



Social Neuro AI: Social Interaction as the “Dark Matter” of AI

Samuele Bolotta¹ and Guillaume Dumas^{2*}

¹ Department of Computer Science and Operations Research, Université de Montréal, Montreal, QC, Canada, ² Department of Psychiatry, CHU Sainte-Justine Research Center, Mila-Quebec Artificial Intelligence Institute, Université de Montréal, Montreal, QC, Canada

This article introduces a three-axis framework indicating how AI can be informed by biological examples of social learning mechanisms. We argue that the complex human cognitive architecture owes a large portion of its expressive power to its ability to engage in social and cultural learning. However, the field of AI has mostly embraced a solipsistic perspective on intelligence. We thus argue that social interactions not only are largely unexplored in this field but also are an essential element of advanced cognitive ability, and therefore constitute metaphorically the “dark matter” of AI. In the first section, we discuss how social learning plays a key role in the development of intelligence. We do so by discussing social and cultural learning theories and empirical findings from social neuroscience. Then, we discuss three lines of research that fall under the umbrella of Social NeuroAI and can contribute to developing socially intelligent embodied agents in complex environments. First, neuroscientific theories of cognitive architecture, such as the global workspace theory and the attention schema theory, can enhance biological plausibility and help us understand how we could bridge individual and social theories of intelligence. Second, intelligence occurs in time as opposed to over time, and this is naturally incorporated by dynamical systems. Third, embodiment has been demonstrated to provide more sophisticated array of communicative signals. To conclude, we discuss the example of active inference, which offers powerful insights for developing agents that possess biological realism, can self-organize in time, and are socially embodied.

Keywords: social interaction, cognitive architecture, virtual agents, social learning, Neuro-AI, neurodynamics, self-organization, Alan Turing

OPEN ACCESS

Edited by:

Adam Safron,
Johns Hopkins Medicine,
United States

Reviewed by:

James Derek Lomas,
Delft University of Technology,
Netherlands
Emily S. Cross,
Macquarie University, Australia

*Correspondence:

Guillaume Dumas
guillaume.dumas@centraliens.net

Specialty section:

This article was submitted to
Human-Media Interaction,
a section of the journal
Frontiers in Computer Science

Received: 31 December 2021

Accepted: 11 April 2022

Published: 06 May 2022

Citation:

Bolotta S and Dumas G (2022) Social
Neuro AI: Social Interaction as the
“Dark Matter” of AI.
Front. Comput. Sci. 4:846440.
doi: 10.3389/fcomp.2022.846440

1. THE IMPORTANCE OF SOCIAL LEARNING

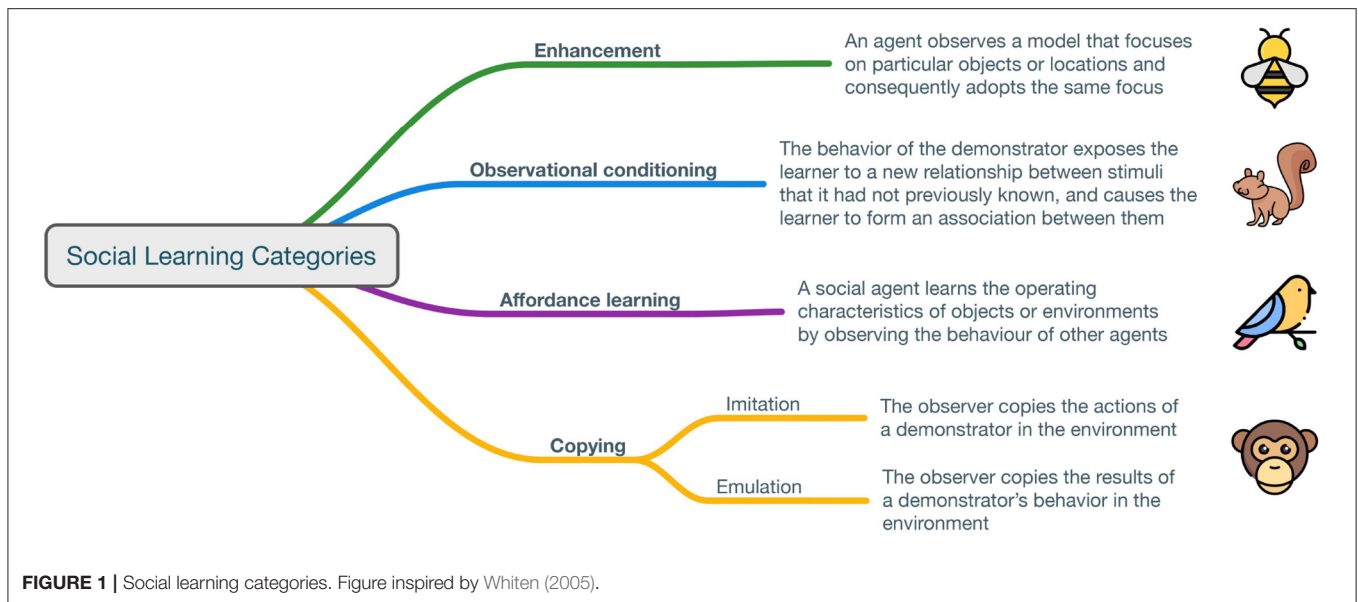
1.1. Social Learning Categories

Various approaches have been proposed in order to reach a human-like level of intelligence. For example, some argue that scaling foundational models (self-supervised pretrained deep network models), data and compute can lead to such kind of intelligence (Bommasani et al., 2021; Yuan et al., 2022). Others argue that attention, understood as a dynamical control of information flow (Mittal et al., 2020), is all we need. Transformers have proposed a general purpose architecture where inductive biases shaping the flow of information are learned from the data itself (Vaswani et al., 2017); this architecture can be applied to various domains ranging from sequence learning to visual processing and time-series forecasting. Others argue that by having a complex enough environment, any

reward should be enough to elicit some complex behavior and end up in intelligent behavior that subserves the maximization of such reward. This therefore discards the idea that specialized problem formulations are needed for each ability (Silver et al., 2021). Our proposal stems from the idea that human cognitive functions such as theory of mind (the capacity to understand other people by ascribing mental states to them) and explicit metacognition (the capacity to reflect on and justify our behavior to others) are not genetically programmed, but rather constructed during development through social interaction (Heyes, 2018). Since their birth, social animals use their conspecifics as vehicles for gathering information that can potentially help them respond efficiently to challenges in the environment, avoiding harm and maximizing rewards (Kendal et al., 2018). Learning adaptive information from others results in better regulation of task performance, especially by gaining fitness benefits and in avoiding some of the costs associated with asocial, trial-and-error learning, such as time loss and energy loss as well as exposure to predation (Clark and Dumas, 2016). Importantly, cultural inheritance permeates a broad array of behavioral domains, including migratory pathways, foraging techniques, nesting sites and mates (Whiten, 2021). The spread of such information across generations gives social learning a unique role in the evolution of culture and therefore makes it a crucial candidate to investigate the biological bases of human cognition (Gariépy et al., 2014). In the current paper, we do not focus extensively on the differences between social learning in humans and in other animals as the cognitive processes used in acquiring behavior seem to be very similar across a wide range of species (Heyes, 2012). What sets humans apart from other animals, however, is: a) social learning in humans is highly rewarded from early infancy (Nielsen et al., 2012), b) the nature of the inputs surrounding humans is way more complex than for other animals (Heyes, 2012). According to the ontogenetic adaptation hypothesis (Tomasello, 2020), human infant's unique social-cognitive skills are the result of shared intentionality (capacity to share attention and intention) and are adaptations for life in a cultural group—with individuals coordinating, communicating and learning from each other in several ways. Recent reviews have identified four main categories of social learning that differ in what is socially learnt and in the cognitive skills that are required (Hoppitt and Laland, 2008; Whiten, 2021) (**Figure 1**). These categories have been developed through the approach of behaviorism. While we acknowledge that there is more to social learning than mere behavior (the affective and cognitive dimensions are equally crucial Gruber et al., 2021), we keep it as the focus of this short article because it is an empirically solid starting point with clarified mechanisms. The purpose of this section, then, is to give an example of social learning mechanisms that are common across multiple species and can be understood as a natural form of Social Neuro-AI. Moreover, this section aims at demonstrating how social interactions are a key component of biological intelligence; we make the case that they might be of inspiration for the development of socially intelligent artificial agents that can cooperate efficiently with humans and with each other. In other words, although there are examples of social agents (chatbots, non-player characters in video games,

