of performance, including human–machine trust.

The impact of feedback from an adaptive prosthetic is quantified in work by Parker et al. [81]. In their work, three different kinds of feedback were used to supply a human with information about how best to control the movements of a wearable robot in the form of a supernumerary limb (see Fig. 1c)—no feedback, mechanistic feedback, and adaptive feedback in the form of predictions. The human needed to move the robot in a confined work space, coming as close as possible to the work space's walls without making physical contact. The human was blindfolded and was acoustically isolated by way of noise-cancelling headphones, so that they only received information about the world via the machine's feedback.

The two capacity functions of interest in Parker et al. [81] measured: the current drawn by the motors due to impacts with the work space walls, and the number of times the human was able to use the arm to fully traverse the work space in the given time. On different trials, feedback from the device was either absent, delivered mechanistically upon contact with the walls, or delivered proportional to learned predictions about impacts with the walls. Realized capacity in terms of current draw was found to increase for the case where the human was paired with the adaptive machine, but was found to approach a reduced maximum capacity for the case of mechanistic feedback from the device (see Fig. 6). This work provides insight into how developing communicative capital, specifically through explicitly modelling and increased agency in the delivery of feedback, can influence the maximum capacity possible for a human-prosthetic partnership.

6 Discussion

This article has discussed the setting of human–machine interaction, specifically the interactions between a human and their prosthetic technology. However, the ideas presented above regarding agency and communicative capital can be identified and analyzed in the interactions between any two or more intelligent systems. In this section, we provide supporting context from both biological and non-biological examples of how agency plays a role in the interactions of multiple agents to achieve a goal.

6.1 Guide dogs and intelligent assistants

A guide dog could be the oldest documented example of an assistive technology with agency, with an early depiction on the wall of a house excavated in Pompeii dated from c. 79 CE [53, 60]. A guide dog needs to be part of an active partnership—it must have the capability to willingly disobey an instruction when it perceives a danger. The agent in charge of the interaction, human or dog, needs to be able to change from moment-to-moment in order for the partnership to be effective. Because of these desired and atypical behaviours, both the dog and the future owner must be explicitly trained. The human must be taught not only the precise vocabulary understood by the guide dog, but what to expect in response. This requires both parties, human and dog, to invest in communicative capital and learn each others' idiosyncrasies in order to approach an effective partnership [82].

Computers, whether desktops, tablets, or smartphones, all augment our cognitive abilities. At present, there is significant effort to develop virtual assistants on such devices. Such assistants may have some level of agency; these assistants may be adaptive, changing their behaviour and suggestions to meet the user's needs [83]. To date, existing computer interfaces have largely remained fixed and unadaptive. However, thanks in part to increases in available computation, computers are now improving in their ability to predict user needs and to provide users with the information and interfaces that are most needed at any given moment [83, 84]. With increased agency, these systems now begin to demonstrate some of the hallmarks of human-human joint action established by the related literature [39, 40, 85].

6.2 Interactive approaches to instruction, communication, and control

There are multiple ways that a human and a machine—e.g., an assistive robot like a prosthesis—can beneficially interact to achieve the human's objectives [70, 72, 86]. A pertinent family of methods, broadly classified as *interactive machine learning* (IML), has demonstrated the potential to increase the capabilities of decision making systems in complex, dynamic, and novel environments.¹ In much of the existing IML literature, feedback channels are used as a means by which a non-expert can train, teach, and interact with a system without explicitly programming it. Shaping allows for the human to learn how the system

¹ Though, some argue that **all machine learning is interactive machine learning** because humans interact with machines through every step of the design, development, deployment, and dissemination of such systems [87].

