The Morphospace of Consciousness

**Xerxes D. Arsiwalla¹, Clément Moulin-Frier¹, Ivan Herreros¹,
Martí Sánchez-Fibla¹, Paul Verschure^{1,2}**

¹Synthetic Perceptive Emotive and Cognitive Systems (SPECS) Lab, Center of Autonomous Systems and Neurorobotics, Universitat Pompeu Fabra, Barcelona, Spain.

²Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain.

E-mail: x.d.arsiwalla@gmail.com

Abstract. Given recent proposals to synthesize consciousness, how many forms of conscious machines can one distinguish and on what grounds? Based on current clinical scales of consciousness, that measure cognitive awareness and wakefulness, we take a perspective on how contemporary artificially intelligent machines and synthetically engineered life forms would measure on these scales. To do so, we argue that awareness and wakefulness can be associated to computational and autonomous complexity respectively. Then, building on insights from cognitive robotics, we ask what function consciousness serves, and interpret it as an evolutionary game-theoretic strategy. We make the case for a third type of complexity necessary for describing consciousness, namely, social complexity. Having identified these complexity types, allows us to represent both, biological and synthetic systems in a common morphospace. This suggests an embodiment-based taxonomy of consciousness. In particular, we distinguish four forms of consciousness, based on embodiment: biological, synthetic, group (resulting from group interactions) and simulated consciousness (embodied by virtual agents within a simulated reality). Such a taxonomy is useful for studying comparative signatures of consciousness across domains, in order to highlight design principles necessary to engineer conscious machines. This is particularly relevant in the light of recent developments at the crossroads of neuroscience, biomedical engineering, artificial intelligence and biomimetics.

Keywords: Consciousness, Artificial Intelligence, Synthetic Biology, Cognitive Robotics, Complex Systems.

1. Introduction

Can one build a taxonomy of consciousness based on evidence from clinical neuroscience, synthetic biology, artificial intelligence (AI) and cognitive robotics? This article explores current biologically motivated metrics of consciousness and through that takes a perspective on how contemporary AI and synthetic systems measure up on these scales. In what follows, we refer to phenomenological consciousness, which can be

described in epistemically objective terms even though the problem of consciousness itself may well have an ontologically subjective element. Drawing from what is known about the phenomenology of consciousness in biological systems, we build an analogous argument for artificial, collective and simulated systems. For example, in clinical diagnosis of disorders of consciousness, two widely used scales are patient awareness and wakefulness (also referred to as arousal), both of which can be assessed using neurophysiological recordings [47], [45]. We will use these scales to construct a morphospace of consciousness. The idea of morphospaces has been discussed before in the context of complex networks and synthetic biology [12], [61]. A morphospace commits one to embodiment. In the context of consciousness, embodiment can be both, physical and virtual. Moreover, a morphospace serves as a useful tool to gain insights on design principles and evolutionary constraints, when looking across a large class of systems (or species) that display complex variations in traits. For the problem of consciousness, we construct this morphospace based on three distinct complexity types. These considerations suggest an embodiment-based taxonomy of consciousness [7].

For practical reasons, many experimental paradigms testing consciousness are designed for humans or higher-order primates (see [14], [43], [76] for an overview of the field). In this article, we argue that metrics commonly associated to biological consciousness can also be meaningfully used for conceptualizing behaviors of synthetic and artificially intelligent systems. This is insightful not only for understanding parallels between biological and potential synthetic consciousness, but more importantly for unearthing design principles necessary for building biomimetic technology that could potentially acquire consciousness. As evidenced by several historical precedents, bio-inspired design thinking has been at the core of some of the greatest scientific breakthroughs. For instance, early attempts at aviation in the 19th century were inspired by studying flight mechanics in birds and insects (the term aviation itself is derived from the Latin "avis" for "bird"). In fact, biological flight mechanisms are so sophisticated that their biomimetic implementations are still being actively studied within the field of soft robotics [55]. However, it so happened that rather than coming around to mimicking nature exactly, humanity learnt the basic laws of aerodynamics based on observations from nature and looked for other embodiments of those principles. This in fact, led to the invention of the modern aircraft by the Wright brothers in 1903, leading to a completely new way to build machines that fly than those that exactly mimic nature. Another paradigm-changing example of bio-inspired thinking leading to modern-day technological innovation can be seen in artificial neural networks, which dates back to the 1930s with the first model of neural networks by Nicolas Rashevsky [64], followed by the seminal work of Walter Pitts and Warren McCulloch in 1943 [54] and Frank Rosenblatt's perceptron in 1958 [66]. The field began as a modest attempt to understand cognition and brain function. Even though artificial neural networks did not quite solve the problem of how the brain works, they led to the discovery of brain-inspired computing technologies such as deep learning systems and powerful technologies for computational intelligence such as IBM's Watson. These machines process massive