social robots), we argue that social interactions still remain the “dark matter” of the field. These social behaviors often emerge from a Piagetian perspective on human intelligence. As argued by Kovač et al. (2021), mainstream Deep Reinforcement Learning research sees intelligence as the product of the individual agent's exploration of the world; it mainly focuses on sensorimotor development and problems involving interaction with inanimate objects rather than social interactions with animate agents. This approach can and has given rise to apparent social behaviors, but we argue that this is not the best approach, as it does not involve any focus on the genuine social mechanisms *per se* (Dumas et al., 2014a). Instead, it sees social behaviors as a collateral effect of the intelligence of a solitary thinker. For this reason, as Schilbach et al. (2013) argued a decade ago that social interactions were the “dark matter” of cognitive neuroscience, here we argue that social interactions can also be considered metaphorically as the “dark matter” of AI (Schilbach et al., 2013). Indeed, more than being a rather unexplored topic, social interactions can constitute a critical missing piece for the understanding and modeling of advanced cognitive abilities.

At the most elementary level, enhancement consists of an agent observing a model that focuses on particular objects or locations and consequently adopting the same focus (Thorpe, 1963; Heyes, 1994). For example, it was demonstrated that bees outside the nest land more often on flowers that they had seen preferred by other bees (Worden and Papaj, 2005). This skill requires social agents to perform basic associative learning in relation to other agents' observed actions; it is likely to be the most widespread form of social learning across the animal kingdom. A more complex form of social learning consists of observational conditioning, which exposes a social agent to a relationship between stimuli (Heyes, 1994); this exposure causes a change in the agent. For example, the observation of experienced demonstrators facilitated the opening of hickory nuts by red squirrels, relative to trial-and-error learning (Weigl and Hanson, 1980). This is therefore a mechanism through which agents learn the value of a stimulus from the interaction with other agents. Yet a more complex form of social learning consists of affordance learning, which allows a social agent to learn the operating characteristics of objects or environments by observing the behavior of other agents (Whiten, 2021). For example, pigeons that saw a demonstrator push a sliding screen for food made a higher proportion of pushes than observers in control conditions, thus exhibiting affordance learning (Klein and Zentall, 2003). In other words, the animals perceive the environment partly in terms of the action opportunities that it provides. Finally, at the most complex level, copying another individual can take the shape of pure imitation, where every detail is copied, or emulation, where only a few elements are copied (Byrne, 2002). For example, most chimpanzees mastered a new technique for obtaining food when they were under the influence of a trained expert, whereas none did so in a population lacking an expert (Whiten, 2005). As to what is required for imitation, there are debates in the literature ranging from the distinctions between program-level and production-level imitation (Byrne, 2002) to the necessity of pairing Theory of Mind (ToM) with behavioral imitation to obtain “true” imitation (Call et al., 2005).



We refer the reader to Breazeal and Scassellati (2002) for a more detailed discussion of imitation in robots.

1.2. Social Learning Strategies

Crucially, while social learning is widespread, using it indiscriminately is rarely beneficial. This suggests that individuals should be selective in what, when, and from whom they learn socially, by following “social learning strategies” (SLSs; Kendal et al., 2018). Several SLSs might be used by the same population and even by the same individual. The aforementioned categories of social learning have been shown to be refined by modulating biases that can strengthen their adaptive power (Kendal et al., 2018). For example, an important SLS is copying when asocial learning would be costly; research has shown that, when task difficulty increases, various animals are more likely to use social information. Individuals also prefer using social information when they are uncertain about a task; high-fidelity copying is observed among children who lack relevant personal information (Wood et al., 2013). In general, other state-based SLSs can affect the decision to use social information, such as age, social rank, and reproductive state of the learner; for example, low- and mid-ranking chimpanzees are more likely to use social information than high-ranking individuals (Kendal et al., 2015). Model-based biases are another crucial category; for example, children prefer to copy prestigious individuals, where status is evidenced by their older age, popularity and social dominance (Flynn and Whiten, 2012). Multiple evidence also suggests that a conformist transmission bias exists, whereby the behavior of the majority of individuals is more likely to be adopted by others (Kendal et al., 2018).

1.2.1. Social Learning in Neuroscience

We have presented evidence that social learning is a crucial hallmark of many species and it manifests itself across different behavioral domains; without it, animals would lose the possibility

to quickly acquire valuable information from their conspecifics and therefore lose fitness benefits. However, one important question is: how does the brain mediate social processes and behavior? Despite the progress made in social neuroscience and in developmental psychology, only in the last decade, serious efforts have started focusing on the answer to this question—as neural mechanisms of social interaction were seen as the “dark matter” of social neuroscience (Schilbach et al., 2013); recently, a framework for computational social neuroscience has been proposed, in an attempt to naturalize social interaction (Tognoli et al., 2018). At the intra-brain level, it was demonstrated that social interaction is categorically different from social perception and that the brain exhibits different activity patterns depending on the role of the subject and on the context in which the interaction is unfolding (Dumas et al., 2012). At the inter-brain level, functional Magnetic Resonance Imaging (fMRI) or Electroencephalography (EEG) recordings of multiple brains (i.e., hyperscanning) have allowed to demonstrate inter-brain synchronization during social interaction—specifically, while subjects were engaged in spontaneous imitation of hand movements (Dumas et al., 2010). Interestingly, the increase in coupling strength between brain signals was also shown to be present during a two-person turn-taking verbal exchange with no visual contact, in both a native or a foreign language context (Pérez et al., 2019). Inter-brain synchronization is also modulated by the type of task and by the familiarity between subjects (Djalovski et al., 2021). Overall, this shows that, beyond their individual cognition, humans are also coupled in the social dimension. Interestingly, the field of computational social neuroscience has also focused on explaining the functional meaning of such correlations between inter-brain synchronization and behavioral coupling. A biophysical model showed that the similarity of both endogenous dynamics and anatomical structure might facilitate inter-individual synchronization and explain our propensity to

socially bind with others via perception and actions (Dumas et al., 2012). More specifically, the connectome, a wiring diagram that maps all neural connections in the brain, not only facilitates the integration of information within brains, but also between brains. In those simulations, tools from dynamical systems thus suggest that beyond their individual cognition, humans are also dynamically coupled in the social realm (Dumas et al., 2012).