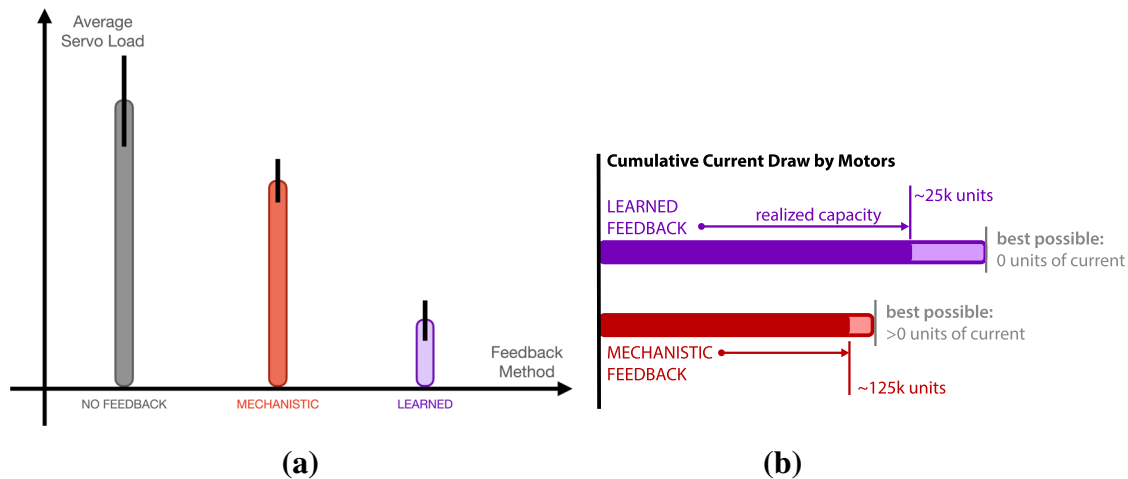


Fig. 6 The difference between learned and mechanistic feedback during control of a supernumerary limb (Fig. 1c). **a** Adaptive machine significantly reduced current drawn by the motors of the robot arm

accepts and interprets feedback and for the system to learn the goals of the human [70].

IML has produced a number of important milestones. With respect to goal-driven systems, trial-and-error machine learning has been shown to be accelerated through the presentation of human-delivered reward and forms of intermediate reinforcement. Examples include the use of shaping signals [88], the delivery of reward from both a human and the environment [74], multi-signal reinforcement [70], and combinations of both direct control and reward-based feedback [48, 75, 76]. As described in Sect. 5.2 above, an agent's learning can be facilitated by a human host through interactive reinforcement learning [74, 89, 90]. Griffith et al. [91] built on the earlier work of Knox and Stone [89] with a framework to maximize the information gained from human feedback. Loftin et al. [92] expanded the space of human interaction through detailed investigation of human teaching strategies and developed systems which model the human feedback. Their systems have been shown to learn faster and with less feedback than other approaches. Interactive learning from demonstrations and instructions have also been shown to help teach different ways of behaving to a learning machine [86, 88, 93–97].

Humans can utilize a number of different approaches to effectively communicate their goals to machine learning agents. Through interactive learning, information from a human can help a machine learner to achieve arbitrary user-centric goals, can improve a system's learning speed, and can increase the overall performance of a learning system. Advances in IML provide a basis for increasing the rate with which a human-prosthetic partnership may develop communicative capital and thereby realize

(adapted from Parker et al. [81]). **b** Increased machine agency increase realized capacity of the partnership through an investment of communicative capital

capacity, and, in certain cases, can also be expected to increase the maximum capacity of a partnership.

6.3 Limitations

There are challenges and limitations in creating machine agents that can build up communicative capital to collaborate more effectively with their human partners. In this section, we highlight several critical areas of focus that should be addressed in future work. Of particular note are challenges related to safely deploying machine learning algorithms in the real-world, especially when deployed on robots tightly coupled to human users. Future work on these algorithms is needed to empirically demonstrate how they are provably robust to a wide variety of environmental factors. As well, mechanisms to align the goals of the human and the machine are critical in shared agency settings. It has been shown in previous research how increasing agency of the machine increases the cognitive demands placed upon the human [76]. Humans often expect machines to function as mechanisms, unaffected by adaptation. There can be significant implications on their cognitive load once they are required to carry out their own actions as well as model the learning agent [98]. Finally, algorithms deployed in human-machine partnerships will need to adapt quickly to information and signals from the human. Both for reasons of safety, but also because a lack of quick adaptation could lead to human disengagement if the human doesn't perceive the machine as learning fast enough. Future work on safety, alignment, rapid adaptation, understanding human expectations, and making connections between these systems and modern theories of agency is needed as human-machine partnerships move from the laboratory and into the world. This is true for both

prosthetic devices and for collaborative machines more generally.