volumes of data and are built for intensive computational tasks that the brain is not even designed for. In that spirit, the next frontier is understanding the governing principles of biological consciousness and its various embodiments, which could potentially lead to the growth of next-generation sentient technologies. Recent work in this direction can be found in [70].

Metrics of consciousness are also the right tools to quantitatively study how human intelligence differs from current machine intelligence. Once again, it is instructive to take a historical perspective on human intelligence as laid out by one of the founders of AI, Allen Newell in 1994 in his seminal work, "Unified Theories of Cognition" [59]. Newell proposed the following thirteen criteria necessary for building human-level cognitive architectures:

- Behave flexibly as a function of the environment
- Exhibit adaptive (rational, goal-oriented) behavior
- Operate in real-time
- Operate in a rich, complex, detailed environment (that is, perceive an immense amount of changing detail, use vast amounts of knowledge, and control a motor system of many degrees of freedom)
- Use symbols and abstractions
- Use language, both natural and artificial
- Learn from the environment and from experience
- Acquire capabilities through development
- Operate autonomously, but within a social community
- Be self-aware and have a sense of self
- Be realizable as a neural system
- Be constructible by an embryological growth process
- Arise through evolution

Current AI architectures still do not meet all these criteria. On the other hand, though Newell did not discuss consciousness back then, the above criteria are very relevant in the light of current research on neural mechanisms of consciousness [43]. While Newell's criteria list signatures that are the consequence of human intelligence, for consciousness it is more useful to have a list of functional criteria that results in the process of consciousness. In this article, we shall discuss prospective functional criteria for consciousness.

2. Biological Consciousness: Insights from Clinical Neuroscience

We begin this discussion reviewing clinical scales used for assessing consciousness in patients with disorders of consciousness. In subsequent sections, we generalize complexity measures pertinent to these biological scales and discuss how current synthetic systems measure up on these.

2.1. Clinical Consciousness and its Disorders

In patients with disorders of consciousness ranging from coma, locked-in syndrome to those in vegetative states, levels of consciousness are assessed through a battery of behavioral tests as well as physiological recordings. Cognitive awareness in patients is assessed by testing several cognitive functionalities using behavioral and neurophysiological (fMRI or EEG) protocols [47]. Assessments of wakefulness/arousal in patients are based on metabolic markers (in cases where reporting is not possible) such as glucose uptake in the brain, captured using PET scans. More generally, in [47] and [45], awareness and wakefulness have been proposed as a two dimensional operational definition of clinical consciousness, shown in figure 1 below. While awareness concerns higher and lower-order cognitive functions enabling complex behavior; wakefulness results from biochemical homeostatic mechanisms regulating survival drives and is clinically measured in terms of glucose metabolism in the brain. In fact, in all known organic life forms, biochemical arousal is a necessary precursor supporting the hardware necessary for cognition. In turn, evolution has shaped cognition in such a way so as to support the organism’s basic survival (using wakefulness/arousal) as well as higher-order drives (using awareness) associated to cooperation and competition in a multi-agent environment [80]. Awareness and wakefulness thus taken together, form the clinical markers of consciousness.