1.2.2. Social Learning and Language Development

Regarding language development in humans, cognitive and structural accounts of language development have often conceptualized linguistic abilities as static and formal sets of knowledge structures, ignoring the contextual nature of language. However, good communication must be tailored to the characteristics of the listener and of the context—language can also be explained as a social construct (Whitehurst, 1978). For example, evidence shows that the language outcome of children with cochlear implants is heavily influenced by parental linguistic input during the first years after cochlear implantation (Holzinger et al., 2020). In terms of specific social learning variables, imitation has also been shown to play a major role in boosting language development, usually in the form of selective imitation (Whitehurst et al., 1974; Whitehurst and Vasta, 1975). Moreover, in children with autism spectrum disorder, social learning variables such as joint attention, immediate imitation, and deferred imitation have been shown to be the best predictors of language ability and rate of communication development (Toth et al., 2006). These results clearly suggest that social learning skills have an influence on language acquisition in humans.

2. STEPS TOWARD SOCIAL NEURO AI

How Could Social Learning Be Useful for AI?

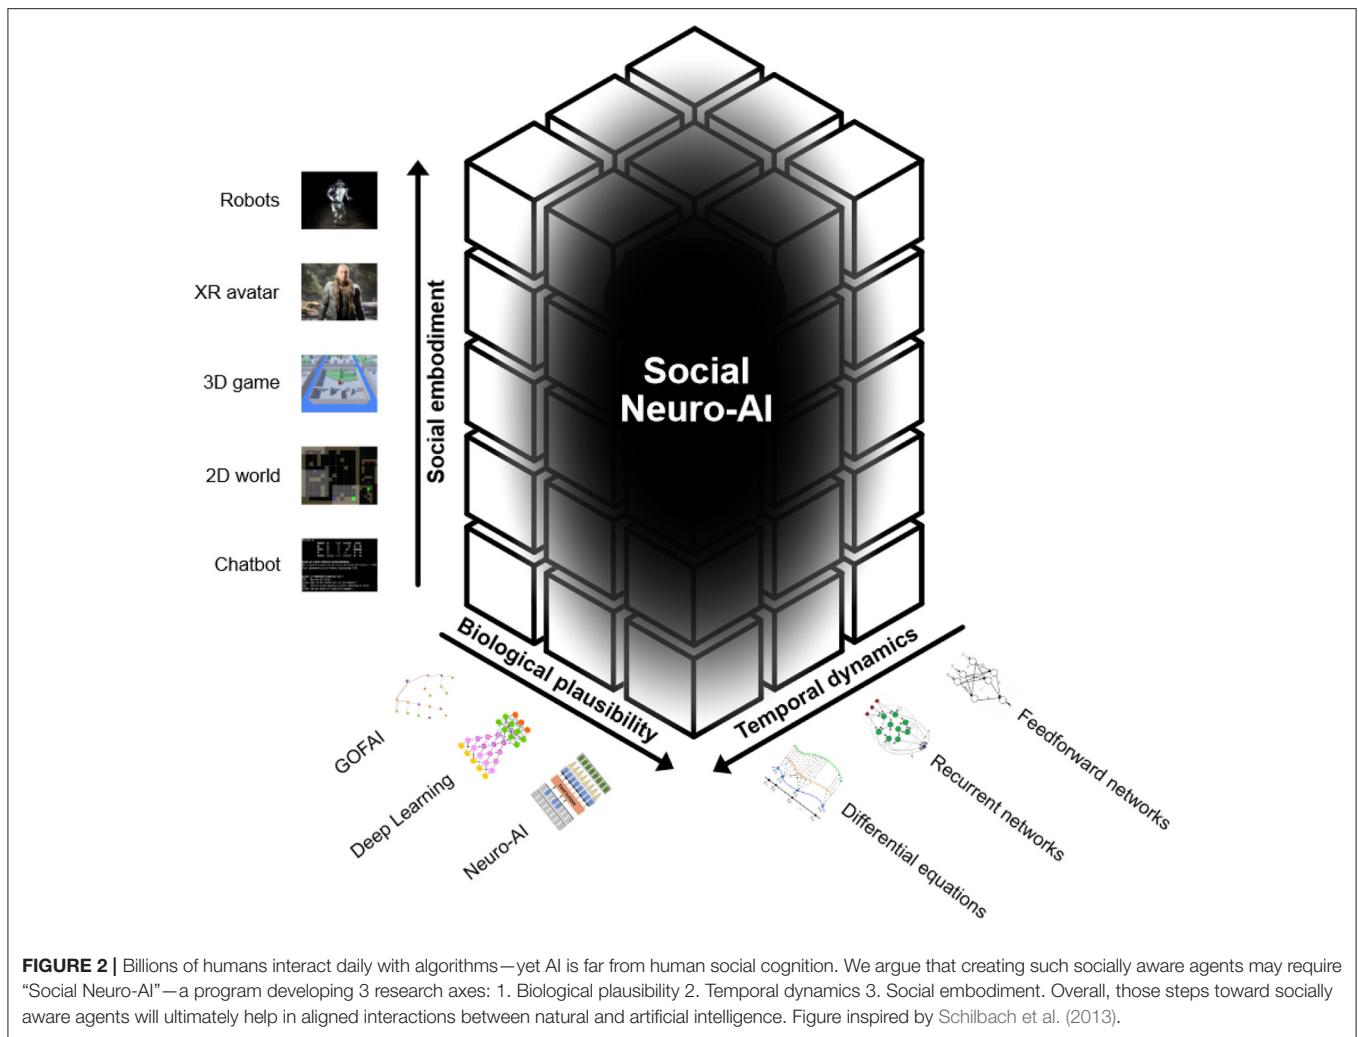
In the previous sections, we have provided convincing evidence that interpersonal intelligence enhances intrapersonal intelligence through the mechanisms and biases of social learning. It is a crucial aspect of biological intelligence that possesses a broad array of modulating biases meant to strengthen its adaptive power. Recent efforts in computational social neuroscience have paved the way for a naturalization of social interactions.

Multi-agent reinforcement learning (MARL) is the best subfield of AI to investigate the interactions between multiple agents. Such interactions can be of three types: cooperative games (all agents working for the same goal), competitive games (all agents competing against each other), and mixed motive games (a mix of cooperative and competitive interactions). At each timestep t , each agent is attempting to maximize its own reward by learning a policy that optimizes the total expected discounted future reward. We refer the reader to high-quality reviews that have been written on MARL (Hernandez-Leal et al., 2019; Nguyen et al., 2020; Wong et al., 2021). Here, we highlight that, among others, low sample efficiency is one of the greatest challenges for MARL, as millions of interactions with the environment are usually needed for agents to learn. Moreover, multi-agent joint action space increases exponentially with the

number of agents, leading to problems that are often intractable. In the last few years, part of the AI community has already started demonstrating that these problems can be alleviated by mechanisms that allow for social learning (Jaques, 2019; Ndousse et al., 2021). For example, rewarding agents for having a causal influence over other agents' actions leads to enhanced coordination and communication in challenging social dilemma environments (Jaques et al., 2019) and rewarding agents for coordinating attention with another agent improves their ability of coordination, by reducing the cost of exploration (Lee et al., 2021). More in general, concepts from complex systems such as self-organization, emergent behavior, swarm optimization and cellular systems suggest that collective intelligence could produce more robust and flexible solutions in AI, with higher sample efficiency and higher generalization (Ha and Tang, 2021). In the following sections, we argue that to exploit all benefits that social learning can offer AI and robotics, more focus on biological plausibility, social embodiment and temporal dynamics is needed. Studies have focused on the potential of conducting research at the intersection of some of these three axes (Kerzel et al., 2017; Husbands et al., 2021). Moreover, it is worth noticing that (Dumas et al., 2012; Heggli et al., 2019) offer a tentative glimpse of what the intersection of the three axes would look like—both using dynamical systems with computational simulations to address falsifiable scientific questions associated with the idea of social embodiment.