7 Paradigms for evaluation

We expect that increasing the agency of a prosthetic device and investing in communicative capital will allow a collaborative partnership to accomplish tasks faster, easier, more safely, and more efficiently. Work is now needed to test this hypothesis and identify the contributions and practical utility of agency and goal-seeking behaviour on the part of machine learning partner agents. It is our recommendation that researchers design experiments varying the level of agency of both human and the machine in a controlled fashion to assess the contributions from each component of agency. As described in Sect. 5.2, increased agency on the part of the machine enables increased shared agency. This increase is depicted as relative changes in the agency and capacity of both agents.

One means by which to test agent contributions is through the conventional outcome measures used to assess the impact of rehabilitation interventions [99–102]. Outcome measures provide a clearly defined notion of capacity. Further, prosthetic outcome measures are already used to study the benefits of pairing patients to systems with different mechanistic levels of agency (e.g., during prosthetic fitting and patient assessment). In the majority of clinically deployed prostheses, the control approach and system design of the device is fixed. The communicative capital of the mechanism—how it interprets body signals and maps them to actuators—provides immediate realized capacity at a level determined by the mechanism’s designers. Measures like the Southampton Hand Assessment Procedure, the Box-and-Blocks Task, and others are used to provide a quantitative assessment of the impact of these prosthetic mechanisms [102, 103]. Recent developments in the assessment of gaze and movement have further shown concrete, capacity-related metrics that evaluate user-prosthesis abilities via changes in the relationship between biomechanics and visual attention, as well as other measurable correlates of perceived control and agency [104–107]. Some of these measures have been shown to serve as proxies for the state of human predictive models of their machine partner, and thus may provide a way to quantify communicative capital as it is built by the human side of a human–machine partnership [104]. Rigorous, incremental testing of agency is therefore highly compatible with existing approaches, and will be significantly extended as more comprehensive motor, sensory, and cognitive outcome measures are developed.

One fruitful avenue for experimentation, as explored in Parker et al. [29], is to deliberately reduce the agency of the

human by removing control options and/or sensory inputs as they complete a task. In this way, the authors were able to elucidate how different levels of agency in the machine contribute to the performance of the partnership. A second, complementary paradigm is to dramatically increase the agency of the machine beyond what is technically possible, so as to study the outcomes and conditions that support shared agency. One way to do this is a type of sham trial known as a *Wizard-of-Oz* experiment (e.g. Viswanathan et al. [18]). Paradigms for evaluating human–machine partnerships will continue to develop as technology supporting shared agency evolves. We now conclude with several brief reflections.

8 Conclusions

We argue that tightly coupled human–machine partnerships, such as humans and prostheses, should be thought of as adaptive multi-agent systems where the agency of human and machine combine to achieve more capacity than either could independently. We present an agency-capacity schema that relates shared agency to the capacity of human–machine partnerships, and we show how communicative capital is the key resource that a partnership needs to invest in to access the full capacity of the combined agency of the pairing. Using examples from the literature, we illustrate how increases in the agency of a prosthesis can tangibly improve the capabilities of its human user. We highlight three main conclusions from this work as novel contributions supporting human-prosthesis interaction: (1) we propose that designing assistive devices as goal-seeking agents improves the range of possibilities for robust and flexible interaction, (2) we argue that an agent-based viewpoint of human–machine interaction enables a structured progression toward more capable partnerships between people and devices, and (3) we describe how communicative capital is a resource built through ongoing human–machine interaction which enables a partnership to eventually perform tasks at a capacity greater than either could individually. Machine intelligence enables the acquisition and use of communicative capital in human-prosthesis partnerships to more effectively and more efficiently accomplish tasks. We believe the agency-based viewpoint on assistive technology proposed in this work contributes unique and complementary ideas to the development of highly functional human–machine partnerships. Designers and developers should construct systems which actively invest in communicative capital as such investment will lead to increases in shared agency to achieve more capacity than they would be able to otherwise.

Acknowledgements The authors thank the other members of the Bionic Limbs for Natural Improved Control Laboratory and the Reinforcement Learning and Artificial Intelligence Laboratory for many helpful thoughts and comments. They specifically thank Michael Rory Dawson, Jacqueline Hebert, Nolan Bard, Doina Precup, Heather Williams, Quinn Boser, and Jaden Travnik for a number of fruitful discussions and their detailed feedback.