This clinical definition of consciousness enables a practical classification of closely associated states/disorders of consciousness into clusters on a bivariate scale with awareness and wakefulness on orthogonal axes. Under healthy conditions, these two levels are almost linearly correlated, as in conscious wakefulness (high arousal and high awareness) or in deep sleep (low arousal and low awareness). In pathological states, wakefulness without awareness can be observed in the vegetative state [47], while transiently reduced awareness is observed following seizures [17]. Patients in the minimally conscious state show intermittent and limited non-reflexive and purposeful behavior [29], [28], whereas patients with hemispatial neglect display reduced awareness of stimuli contralateral to the side where brain damage has occurred [62].

The question is how can one generalize wakefulness/arousal and awareness for non-biological systems in order to obtain homologous scales of consciousness that can be mapped to artificial systems? As noted above, wakefulness/arousal results from autonomous homeostatic mechanisms necessary for the self-preservation of an organism’s germ line in a given environment. In other words, arousal results from self-sustaining life processes necessary for basic survival, whereas awareness refers to functionalities pertaining to estimating or predicting states of the world and optimizing the agent’s own actions with respect to those states. If biological consciousness as we know it, is a synergy between metabolic and cognitive processes, the question one can ask is how should this insight be extended to conceive a functional notion of consciousness in synthetic systems? One way of doing so might be generalizing wakefulness/arousal to scales of autonomous functioning and awareness to scales of computational or

Figure 1. Clinical scales of consciousness. A clustering of disorders of consciousness in humans represented on scales of awareness and wakefulness. Adapted from [45]. In neurophysiological recordings, signatures of awareness have been found in cortico-thalamic activity, whereas wakefulness corresponds to activity in the brainstem and associated systems [47], [45]. Abbreviated legends: VS/UWS (vegetative state/unresponsive wakefulness state) [46]; MCS(+/-) (minimally conscious state plus/minus), EMCS (emergence from minimally conscious state) [20].

informational processes.

2.2. Measures of Consciousness

Specific measures of autonomy and computation/information processing have been discussed in psychometric [83] respectively neurophysiological studies [85]. However, applying these measures to artificial systems and comparing those values to biological systems is not always so straightforward (due to completely different processing substrates). Nonetheless, these measures offer a first step in this direction. For example, [83] introduce an "Index of Autonomous Functioning", tested on healthy human subjects (via psychometric questionnaires). This index aims to assess psychological ownership, interest-taking and susceptibility to external controls. This is similar to the concept of volition (or agency), introduced in the cognitive neurosciences [32], which seeks to determine the neural correlates of self-regulation, referring to actions regulated by internal drives rather than exclusively driven by external contingencies.

Attempts to quantify awareness have appeared in [24], discussed in the context of a unified psychological theory of self-functioning. However, in consciousness research, a measure of awareness that has gained a lot of traction is integrated information [78]

(often denoted as Φ). This is an information-theoretic complexity measure. It was first introduced in neuroscience as a measure applicable to neural networks. Based on mutual information, Φ has been touted as a correlate of consciousness [78]. Integrated information is loosely defined as the quantity of information generated by a network as a whole, due to its causal dynamical interactions, over and above the information generated independently by the disjoint sum of its parts. As a complexity measure, Φ seeks to operationalize the intuition that complexity arises from simultaneous integration and differentiation of the network's structure and dynamics, thus enabling the emergence of the system's collective states. The interplay between integration and differentiation generates information that is highly diversified yet integrated, creating patterns of high complexity. Following initial proposals [75], [77], [78], several approaches have been developed to compute integrated information [3], [8], [9], [10], [13], [15], [16], [31], [60], [63], [74], [84]. Notably the work of [10] is of particular significance in the context of this discussion as it develops large-scale network computations of integrated information, applied to the human brain's connectome data. The human connectome data consists of structural connectivity of white matter fiber tracts in the cerebral cortex, extracted using diffusion spectrum imaging and tractography [33], [38] (see [4], [11], [5] for neurodynamical models used on this network). Compared to a randomly re-wired network, it was seen that the particular topology of the human brain generates greater information complexity for all allowed couplings associated to the network's attractor states, as well as to its non-stationary dynamical states [10]. However, the formulation of Φ is not specific to biological systems and can equally well be applied to artificial dynamical systems and serves as a measure of their computational or information processing complexity (which we interpret as cognitive complexity or awareness in biological agents).