2.1. Biological Plausibility

Biological plausibility refers to the extent to which an artificial architecture takes inspiration from empirical results in neuroscience and psychology. The social learning skills and biases that we have shown so far are boosted in humans by their advanced cognitive architecture (Whiten, 2021). Equipping artificial agents with complex social learning abilities will therefore require more complex architectures that can handle a great variety of information efficiently. This is exactly what "Neuro-AI" aims at: drawing on how evolution has shaped the brain of humans and of other animals in order to create more robust agents (Figure 2). While the human unconscious brain aligns well with the current successful applications of deep learning, the conscious brain involves higher-order cognitive abilities that perform much more complex computations than what deep learning can currently do (Bengio, 2019). More specifically, "unconsciousness" is where most of our intelligence lies and involves unconscious abilities related to view-invariance, meaning extraction, control, decision-making and learning; "i-consciousness" is the part of human consciousness that is focused on integrating all available evidence to converge toward a single decision; "m-consciousness" is the part of human consciousness that is focused on reflexively representing oneself, utilizing error detection, meta-memory and reality monitoring (Graziano, 2017). Notably, recent efforts in the deep learning community have indeed focused on Neuro-AI: building advanced cognitive architectures that are inspired from neuroscience. In particular, the global workspace theory (GWT) is the most widely accepted theory of consciousness, and it postulates that when a piece of information is selected by attention, it may non-linearly



achieve “ignition,” enter the global workspace (GLW) and be shared across specialized cortical modules, therefore becoming conscious (Baars, 1993; Dehaene et al., 1998). The use of such a communication channel in the context of deep learning was explored for modeling the structure of complex environments. This architecture was demonstrated to encourage specialization and compositionality and to facilitate the synchronization of otherwise independent specialists (Goyal et al., 2021). Moreover, inductive biases inspired by higher-order cognitive functions in humans have been shown to improve OOD generalization. Overall, this section proposes that we draw inspiration from one structure we know is capable of comprehensive intelligence capable of perception, planning, and decision making: the human brain (Figure 2). For a more extensive discussion on biological plausibility in AI, we refer the reader to Hassabis et al. (2017) and Macpherson et al. (2021).

2.2. Temporal Dynamics

Figure 2 more specifically, FFNs allow signals to travel only from input to output, whereas RNNs can have signals traveling in both directions and therefore introduce loops in the network.

Incorporating differential equations in a RNN (continuous-time recurrent neural network) can help learn long-term dependencies (Chang et al., 2019) and model more complex phenomena, such as the effects of incoming inputs on a neuron. Moreover, viewing RNNs as a discretization of ordinary differential equations (ODEs) driven by input data has led to gains in reliability and robustness to data perturbations (Lim et al., 2021b). This becomes clear when one notices that many fundamental laws of physics and chemistry can be formulated as differential equations. In general, differential equations are expected to contribute to shifting the perspective from representation-centered to self-organizing agents (Brooks, 1991). The former view has been one predominant way of thinking about autonomous systems that exhibit intelligent behavior: such autonomous agents use their sensors to extract information about the world they operate in and use it to construct an internal model of the world and therefore rationally perform optimal decision making in pursuit of some goal. In other words, autonomous agents are information processing systems and their environment can be abstracted away as the source of answers to questions raised by the ongoing agents’ needs.

Cognitive processes are thought to incorporate representational content and to acquire such contents via inferential processes instantiated by the brain. Importantly, according to this view, the sensorimotor connections of the agents to the environment are still relevant to understand their behavior, but there is no focus on what such connections involve and how they take place (Newell and Simon, 1976). The latter view, in line with the subsumption architecture introduced by Brooks (1991), shows how the representational approach ignores the nonlinear dynamical aspect of intelligence, that is, the temporal constraints that characterize the interactions between agent and environment. Instead, dynamics is a powerful framework that has been used to describe multiple natural phenomena as an interdependent set of coevolving quantitative variables (van Gelder, 1998) and a crucial aspect of intelligence is that it occurs in time and not over time. If we abstract away the richness of real time, then we also change the behavior of the agents (Smithers, 2018). In other words, one should indeed focus on the structural complexity and on the algorithmic computation the agents need to carry out, but without abstracting away the dynamical aspects of the agent-environment interactions: such dynamical aspects are pervasive and, therefore, necessary to explain the behavior of the system (van Gelder, 1998; Barandiaran, 2017; Smithers, 2018).

2.3. Social Embodiment

There has been a resurgence of enactivism in cognitive neuroscience over the past decade, emphasizing the circular causality induced by the notion that the environment is acting upon the individual and the individual is acting upon the environment. To understand how the brain works, then one has to acknowledge that it is embodied (Clark, 2013; Hohwy, 2013). Evidence for this shows that embodied intelligence in human children arises from the interaction of the child with the environment through a sensory body that is capable of recognizing the statistical properties of such interaction (Smith and Gasser, 2005). Moreover, higher primates interpret each other as psychological subjects based on their bodily presence; social embodiment is the idea that the embodiment of a socially interactive agent plays a significant role in social interactions. It refers to “states of the body, such as postures, arm movements, and facial expressions, that arise during social interaction and play central roles in social information processing.” Thompson and Varela (2001) and Barsalou et al. (2003). This includes internal and external structures, sensors, and motors that allow them to interact actively with the world. We argue that robots are more socially embodied than digital avatars for a simple reason: they have a higher potential to use parts of their bodies to communicate and to coordinate with other agents (Figure 2). At a high level, sensorimotor capabilities in the avatar and robots are meant to model their role in biological beings: the agent now has limitations in the ways they can sense, manipulate, and navigate its environments. Importantly, these limitations are closely tied to the agent’s function (Deng et al., 2019). The idea of social embodiment in artificial agents is supported by evidence of improvements in the interactions between embodied agents and humans (Zhang et al., 2016). Studies have shown positive