Author contributions All authors contributed to the conception and design of the work, the interpretation of data and studies, drafting and revising the work for important intellectual content, the approval of the version to be submitted/published, and agree to be accountable for all aspects of the work.

Funding This research was undertaken, in part, thanks to funding from the Canada Research Chairs program, the Canada Foundation for Innovation, the Alberta Innovates Centre for Machine Learning / the Alberta Machine Intelligence Institute, Alberta Innovates – Technology Futures, DeepMind, and the Natural Sciences and Engineering Research Council.

Data availability Not applicable.

Declarations

Conflict of interest KWM, RSS, and PMP are employees of DeepMind and CS is an employee of Sony AI.

Ethics approval Not applicable.

Consent to participate Not applicable.

Code availability Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Geary J (2002) *The body electric: an anatomy of the new bionic senses*. Rutgers University Press, New Brunswick, N.J
- Doidge N (2007) *The brain that changes itself: stories of personal triumph from the frontiers of brain science*. Viking, New York
- Dewdney C (1998) *Last flesh: life in the transhuman era*. HarperCollins, Toronto
- Brooks R (2002) *Flesh and machines: how robots will change us*. Pantheon Books, New York
- Belfiore M (2010) *The department of mad scientists: how DARPA is remaking our world, from the internet to artificial limbs*. Harper, New York
- Moss F (2011) *The sorcerers and their apprentices: how the digital magicians of the mit media lab are creating the innovative technologies that will transform our lives*. Crown Business, New York
- Clark A (2008) *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford University Press, Oxford New York
- Risko EF, Gilbert SJ (2016) Cognitive offloading. *Trends Cogn Sci* 20:676–688
- Osiurak F, Badets A (2016) Tool use and affordance: Manipulation-based versus reasoning-based approaches. *Psychol Rev* 123(5):534–68
- World Health Organization (2002) Towards a common language for functioning, disability and health: ICF. In: *The International Classification of Functioning, Disability and Health*
- Jette AM (2006) Toward a common language for function, disability, and health. *Phys Ther* 86(5):726–734
- Millán, JdR, Rupp R, Mueller-Putz G, Murray-Smith R, Giugliemma C, Tangermann M, Vidaurre C, Cincotti F, Kubler A, Leeb R, Neuper C, Mueller K, Mattia D (2010) Combining brain-computer interfaces and assistive technologies: State-of-the-art and challenges. *Front Neurosci* 4
- Castellini C, Artemiadis P, Winger M, Ajoudani A, Alimusaj M, Bicchi A, Caputo B, Craelius W, Dosen S, Englehart K, Farina D, Gijssberts A, Godfrey SB, Hargrove L, Ison M, Kuiken T, Marković M, Pilarski PM, Rupp R, Scheme E (2014) Proceedings of the first workshop on peripheral machine interfaces: going beyond traditional surface electromyography. *Front Neurobot*, 8
- Carmena JM (2012) Becoming bionic. *IEEE Spectr* 49(3):24–29
- Licklider JCR (1960) Man-computer symbiosis. *IRE Trans Hum Fact Electron HFE* 1(1):4–11
- Sutton RS, Barto AG (2018) *Reinforcement learning: an introduction*, 2nd edn. MIT Press, Cambridge, MA
- Parker ASR, Pilarski, (2021) PM Position statement: Assistive technology as partners through machine-learned communication. In: *Workshop on Reinforcement Learning for Humans, Computer, and Interaction (RL4HCI)*, ACM CHI 2021
- Viswanathan P, Bell J, Wang RH, Adhikari B, Mackworth AK, Mihailidis A, Miller WC, Mitchell IM (2014) A wizard-of-oz intelligent wheelchair study with cognitively-impaired older adults: attitudes toward user control. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Workshop on Assistive Robotics for Individuals with Disabilities: HRI Issues and Beyond*, September 14, 2014, Chicago, Illinois, USA
- Ortiz-Catalan M, Mastinu E, Sassu P, Aszmann O, Brånemark R (2020) Self-contained neuromusculoskeletal arm prostheses. *N Engl J Med* 382(18):1732–1738
- Herr, H (2009) Exoskeletons and orthoses: classification, design challenges and future directions. *J NeuroEng Rehabil* 6(1)
- Rashidi P, Mihailidis A (2013) A survey on ambient-assisted living tools for older adults. *IEEE J Biomed Health Inform* 17(3):579–590
- Feil-Seifer D, Matarić MJ (2011) Socially assistive robotics. *IEEE Robot Autom Mag* 18(1):24–31
- Pilarski PM, Hebert JS (2017) Upper and lower limb robotic prostheses. In: *Encarnacao P, Cook AM (eds) Robot Assist Technol Princ Pract*. CRC Press, Florida, pp 99–144
- Ziegler-Graham K, MacKenzie EJ, Ephraim PL, Trivison TG, Brookmeyer R (2008) Estimating the prevalence of limb loss in the united states: 2005 to 2050. *Arch Phys Med Rehabil* 89(3):422–429