3. Synthetic Consciousness: Insights from Synthetic Biology and Artificial Intelligence

Generalizing arousal and awareness for artificial life and artificially intelligent systems is useful in view of the fact that there have recently been some remarkable advances in seemingly unrelated fields of artificial intelligence and synthetic life forms. A prominent example of the former is the AlphaGo system [69]. Important examples of the latter include the synthesis of artificial DNA with six nucleotide bases (the naturally occurring adenine (A), cytosine (C), guanine (G) and thymine (T) plus a new synthetic pair), which was engineered into the cell of the bacterium *Escherichia Coli* and successfully passed on from one generation to the next [53]; synthetic protocells capable of replicating themselves [44]; and the synthesis of a fully functioning artificial genome implanted into the cell of the bacterium *Mycoplasma genitalium* that converts it into a novel bacterium species [41].

Of course the clinical criteria we have defined above for consciousness, based on awareness and arousal, is not even remotely satisfied by any of these systems as

they either have some limited form of intelligence or life but not yet both. However, AlphaGo's feat in beating the top human Go champion was noteworthy for a couple of reasons. Unlike Chess, possible combinations in Go run into the millions and when played using a timer any brute-force algorithm trying to scan the entire search space would simply run out of computational capacity or time. Hence, an efficient pattern recognition algorithm was crucial to the development of AlphaGo, where using deep reinforcement learning the system was trained on a large number of games after which it was made to play itself over and over again (this aspect of playing itself is akin to training via social interactions as described later on) while reinforcing successful sequence of plays through the weights of its deep neural networks [69]. Most remarkably, it played counterintuitive moves that shocked the best human players and the sole game of the series that Lee Sedol, the human champion won out of five, itself was only possible after he himself adopted a brilliant counterintuitive strategy. Thus, AlphaGo demonstrates a form of domain-specific intelligence. In contrast, biological awareness spans across domains. Moreover, AlphaGo is not equipped with any form of arousal mechanisms coupled to its computational capabilities. The same can be said for other state-of-the-art AI systems including deep convolutional neural networks, or deep recurrent networks. Both these latter architectures were inspired from Hubel and Weisel's groundbreaking work on the coding properties of the visual system, which led to the realization of a hierarchical processing architecture [39]. Today deep convolutional networks are widely used for image classification [23] and recurrent neural networks for speech recognition [68], among countless other applications. The current interest of deep learning has been anticipated in computational neuroscience using objective functions from which physiologically plausible perceptual hierarchies can be learnt [87]. Recent developments have advanced this by virtue of larger data sets and more computational power. For example, there have been attempts to build biologically-plausible models of learning in the visual cortex using recurrent neural networks [51]. In summary, deep architectures have made remarkable progress in domain-specific AI.

However, asking whether a machine can be conscious in exactly the same way that a human is, is similar to asking whether a submarine can swim. It just does it differently. If the goal of a system is to learn and solve complex tasks close to human performance or better, current machines are already doing that in specific domains. However, these machines are still far from learning and solving problems in generic domains and in ways that would couple its problem-solving capabilities to its autonomous survival drives. On the other hand, neither have any of the synthetic life systems discussed above been used to build architectures with complex computing or cognitive capabilities. Nevertheless, this does suggest that a future synthesis between artificial life forms and AI could be evaluated using homologous scales of consciousness to the ones currently used for biological life forms. This form of synthetic consciousness, if based on a life form with different survival drives/mechanisms and non-human forms of intelligence or computation, would also likely lead to non-human behavioral outcomes.

These phenomenological considerations suggest at least two generic types

of complexities to label states of consciousness, those associated to computational/informational capabilities and those referring to autonomous functioning. In the following section, we argue for a third complexity type, necessary to build the morphospace of consciousness, namely, social complexity.

4. The Function of Consciousness: Insights from Cognitive Robotics

Based on insights from cognitive robotics, this section takes a functional perspective on consciousness [80], [35], [79], [6], [57] eventually interpreting it as a game-theoretic strategy. In [79], it was suggested that rather than being the problem itself, consciousness might in fact be a solution to the problem of autonomous goal-oriented action with intentionality, when agents are faced with a multi-agent social environment. The latter was formulated as the *H5W problem*.