effects of physical embodiment on the feeling of an agent’s social presence, the evaluation of the agent, the assessment of public evaluation of the agent, and the evaluation of the interaction with the agent (Kose-Bagci et al., 2009; Gupta et al., 2021). In robots, social presence is a key component in the success of social interactions and it can be defined as the combination of seven abilities that enhance a robot/s social skills: 1. Express emotion, 2. Communicate with high-level dialogue, 3. Learn/recognize models of other agents, 4. Establish/maintain social relationships, 5. Use natural cues, 6. Exhibit distinctive personality and character, and 7. Learn/develop social competencies (Lee, 2006). Social embodiment thus equips artificial agents with a more articulated and richer repertoire of expressions, ameliorating the interactions with it (Jaques, 2019). For instance, in human-robot interaction, a gripper is not limited to its role in the manipulation of objects. Rather, it opens a broad array of movements that can enhance the communicative skills of the robot and, consequently, the quality of its possible interactions (Deng et al., 2019). The embodied agent is therefore the best model of the aspects of the world relevant to its surviving and thriving, through performing situationally appropriate actions (Ramstead et al., 2020) (Figure 2). Therefore, it will be crucial to scale up the realism of what the agents perceive in their social context, going from simple environments like GridWorld to more complex ones powered by video-game engines and, finally, to extremely realistic environments, like the one offered by the MetaHuman Creator of Unreal Engine. In parallel, greater focus is needed on the mental processes supporting our interactions with social machines, so as to develop a more nuanced understanding of what is ‘social’ about social cognition (Cross and Ramsey, 2021) and to gather insights critical for optimizing social encounters between humans and robots (Henschel et al., 2020). For a more extensive discussion on embodied intelligence, we refer the reader to Roy et al. (2021). These advancements will hopefully result in more socially intelligent agents and therefore in more fruitful interactions between humans and virtual agents.

3. ACTIVE INFERENCE

The active inference framework represents a biologically realistic way of moving away from rule-governed manipulation of internal representations to action-oriented and situationally appropriate cognition (Friston et al., 2006). More specifically, active inference can be seen as a self-organizing process of action policy selection (Ramstead et al., 2020), which a) concerns the selective sampling of the world by an embodied agent and b) instantiates in a generative model the goal of minimizing their surprise through perception and action (Ramstead et al., 2020). In other words, generative models do not encode exploitable and symbolic structural information about the world, because cognition does not perform manipulation of internal representations, but rather instantiates control systems that are expressed in embodied activity and utilize information encoded in the approximate posterior belief (Ramstead et al., 2020). Interestingly, by grounding GWT within the embodied perspective of the active inference framework, the Integrated

World Modeling Theory (IWMT) suggests that conscious experience can only result from autonomous embodied agents with global workspaces that generate integrative models of the world with spatial, temporal and causal coherence (Safron, 2020).

Active inference models are still very discrete in their architectures, especially regarding high-level cognitive aspects, but they may be a good class of models to raise the tension between computation and implementation (Figure 2). Therefore, they only have been able to handle small policies and state-spaces, while also requiring the environmental dynamics to be well known. However, using deep neural networks to approximate key densities, the agent can scale to more complex tasks and obtain performance comparable to common reinforcement learning baselines (Millidge, 2020). Moreover, one advantage of active inference is that the associated biologically inspired architectures predict future trajectories of the agent N steps forward in time, rather than just at the next step. By sampling from these trajectories, the variance of the decision is reduced (Millidge, 2020).

Interestingly, by grounding GWT within the embodied perspective of the active inference framework, the Integrated World Modeling Theory (IWMT) suggests that complexes of integrated information and global workspaces can entail conscious experiences if (and only if) they are capable of generating integrative world models with spatial, temporal, and causal coherence. These ways of categorizing experience are increasingly recognized as constituting essential “core knowledge” at the foundation of cognitive development (Spelke and Kinzler, 2007). In addition to space, time, and cause, IWMT adds embodied autonomous selfhood as a precondition for integrated world modeling.

4. A DETAILED PROPOSAL: HOW CAN INCREASED BIOLOGICAL PLAUSIBILITY ENHANCE SOCIAL AFFORDANCE LEARNING IN ARTIFICIAL AGENTS?

Attention has become a common ingredient in deep learning architectures. It can be understood as a dynamical control of information flow (Mittal et al., 2020). In the last decade, transformers have demonstrated how *attention may be all we need*, obtaining excellent performances in sequence learning (Vaswani et al., 2017), visual processing (Dosovitskiy et al., 2020) and time-series forecasting (Lim et al., 2021a). While transformers proposed a general purpose architecture where inductive biases shaping the flow of information are learned from the data itself, we can imagine a higher-order informational filter built on top of attention: an *Attention Schema (AS)*, namely a descriptive and predictive model of attention. In this regard, the attention schema theory (AST) is a neuroscientific theory that postulates that the human brain, and possibly the brain of other animals, does construct a model of attention: an attention schema (Graziano and Webb, 2015). Specifically, the proposal is that the brain constructs not only a model of the physical body but also a coherent, rich, and descriptive model of attention. The body schema contains layers of valuable

information that help control and predict stable and dynamic properties of the body; in a similar fashion, the attention schema helps control and predict attention. One cannot understand how the brain controls the body without understanding the body schema, and in a similar way one cannot understand how the brain controls its limited resources without understanding the attention schema (Graziano, 2017). The key reason a higher-order filter on top of attention seems a promising idea for deep learning comes from control engineering: a good controller contains a model of the item being controlled (Conant and Ross Ashby, 1970). More specifically, a descriptive and predictive model of attention could help the dynamical control of attention and therefore maximize the efficiency with which resources are strategically devoted to different elements of an ever-changing environment (Graziano, 2017). Indeed, the performance of an artificial agent in solving a simple sensorimotor task is greatly enhanced by an attention schema, but its performance is greatly reduced when the schema is not available (Wilterson and Graziano, 2021). Therefore, the study of consciousness in artificial intelligence is not a mere pursuit of metaphysical mystery; from an engineering perspective, without understanding subjective awareness, it might not be possible to build artificial agents that intelligently control and deploy their limited processing resources. It has also been argued that, without an attention schema, it might be impossible to build artificial agents that are socially intelligent. This idea stems from the evidence that points at an overlap of social cognition functions with awareness and attention functions in the right temporoparietal junction of the human brain (Mitchell, 2008). It was then proposed that an attention schema might also be used for social cognition, giving rise to an overlap between modeling one’s own attention and modeling others’ attention. In other words, when we attribute to other people an awareness of their surroundings, we are constructing a simplified model of their attention—a schema of others’ attention (Graziano and Kastner, 2011). Indeed, such a model would enhance the ability of the agent to predict social affordances in real time, which is a goal the field has been trying to achieve in different ways (Shu et al., 2016; Ardón et al., 2021). Without a model of others’ attention, even if we had detailed information about them, we could not predict their behavior on a moment-by-moment basis. However, with a component that tracks how and where other agents are focusing their resources in the environment, the probabilities for many affordances in the environment become computable in real time (Graziano, 2019). Specifically, there are three predictions that are investigated in this proposal. The first prediction is that, without an attention schema, attention is still possible, but it suffers deficits in control and thus leads to worse performance. The second prediction is that an attention schema is useful for modeling the attention of other agents as well—as the machinery that computes information about other people’s attention is the same machinery that computes information about our own attention (Graziano and Kastner, 2011). The third prediction is that an agent equipped with an attention schema is going to have better OOD generalization than a classic Proximal Policy Optimization agent (Schulman et al., 2017), especially in environments in which the ability to

intelligently control and deploy limited processing resources is necessary.