25. Peerdeman B, Boere D, Witteveen H, Huisin't Veld R, Hermens H, Stramigioli S, Rietman H, Veltink P, Misra S (2011) Myoelectric forearm prostheses: state of the art from a user-centered perspective. *J Rehabil Res Dev* 48(6):719
26. Williams III TW (2011) Progress on stabilizing and controlling powered upper-limb prostheses. *J Rehabil Res Dev* 48(6)
27. Zuo KJ, Olson JL (2014) The evolution of functional hand replacement: From iron prostheses to hand transplantation. *Plast Surg* 22(1):44–51
28. Dawson MR, Sherstan C, Carey JP, Hebert JS, Pilarski PM (2014) Development of the bento arm: An improved robotic arm for myoelectric training and research. MEC'14: Myoelectric Controls Symposium. Fredericton, New Brunswick, Canada 14:60–64
29. Parker ASR, Edwards AL, Pilarski PM (2014) Using learned predictions as feedback to improve control and communication with an artificial limb: preliminary findings. *CoRR abs/1408.1913*
30. 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(5):499–511
31. Saridis GN, Stephanou HE (1977) A hierarchical approach to the control of a prosthetic arm. *IEEE Trans Syst Man Cybern* 7(6):407–420
32. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442(7099):164–171
33. Ortiz-Catalán M, Håkansson B, Brånemark R (2014) An osseointegrated human-machine gateway for long-term sensory feedback and motor control of artificial limbs. *Sci Transl Med* 6(257):257–62576
34. Marković M, Dosen S, Cipriani C, Popović DB, Farina D (2014) Stereovision and augmented reality for closed-loop control of grasping in hand prostheses. *J Neural Eng* 11(4):046001
35. Kollar T, Tellex S, Roy D, Roy N (2010) Toward understanding natural language directions. In: 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp 259–266
36. Hebert JS, Elzinga KE, Chan KM, Olson J, Morhart MJ (2014) Updates in targeted sensory reinnervation for upper limb amputation. *Curr Surg Rep* 2:1–9
37. 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(1)
38. Vu PP, Vaskov AK, Irwin ZT, Henning PT, Lueders DR, Laidlaw AT, Davis AJ, Nu CS, Gates DH, Gillespie RB, Kemp SWP, Kung TA, Chestek CA, Cederna PS (2020) A regenerative peripheral nerve interface allows real-time control of an artificial hand in upper limb amputees. *Sci Transl Med* 12(533):2857
39. Günther Knoblich SB, Sebanz N (2011) Psychological research on joint action: theory and data. In: Ross B (ed) *The psychology of learning and motivation*, vol 54. Academic Press, Burlington, pp 59–101
40. Pesquita A, Whitwell RL, Enns JT (2018) Predictive joint-action model: a hierarchical predictive approach to human cooperation. *Psychon Bull Rev* 25(5):1751–1769
41. Pezzulo G, Donnarumma F, Dindo H (2013) Human sensorimotor communication: a theory of signaling in online social interactions. *PLoS ONE* 8
42. Sebanz N, Bekkering H, Knoblich G (2006) Joint action: bodies and minds moving together. *Trends Cogn Sci* 10:70–76
43. Candidi M, Curioni A, Donnarumma F, Sacheli LM, Pezzulo G (2015) Interactional leader-follower sensorimotor communication strategies during repetitive joint actions. *J Royal Soc Interface* 12