4.1. The H5W Problem

What does an agent operating in a social world need to do in order to optimize its fitness? It needs to perceive the world, to act and, through time, to understand the consequences of its actions so it can start to reason about its goals and how to achieve them. This requires building a representation of the world grounded on the agent's own sensorimotor history and use that to reason and act. It will witness a scene of agents, including itself, and objects interacting in various manners, times and places. This comprises the six fundamental problems that the agent is faced with, together referred to as the H5W problem [79]:

- In order to act in the physical world an agent needs to determine a behavioral procedure to achieve a goal state; that is, it has to answer the HOW of action.

In turn this requires the agent to:

- Define the motivation for action in terms of its needs, drives and goals; that is, the WHY of action.
- Determine knowledge of objects it needs to act upon and their affordances in the world, pertaining to the above goals; that is, the WHAT of action.
- Determine the location of these objects, the spatial configuration of the task domain and the location of the self; that is, the WHERE of action.
- Determine the sequencing and timing of action relative to dynamics of the world and self; that is, the WHEN of action.
- Estimate hidden mental states of other agents when action requires cooperation or competition; that is, the WHO of action.

While the first four of the above questions suffices for generating simple goal-oriented behaviors, the last of the Ws (the WHO) is of particular significance as it involves intentionality, in the sense of estimating the future course of action of other agents based on their social behaviors and psychological states. However, because mental states of other agents that are predictive of their actions are hidden, they can at best

be inferred from incomplete sensory data such as location, posture, vocalizations, social salience, etc. As a result the acting agent faces the challenge to univocally assess, in a deluge of sensory data those exteroceptive and interoceptive states that are relevant to ongoing and future action and therefore has to deal with the ensuing credit assignment problem in order to optimize its own actions. Furthermore, this results in a reciprocity of behavioral dynamics, where the agent is now acting on a social and dynamical world that is in turn acting upon itself. It was proposed in [79] that consciousness is associated to the ability of an agent to maintain a transient and autonomous memory of the virtualized agent-environment interaction, that captures the hidden states of the external world, in particular, the intentional states of other agents and the norms that they implicitly convey through their actions.

4.2. Social Game Theory

Hence, the function of consciousness is to enable an acting agent to solve its H5W problem while being engaged in social cooperation and competition with other agents, who are trying to solve their own H5W problem in a world with limited resources. This leads our discussion precisely within the setting of social game theory. In a scenario with only a small number of other agents, a given agent might use statistical learning approaches to learn and classify behaviors of the few others agents in that game. For example, multiple robots interacting to learn naming conventions of perceptual aspects of the world [71]. Here the multi-agent interaction has to be embodied so that one agent can interpret which specific perceptual aspect the other agent is referring to (by pointing at objects) [72]. Another example is the emergence of signaling languages in sender-receiver games based on replicator dynamics described by David Lewis in 1969 in his seminal work, *Convention* [49], [50]. However, in both these examples, strategies that evolutionarily succeed when only few players are involved, are no longer optimal in the event of an explosion in the number of players [37]. Likewise in a social environment comprising a large number of agents trying to solve the H5W problem, machine learning strategies for reward and punishment valuations may soon become computationally unfeasible for an agent's processing capacities and memory storage. Therefore, for a large population to sustain itself in an evolutionary game involving complex forms of cooperation and competition would require strategies other than simple machine learning algorithms. One such strategy involves modeling and inferring intentional states of itself and that of other agents. Emotion-driven flight or fight responses depend on such intentional inferences and so do higher-order psychological drives. The mechanisms of consciousness enable such strategies, whereas, contemporary AI systems such as AlphaGo do not possess such capabilities.

In summary, interpreting consciousness as a game-theoretic strategy highlights the role of complex social behaviors inevitable for survival in a multi-agent world. From an evolutionary standpoint, social behaviors result from generations of cooperation-competition games, with natural selection filtering out unfavorable

strategies. Presumably, winning strategies were eventually encoded as anatomical mechanisms, such as emotional responses. The complexity of these behaviors depends on the ability of an agent to make complex social inferences. This suggests a third dimension in the morphospace of consciousness (shown in figure 3), namely, social complexity, which serves as a measure of an agent’s social intelligence.