5. CONCLUSION

At the crossroads of robotics, computer science, psychology, and neuroscience, one of the main challenges for humans is to build autonomous agents capable of participating in cooperative social interactions. This is important not only because AI will play a crucial role in daily life well into the future, but also because, as demonstrated by results in social neuroscience and evolutionary psychology, intrapersonal intelligence is tightly connected with interpersonal intelligence, especially in humans (Dumas et al., 2014b). In this opinion article, we have proposed an approach that unifies three lines of research that, at the moment, are separated from each other; in particular, we have proposed three research directions that are expected to enhance efficient exchange of information between agents. Biological plausibility attempts to increase the robustness and OOD generalization of algorithms by drawing on knowledge about biological brains; temporal dynamics attempts to better capture long-term temporal dependencies; social embodiment proposes that states of the body that arise during social interaction play central roles in social information

processing. Unifying these axes of research would contribute to creating agents that are able to cooperate efficiently in extremely complex and realistic environments (Dennis et al., 2021), while interacting with other embodied agents and with humans.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

GD designed the theoretical framework and initial outline. SB wrote the first draft. SB and GD wrote the final version. Both authors contributed to the article and approved the submitted version.

FUNDING

GD is funded by the Institute for Data Valorization (IVADO; CF00137433), Montreal, and the Fonds de recherche du Québec (FRQ; 285289).

REFERENCES

- Ardón, P., Pairet, R., Lohan, K. S., Ramamoorthy, S., and Petrick, R. P. A. (2021). Building affordance relations for robotic agents—a review. *arXiv:2105.06706 [cs]*. arXiv: 2105.06706. doi: 10.24963/ijcai.2021/590
- Baars, B. J. (1993). *A Cognitive Theory of Consciousness*. Cambridge; Cambridge University Press. Google-Books-ID: 7w6IYeJRqyoC.
- Barandiaran, X. E. (2017). Autonomy and enactivism: towards a theory of sensorimotor autonomous agency. *Topoi* 36, 409–430. doi: 10.1007/s11245-016-9365-4
- Barsalou, L. W., Niedenthal, P. M., Barbey, A. K., and Ruppert, J. A. (2003). “Social embodiment,” in *Psychology of Learning and Motivation-Advances in Research and Theory* (San Diego, CA), 43–92.
- Bengio, Y. (2019). The consciousness prior. *arXiv:1709.08568 [cs, stat]*. arXiv: 1709.08568.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. *arXiv [Preprint]*. arXiv: 2108.07258. Available online at: <http://arxiv.org/abs/2108.07258>
- Breazeal, C., and Scassellati, B. (2002). Robots that imitate humans. *Trends Cogn. Sci.* 6, 481–487. doi: 10.1016/S1364-6613(02)02016-8
- Brooks, R. A. (1991). Intelligence without representation. *Artif. Intell.* 47, 139–159. doi: 10.1016/0004-3702(91)90053-M
- Byrne, R. W. (2002). *Imitation of Novel Complex Actions: What Does the Evidence From Animals Mean?* San Diego, CA: Academic Press.
- Call, J., Carpenter, M., and Tomasello, M. (2005). Copying results and copying actions in the process of social learning: chimpanzees (pan troglodytes) and human children (homo sapiens). *Anim. Cogn.* 8, 151–163. doi: 10.1007/s10071-004-0237-8
- Chang, B., Chen, M., Haber, E., and Chi, E. H. (2019). Antisymmetricrnn: a dynamical system view on recurrent neural networks. *arXiv:1902.09689 [cs, stat]*. arXiv: 1902.09689.
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204. doi: 10.1017/S0140525X12000477
- Clark, I., and Dumas, G. (2016). The regulation of task performance: a trans-disciplinary review. *Front. Psychol.* 6, 1862. doi: 10.3389/fpsyg.2015.01862
- Conant, R. C., and Ross Ashby, W. (1970). Every good regulator of a system must be a model of that system †. *Int. J. Syst. Sci.* 1, 89–97. doi: 10.1080/00207727008920220
- Cross, E. S., and Ramsey, R. (2021). Mind meets machine: Towards a cognitive science of human-machine interactions. *Trends Cogn. Sci.* 25, 200–212. doi: 10.1016/j.tics.2020.11.009
- Dehaene, S., Kerszberg, M., and Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proc. Natl. Acad. Sci. U.S.A.* 95, 14529–14534. doi: 10.1073/pnas.95.24.14529
- Deng, E., Mutlu, B., and Mataric, M. (2019). Embodiment in socially interactive robots. *Foundat. Trends Rob.* 7, 251–356. doi: 10.1561/23000000056
- Dennis, M., Jaques, N., Vinitzky, E., Bayen, A., Russell, S., Critch, A., et al. (2021). Emergent complexity and zero-shot transfer via unsupervised environment design. *arXiv:2012.02096 [cs]*. arXiv: 2012.02096.
- Djalovski, A., Dumas, G., Kinreich, S., and Feldman, R. (2021). Human attachments shape interbrain synchrony toward efficient performance of social goals. *Neuroimage* 226, 117600. doi: 10.1016/j.neuroimage.2020.117600
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., et al. (2020). An image is worth 16x16 words: transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dumas, G., Chavez, M., Nadel, J., and Martinerie, J. (2012). Anatomical connectivity influences both intra- and inter-brain synchronizations. *PLoS ONE* 7, e36414. doi: 10.1371/journal.pone.0036414
- Dumas, G., de Guzman, G. C., Tognoli, E., and Kelso, J. S. (2014a). The human dynamic clamp as a paradigm for social interaction. *Proc. Natl. Acad. Sci. U.S.A.* 111, E3726–E3734. doi: 10.1073/pnas.1407486111
- Dumas, G., Laroche, J., and Lehmann, A. (2014b). Your body, my body, our coupling moves our bodies. *Front. Hum. Neurosci.* 8, 1004. doi: 10.3389/fnhum.2014.01004
- Dumas, G., Nadel, J., Soussignan, R., Martinerie, J., and Garnero, L. (2010). Inter-brain synchronization during social interaction. *PLoS ONE* 5, e12166. doi: 10.1371/journal.pone.0012166