44. Misselhorn C (2015) Collective agency and cooperation in natural and artificial systems. In: *Collective Agency and Cooperation in Natural and Artificial Systems*, pp 3–24. Springer
45. Brady JV, Jonsen AR (2014) The Belmont report: ethical principles and guidelines for the protection of human subjects of research. *J Am Coll Dent* 81–3:4–13
46. Tomic PT, Agha GA (2004) Towards a hierarchical taxonomy of autonomous agents. In: 2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583) 4: 3421–34264
47. Parker PA, Englehart KB, Hudgins B (2006) Myoelectric signal processing for control of powered limb prostheses. *J Electromyogr Kinesiol Off J Int Soc Electrophysiol Kinesiol* 16(6):541–8
48. Pilarski PM, Dawson MR, Degris T, Carey JP, Chan KM, Hebert JS, Sutton RS (2013) Adaptive artificial limbs: a real-time approach to prediction and anticipation. *IEEE Robot Autom Mag* 20:53–64
49. Heyes CM (1998) Theory of mind in nonhuman primates. *Behav Brain Sci* 21(1):101–114
50. Cultural General Intelligence Team A, Bhoopchand B, Brownfield A, Collister AD, Lago A, Edwards R, Everett A, Frechette YG, Oliveira E, Hughes KW, Mathewson P, Mendolicchio J, Pawar M, Pislar A, Platonov E, Senter S, Singh A, Zacherl LM, Zhang (2022) Learning robust real-time cultural transmission without human data. *CoRR* <https://doi.org/10.48550/arXiv.2203.00715abs/2203.00715>
51. Rabinowitz N, Perbet F, Song F, Zhang C, Eslami SA, Botvinick M (2018) Machine theory of mind. In: *International Conference on Machine Learning*. PMLR, pp 4218–4227
52. Zhu H, Neubig G, Bisk Y (2021) Few-shot language coordination by modeling theory of mind. In: *International Conference on Machine Learning*. PMLR, pp 12901–12911
53. Fishman GA (2003) When your eyes have a wet nose: the evolution of the use of guide dogs and establishing the seeing eye. *Surv Ophthalmol* 48(4):452–8
54. Pezzulo G, Dindo H (2011) What should I do next? Using shared representations to solve interaction problems. *Exp Brain Res* 211:613–630
55. Sebanz N, Knoblich G (2009) Prediction in joint action: what, when, and where. *Top Cogn Sci* 1(2):353–67
56. Pilarski PM, Sutton RS, Mathewson KW (2015) Prosthetic devices as goal-seeking agents. In: *2nd Workshop on Present and Future of Non-Invasive Peripheral-Nervous-System Machine Interfaces* 7: 48
57. Santoro A, Lampinen AK, Mathewson KW, Lillicrap TP, Raposo D (2021) Symbolic behaviour in artificial intelligence. *CoRR abs/2102.03406*
58. Scott-Phillips T (2014) Speaking our minds: why human communication is different, and how language evolved to make it special. Palgrave Macmillan, UK
59. Scott-Phillips TC, Kirby S, Ritchie GRS (2009) Signalling signalhood and the emergence of communication. *Cognition* 113(2)
60. Pfaffenberger C (1976) *Guide dogs for the blind, their selection, development, and training*. Elsevier Scientific Pub. Co. Distributors for the U.S. and Canada, Elsevier/North Holland, Amsterdam New York
61. Micera S, Carpaneto J, Raspovic S (2010) Control of hand prostheses using peripheral information. *IEEE Rev Biomed Eng* 3:48–68
62. Scheme E, Englehart K (2011) Electromyogram pattern recognition for control of powered upper-limb prostheses: state of the art and challenges for clinical use. *J Rehabil Res Dev* 48(6):643