- Flynn, E., and Whiten, A. (2012). Experimental “microcultures” in young children: identifying biographic, cognitive, and social predictors of information transmission: Identifying predictors of information transmission. *Child Dev.* 83, 911–925. doi: 10.1111/j.1467-8624.2012.01747.x
- Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *J. Physiol. Paris* 100, 70–87. doi: 10.1016/j.jphysparis.2006.10.001
- Gariépy, J.-F., Watson, K., Du, E., Xie, D., Erb, J., Amasino, D., et al. (2014). Social learning in humans and other animals. *Front. Neurosci.* 8, 58 doi: 10.3389/fnins.2014.00058
- Goyal, A., Didolkar, A., Lamb, A., Badola, K., Ke, N. R., Rahaman, N., et al. (2021). Coordination among neural modules through a shared global workspace. *arXiv:2103.01197 [cs, stat]*. arXiv: 2103.01197.
- Graziano, M. (2019). Attributing awareness to others: the attention schema theory and its relationship to behavioural prediction. *J. Consciousness Stud.* 26, 17–37. Available online at: <https://grazianolab.princeton.edu/publications/attribution-awareness-others-attention-schema-theory-and-its-relationship>
- Graziano, M. S. A. (2017). The attention schema theory: a foundation for engineering artificial consciousness. *Front. Rob. AI* 4, 60. doi: 10.3389/frobt.2017.00060
- Graziano, M. S. A., and Kastner, S. (2011). Human consciousness and its relationship to social neuroscience: a novel hypothesis. *Cogn. Neurosci.* 2, 98–113. doi: 10.1080/17588928.2011.565121
- Graziano, M. S. A., and Webb, T. W. (2015). The attention schema theory: a mechanistic account of subjective awareness. *Front. Psychol.* 6, 500. doi: 10.3389/fpsyg.2015.00500
- Gruber, T., Bazhydai, M., Sievers, C., Clément, F., and Dukes, D. (2021). The abc of social learning: affect, behavior, and cognition. *Psychol Rev.* doi: 10.1037/rev0000311. [Epub ahead of print].
- Gupta, A., Savarese, S., Ganguli, S., and Fei-Fei, L. (2021). Embodied intelligence via learning and evolution. *Nat. Commun.* 12, 5721. doi: 10.1038/s41467-021-25874-z
- Ha, D., and Tang, Y. (2021). Collective intelligence for deep learning: a survey of recent developments. *arXiv:2111.14377 [cs]*. arXiv: 2111.14377.
- Hassabis, D., Kumaran, D., Summerfield, C., and Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron* 95, 245–258. doi: 10.1016/j.neuron.2017.06.011
- Heggli, O. A., Cabral, J., Konvalinka, I., Vuust, P., and Kringselbach, M. L. (2019). A kuramoto model of self-other integration across interpersonal synchronization strategies. *PLoS Comput. Biol.* 15, e1007422. doi: 10.1371/journal.pcbi.1007422
- Henschel, A., Hortensius, R., and Cross, E. S. (2020). Social cognition in the age of human-robot interaction. *Trends Neurosci.* 43, 373–384. doi: 10.1016/j.tins.2020.03.013
- Hernandez-Leal, P., Kartal, B., and Taylor, M. E. (2019). A survey and critique of multiagent deep reinforcement learning. *Auton. Agents Multi Agent Syst.* 33, 750–797. doi: 10.1007/s10458-019-09421-1
- Heyes, C. (2012). What’s social about social learning? *J. Compar. Psychol.* 126, 193–202. doi: 10.1037/a0025180
- Heyes, C. (2018). *Cognitive Gadgets*. Cambridge, MA: Harvard University Press.
- Heyes, C. M. (1994). Social learning in animals: categories and mechanisms. *Biol. Rev. Camb. Philos. Soc.* 69, 207–231. doi: 10.1111/j.1469-185X.1994.tb01506.x
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press. Google-Books-ID: z7gVDAAAQBAJ.
- Holzinger, D., Dall, M., Sanduvete-Chaves, S., Saldaña, D., Chacón-Moscoso, S., and Fellingner, J. (2020). The impact of family environment on language development of children with cochlear implants: a systematic review and meta-analysis. *Ear. Hear.* 41, 1077–1091. doi: 10.1097/AUD.0000000000000852
- Hoppitt, W., and Laland, K. N. (2008). Chapter 3 social processes influencing learning in animals: a review of the evidence. *Adv. Study Behav.* 38, 105–165. doi: 10.1016/S0065-3454(08)00003-X
- Husbands, P., Shim, Y., Garvie, M., Dewar, A., Domcsek, N., Graham, P., et al. (2021). Recent advances in evolutionary and bio-inspired adaptive robotics: exploiting embodied dynamics. *Appl. Intell.* 51, 6467–6496. doi: 10.1007/s10489-021-02275-9
- Jaques, N. (2019). *Social and Affective Machine Learning* (Ph.D. thesis). Massachusetts Institute of Technology, Cambridge, MA, United States. Available online at: <https://dspace.mit.edu/handle/1721.1/129901>
- Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P. A., Strouse, D. J., et al. (2019). Social influence as intrinsic motivation for multi-agent deep reinforcement learning. *arXiv:1810.08647 [cs, stat]*. arXiv: 1810.08647.
- Kendal, R., Hopper, L. M., Whiten, A., Brosnan, S. F., Lambeth, S. P., Schapiro, S. J., et al. (2015). Chimpanzees copy dominant and knowledgeable individuals: implications for cultural diversity. *Evolut. Hum. Behav.* 36, 65–72. doi: 10.1016/j.evolhumbehav.2014.09.002
- Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., and Jones, P. L. (2018). Social learning strategies: bridge-building between fields. *Trends Cogn. Sci.* 22, 651–665. doi: 10.1016/j.tics.2018.04.003
- Kerzel, M., Strahl, E., Magg, S., Navarro-Guerrero, N., Heinrich, S., and Wermter, S. (2017). “Nico-neuro-inspired companion: a developmental humanoid robot platform for multimodal interaction,” in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Lisbon: IEEE), 113–120.
- Klein, E. D., and Zentall, T. R. (2003). Imitation and affordance learning by pigeons (*columba livia*). *J. Comp. Psychol.* 117, 414–419. doi: 10.1037/0735-7036.117.4.414
- Kose-Bagci, H., Ferrari, E., Dautenhahn, K., Syrdal, D. S., and Nehaniv, C. L. (2009). Effects of embodiment and gestures on social interaction in drumming games with a humanoid robot. *Adv. Rob.* 23, 1951–1996. doi: 10.1163/016918609X12518783330360
- Kovač, G., Portelas, R., Hofmann, K., and Oudeyer, P.-Y. (2021). Socialai: Benchmarking socio-cognitive abilities in deep reinforcement learning agents. *arXiv:2107.00956 [cs]*. arXiv: 2107.00956.
- Lee, D., Jaques, N., Kew, C., Wu, J., Eck, D., Schuurmans, D., et al. (2021). Joint attention for multi-agent coordination and social learning. *arXiv:2104.07750 [cs]*. arXiv: 2104.07750.
- Lee, K. (2006). Are physically embodied social agents better than disembodied social agents?: the effects of physical embodiment, tactile interaction, and people’s loneliness in human-robot interaction. *Int. J. Hum. Comput. Stud.* 64, 962–973. doi: 10.1016/j.ijhcs.2006.05.002
- Lim, B., Arik, S. Ö., Loeff, N., and Pfister, T. (2021a). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* 37, 1748–1764. doi: 10.1016/j.ijforecast.2021.03.012
- Lim, S. H., Erichson, N. B., Hodgkinson, L., and Mahoney, M. W. (2021b). Noisy recurrent neural networks. *arXiv:2102.04877 [cs, math, stat]*. arXiv: 2102.04877.
- Macpherson, T., Churchland, A., Sejnowski, T., DiCarlo, J., Kamitani, Y., Takahashi, H., et al. (2021). Natural and artificial intelligence: a brief introduction to the interplay between ai and neuroscience research. *Neural Netw.* 144, 603–613. doi: 10.1016/j.neunet.2021.09.018
- Millidge, B. (2020). Deep active inference as variational policy gradients. *J. Math. Psychol.* 96, 102348. doi: 10.1016/j.jmp.2020.102348
- Mitchell, J. P. (2008). Activity in right temporo-parietal junction is not selective for theory-of-mind. *Cereb. Cortex* 18, 262–271. doi: 10.1093/cercor/bhm051
- Mittal, S., Lamb, A., Goyal, A., Voleti, V., Shanahan, M., Lajoie, G., et al. (2020). “Learning to combine top-down and bottom-up signals in recurrent neural networks with attention over modules,” in *International Conference on Machine Learning* (PMLR), 6972–6986.
- Ndousse, K. K., Eck, D., Levine, S., and Jaques, N. (2021). “Emergent social learning via multi-agent reinforcement learning,” in *Proceedings of the 38th International Conference on Machine Learning*, eds M. Meila and T. Zhang, Vol. 139 (PMLR), 7991–8004. Available online at: <https://proceedings.mlr.press/v139/ndousse21a.html>
- Newell, A., and Simon, H. A. (1976). Computer science as empirical inquiry: symbols and search. *Commun. ACM* 19, 113–126. doi: 10.1145/360018.360022
- Nguyen, T. T., Nguyen, N. D., and Nahavandi, S. (2020). Deep reinforcement learning for multi-agent systems: a review of challenges, solutions and applications. *IEEE Trans. Cybern.* 50, 3826–3839. doi: 10.1109/TCYB.2020.2977374
- Nielsen, M., Subiaul, F., Galef, B., Zentall, T., and Whiten, A. (2012). Social learning in humans and nonhuman animals: theoretical and empirical dissections. *J. Comp. Psychol.* 126, 109–113. doi: 10.1037/a0027758
- Pérez, A., Dumas, G., Karadag, M., and Duñabeitia, J. A. (2019). Differential brain-to-brain entrainment while speaking and listening in native and foreign languages. *Cortex* 111, 303–315. doi: 10.1016/j.cortex.2018.11.026

- Ramstead, M. J., Kirchoff, M. D., and Friston, K. J. (2020). A tale of two densities: active inference is enactive inference. *Adapt. Behav.* 28, 225–239. doi: 10.1177/1059712319862774
- Roy, N., Posner, I., Barfoot, T., Beaudoin, P., Bengio, Y., Bohg, J., et al. (2021). From machine learning to robotics: challenges and opportunities for embodied intelligence. *arXiv:2110.15245 [cs]*. arXiv: 2110.15245.
- Safron, A. (2020). An integrated world modeling theory (iwmt) of consciousness: Combining integrated information and global neuronal workspace theories with the free energy principle and active inference framework; toward solving the hard problem and characterizing agentic causation. *Front. Artif. Intell.* 3, 30. doi: 10.3389/frai.2020.00030
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., et al. (2013). Toward a second-person neuroscience. *Behav. Brain Sci.* 36, 393–414. doi: 10.1017/S0140525X12000660
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv:1707.06347 [cs]*. arXiv: 1707.06347.
- Shu, T., Ryoo, M. S., and Zhu, S.-C. (2016). Learning social affordance for human-robot interaction. *arXiv:1604.03692 [cs]*. arXiv: 1604.03692.
- Silver, D., Singh, S., Precup, D., and Sutton, R. S. (2021). Reward is enough. *Artif. Intell.* 299, 103535. doi: 10.1016/j.artint.2021.103535
- Smith, L., and Gasser, M. (2005). The development of embodied cognition: six lessons from babies. *Artif. Life* 11, 13–29. doi: 10.1162/1064546053278973
- Smithers, T. (2018). *Are Autonomous Agents Information Processing Systems?, 1st Edn.* Hillsdale, NJ: Routledge.
- Spelke, E. S., and Kinzler, K. D. (2007). Core knowledge. *Develop. Sci.* 10, 89–96. doi: 10.1111/j.1467-7687.2007.00569.x
- Thompson, E., and Varela, F. J. (2001). Radical embodiment: neural dynamics and consciousness. *Trends Cogn. Sci.* 5, 418–425. doi: 10.1016/S1364-6613(00)01750-2
- Thorpe, W. H. (1963). *Learning and Instinct in Animals*. London: Methuen.
- Tognoli, E., Dumas, G., and Kelso, J. A. S. (2018). A roadmap to computational social neuroscience. *Cogn. Neurodyn.* 12, 135–140. doi: 10.1007/s11571-017-9462-0
- Tomasello, M. (2020). The adaptive origins of uniquely human sociality. *Philos. Trans. R. Soc. B Biol. Sci.* 375, 20190493. doi: 10.1098/rstb.2019.0493
- Toth, K., Munson, J., N., Meltzoff, A., and Dawson, G. (2006). Early predictors of communication development in young children with autism spectrum disorder: joint attention, imitation, and toy play. *J. Autism. Dev. Disord.* 36, 993–1005. doi: 10.1007/s10803-006-0137-7
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behav. Brain Sci.* 21, 615–628. doi: 10.1017/S0140525X98001733
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in Neural Information Processing Systems* (Long Beach, CA), 5998–6008.
- Weigl, P. D., and Hanson, E. V. (1980). Observational learning and the feeding behavior of the red squirrel *tamiasciurus hudsonicus*: the ontogeny of optimization. *Ecology* 61, 213–218. doi: 10.2307/1935176
- Whitehurst, G. J. (1978). The contributions of social learning to language acquisition. *Contemp. Educ. Psychol.* 3, 2–10. doi: 10.1016/0361-476X(78)90002-4
- Whitehurst, G. J., Ironsmith, M., and Goldfein, M. (1974). Selective imitation of the passive construction through modeling. *J. Exp. Child Psychol.* 17, 288–302. doi: 10.1016/0022-0965(74)90073-3
- Whitehurst, G. J., and Vasta, R. (1975). Is language acquired through imitation? *J. Psycholinguist. Res.* 4, 37–59. doi: 10.1007/BF01066989
- Whiten, A. (2005). The second inheritance system of chimpanzees and humans. *Nature* 437, 52–55. doi: 10.1038/nature04023
- Whiten, A. (2021). The burgeoning reach of animal culture. *Science* 372, eabe6514. doi: 10.1126/science.abe6514
- Wilterson, A. I., and Graziano, M. S. A. (2021). The attention schema theory in a neural network agent: controlling visuospatial attention using a descriptive model of attention. *Proc. Natl. Acad. Sci. U.S.A.* 118, e2102421118. doi: 10.1073/pnas.2102421118
- Wong, A., Bäck, T., Kononova, A. V., and Plaat, A. (2021). Multiagent deep reinforcement learning: Challenges and directions towards human-like approaches. *arXiv:2106.15691 [cs]*. arXiv: 2106.15691.
- Wood, L. A., Kendal, R. L., and Flynn, E. G. (2013). Copy me or copy you? the effect of prior experience on social learning. *Cognition* 127, 203–213. doi: 10.1016/j.cognition.2013.01.002
- Worden, B. D., and Papaj, D. R. (2005). Flower choice copying in bumblebees. *Biol. Lett.* 1, 504–507. doi: 10.1098/rsbl.2005.0368
- Yuan, S., Zhao, H., Zhao, S., Leng, J., Liang, Y., Wang, X., et al. (2022). A roadmap for big model. *arXiv [Preprint]*. arXiv: 2203.14101. Available online at: <http://arxiv.org/abs/2203.14101>
- Zhang, M., Dumas, G., Kelso, J. S., and Tognoli, E. (2016). Enhanced emotional responses during social coordination with a virtual partner. *Int. J. Psychophysiol.* 104, 33–43. doi: 10.1016/j.ijpsycho.2016.04.001

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Bolotta and Dumas. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.