63. Edwards AL (2016) Adaptive and autonomous switching: shared control of powered prosthetic arms using reinforcement learning. In: Master's Thesis at University of Alberta
64. Edwards AL, Dawson MR, Hebert JS, Sherstan C, Sutton RS, Chan KM, Pilarski PM (2016) Application of real-time machine learning to myoelectric prosthesis control. *Prosthet Orthot Int* 40(5):573–581
65. Smith LH, Kuiken TA, Hargrove LJ (2016) Myoelectric control system and task-specific characteristics affect voluntary use of simultaneous control. *IEEE Trans Neural Syst Rehabil Eng* 24(1):109–116
66. Edwards AL, Hebert JS, Pilarski PM (2016) Machine learning and unlearning to autonomously switch between the functions of a myoelectric arm. 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob), pp 514–521
67. Sherstan C, Modayil J, Pilarski PM (2015) A collaborative approach to the simultaneous multi-joint control of a prosthetic arm. 2015 IEEE International Conference on Rehabilitation Robotics (ICORR), pp 13–18
68. Xu W, Huang J, Wang Y, Cai H (2013) Study of reinforcement learning based shared control of walking-aid robot, *Proceedings of the 2013 IEEE/SICE International Symposium on System Integration*, pp 282–287
69. Veeriah V, Pilarski PM, Sutton RS (2016) Face valuing: training user interfaces with facial expressions and reinforcement learning. *Interactive Machine Learning Workshop at IJCAI 2016*
70. Thomaz AL, Breazeal C (2008) Teachable robots: understanding human teaching behavior to build more effective robot learners. *Artif Intell* 172(6):716–737
71. Chernova S, Thomaz AL (2014) Robot learning from human teachers. *Synth Lect Artif Intell Mach Learn* 8(3):1–121
72. Pilarski PM, Sutton RS (2012) Between instruction and reward: human-prompted switching. In: *AAAI Fall Symposium: Robots Learning Interactively from Human Teachers*
73. Amershi S, Cakmak M, Knox WB, Kulesza T (2014) Power to the people: the role of humans in interactive machine learning. *AI Mag* 35(4):105–120
74. Knox WB, Stone P (2012) Reinforcement learning from simultaneous human and MDP reward. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, Vol 1, AAMAS '12, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC*, pp 475–482
75. Pilarski PM, Dawson MR, Degris T, Fahimi F, Carey JP, Sutton RS (2011) Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. 2011 IEEE International Conference on Rehabilitation Robotics. pp 1–7
76. Mathewson KW, Pilarski PM (2016) Simultaneous control and human feedback in the training of a robotic agent with actor-critic reinforcement learning. *Interactive Machine Learning Workshop at IJCAI 2016*.
77. Sutton RS, Modayil J, Delp M, Degris T, Pilarski PM, White A, Precup D (2011) Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. *The 10th International Conference on Autonomous Agents and Multiagent Systems -, vol 2. AAMAS '11. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC*, pp 761–768
78. Pilarski PM, Sherstan, C (2016) Steps toward knowledgeable neuroprostheses. In *Proc. of the 6th IEEE RAS/EMBS Int. Conf. on Biomedical Robotics and Biomechatronics (BioRob2016)*, June 26-29, 2016, Singapore, pp 220–220
79. Bicho E, Erhagen W, Le Louro, Silva EC (2011) Neuro-cognitive mechanisms of decision making in joint action: a human-robot interaction study. *Hum Mov Sci* 30(5):846–868
80. Liu C, Hedrick JK (2016) Cooperative search using human-UAV teams. In: *AIAA Infotech at Aerospace*, pp 1653. American Institute of Aeronautics and Astronautics, San Diego, California, USA
81. Parker ASR, 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 (ICORR)*, pp 1239–1246
82. Serpell JA, Hsu Y (2001) Development and validation of a novel method for evaluating behavior and temperament in guide dogs. *Appl Anim Behav Sci* 72(4):347–364
83. Markoff J (2015) *Machines of loving grace: the quest for common ground between humans and robots*. ECCO, an imprint of HarperCollinsPublishers, New York, NY
84. Langley P (1997) Machine learning for adaptive user interfaces. In: *Brewka G, Habel C, Nebel B (eds) KI-97 Adv Artif Intell*. Springer, Berlin, Heidelberg, pp 53–62
85. Vesper C, Butterfill S, Knoblich G, Sebanz N (2010) A minimal architecture for joint action. *Neural Netw* 23(8–9):998–1003
86. Argall BD, Chernova S, Veloso M, Browning B (2009) A survey of robot learning from demonstration. *Robot Auton Syst* 57(5):469–483
87. Mathewson KW, Pilarski PM (2022) A brief guide to designing and evaluating human-centered interactive machine learning. In: *Machine Learning Evaluation Standards Workshop at ICLR 2022*
88. Kaplan F, Oudeyer P-Y, Kubinyi E, Miklósi Á (2002) Robotic clicker training. *Robot Auton Syst* 38(3):197–206
89. Knox WB, Stone P (2009) Interactively shaping agents via human reinforcement: The tamer framework. In: *Proceedings of the Fifth International Conference on Knowledge Capture. K-CAP '09*, pp 9–16. Association for Computing Machinery, New York, NY, USA
90. Knox WB, Stone P, Breazeal C (2013) Training a robot via human feedback: a case study. In: *Herrmann G, Pearson MJ, Lenz A, Bremner P, Spiers A, Leonards U (eds) Int Conf Soc Robot*. Springer, USA, pp 460–470
91. Griffith S, Subramanian K, Scholz J, Isbell CL, Thomaz AL (2013) Policy shaping: integrating human feedback with reinforcement learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*
92. Loftin R, MacGlashan J, Peng B, Taylor ME, Littman ML, Huang J, Roberts DL (2014) A strategy-aware technique for learning behaviors from discrete human feedback. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence. AAAI'14*, pp 937–943
93. Chao C, Cakmak M, Thomaz AL (2010) Transparent active learning for robots. In: *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction. HRI '10*, pp 317–324
94. Judah K, Roy S, Fern A, Dietterich T (2010) Reinforcement learning via practice and critique advice. *Proc AAAI Conf Artif Intell* 24:481–486
95. Lin LJ (1991) Programming robots using reinforcement learning and teaching. In: *AAAI*, pp 781–786
96. Lin LJ (1993) Hierarchical learning of robot skills by reinforcement. In: *IEEE International Conference on Neural Networks*, pp 181–1861
97. Abramson J, Ahuja A, Brussee A, Carnevale F, Cassin M, Clark S, Dudzik A, Georgiev P, Guy A, Harley T, Hill F, Hung A, Kenton Z, Landon J, Lillicrap TP, Mathewson KW, Muldal A, Santoro A, Savinov N, Varma V, Wayne G, Wong N, Yan C,

- Zhu R (2020) Imitating interactive intelligence. CoRR abs/2012.05672
98. Hedlund E, Johnson M, Gombolay M (2021) The effects of a robot's performance on human teachers for learning from demonstration tasks. In: Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, pp 207–215
99. Resnik L (2011) Development and testing of new upper-limb prosthetic devices: research designs for usability testing. *J Rehabil Res Dev* 48(6):697–706
100. Resnik L, Meucci MR, Lieberman-Klinger S, Fantini C, Kelty DL, Disla R, Sasson N (2012) Advanced upper limb prosthetic devices: implications for upper limb prosthetic rehabilitation. *Arch Phys Med Rehabil* 93(4):710–717
101. Hebert JS, Wolfe DL, Miller WC, Deathe AB, Devlin M, Pallaveshi L (2009) Outcome measures in amputation rehabilitation: ICF body functions. *Disabil Rehabil* 31(19):1541–1554
102. Light CM, Chappell PH, Kyberd PJ (2002) Establishing a standardized clinical assessment tool of pathologic and prosthetic hand function: normative data, reliability, and validity. *Arch Phys Med Rehabil* 83(6):776–783
103. Mathiowetz VG, Volland G, Kashman N, Weber K (1985) Adult norms for the box and block test of manual dexterity. *Am J occup Ther Off Publ Am Occup Ther Assoc* 39(6):386–91
104. Marasco PD, Hebert JS, Sensinger JW, Beckler DT, Thumser ZC, Shehata AW, Williams HE, Wilson KR (2021) Neuro-robotic fusion of prosthetic touch, kinesthesia, and movement in bionic upper limbs promotes intrinsic brain behaviors. *Sci Robot* 6
105. 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
106. Williams HE, Chapman CS, Pilarski PM, Vette AH, Hebert JS (2021) Myoelectric prosthesis users and non-disabled individuals wearing a simulated prosthesis exhibit similar compensatory movement strategies. *J NeuroEng Rehabil* 18
107. Hebert JS, Boser QA, Valevicius AM, Tanikawa H, Lavoie EB, Vette AH, Pilarski PM, Chapman CS (2019) Quantitative eye gaze and movement differences in visuomotor adaptations to varying task demands among upper-extremity prosthesis users. *JAMA Netw Open* 2

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.