5. Morphospace of Consciousness

As evident from our discussions above, consciousness research draws insights from a variety of disciplines such as clinical neuroscience, synthetic biology, artificial intelligence and cognitive robotics. Taken together, this suggests at least three complexity types that can be associated to consciousness: autonomous complexity, computational complexity and social complexity. As a generic definition of a system’s complexity $\mathcal{C}$, we define

$$\mathcal{C} = \mathcal{I}_{\text{substrate}} - \sum_{\{\text{parts}\}} \mathcal{I}_{\text{part}} \quad (1)$$

which is a measure of information generated by the dynamics of a system as a whole ($\mathcal{I}_{\text{substrate}}$) minus the sum of that generated by its parts. While this is similar in spirit to integrated information discussed in an earlier section, it is generically defined for specifying substrate-specific complexity. This provides a general framework that includes the possibility of several different types of complexity, among which, $\mathcal{C}_{\text{Autonomous}}$, $\mathcal{C}_{\text{Computational}}$ and $\mathcal{C}_{\text{Social}}$ will be relevant for labelling states of consciousness.

Autonomous complexity $\mathcal{C}_{\text{Autonomous}}$ measures the complexity of autonomous actions. In eukaryotes, autonomous action refers to arousal mechanisms resulting from coordinated nervous system activity; in prokaryotes, autonomous action refers to reactive behaviors such as chemotaxis, stress responses to temperature, toxins, mechanical damage, etc., all of these resulting from coordinated cellular signaling processes; in robotics, autonomous action refers to homeostatic mechanisms driving reactive behaviors. Therefore, autonomous complexity is the information generated by the collective dynamics of the complex system driving autonomous actions, over the information generated by a (hypothetical) uncoordinated copy of this system. On the other hand, computational complexity $\mathcal{C}_{\text{Computational}}$ refers to the ability of an agent to integrate information over space and time across computational or cognitive tasks. In complex biological systems, this complexity is typically associated to neural processes, in artificial computational systems, it refers to microprocessor signaling. The distinction between $\mathcal{C}_{\text{Computational}}$ and $\mathcal{C}_{\text{Autonomous}}$ is specified by the tasks that they refer to, rather than the specific substrate. Finally, social complexity $\mathcal{C}_{\text{Social}}$ refers to the information generated by the population as a whole, during the course of social interactions, over the information generated additively by individual agents of the population. Unlike $\mathcal{C}_{\text{Autonomous}}$ or $\mathcal{C}_{\text{Computational}}$, $\mathcal{C}_{\text{Social}}$ is not assigned to an individual, but rather to a specific population (its own species) with which the individual has been interacting. Nonetheless, as discussed above, these interactions are believed to have contributed to the consciousness of an individual on an evolutionary time-scale, by way of social

games. Note that $\mathcal{C}_{\text{Social}}$ as defined here, does not refer to group consciousness (we shall discuss that in the following section). Table 1 summarizes the three complexity types and specifies the corresponding substrates along with the associated emergent property.

	$\mathcal{C}_{\text{Autonomous}}$	$\mathcal{C}_{\text{Computational}}$	$\mathcal{C}_{\text{Social}}$
Substrate	Organism, nervous system, bots	Cognitive systems (brains, microprocessors)	Interacting population of agents
Parts	Sensors, actuators, signalling cascades	Neurons, transistors	Individual agents
Emergent Property	Self-regulated real-time behavior	Problem solving capabilities	Signaling systems, language, social norms, conventions, art, science, culture

Table 1. The three complexity types along with their respective substrates, components and emergent properties.

Figure 2. Schematic representation of autonomous, computational and social complexity. Each complexity measure is illustrated as a whole (the large circles) constituted of its parts (the inner circles), their interactions (the arrows) and the emerging properties resulting from these interactions (the inner space within the large circles, in light grey). Autonomous complexity (left) refers to the collective phenomena resulting from the interactions between typical components of reactive behavior such as sensors (illustrated by whiskers in the top inner circle), actuators (illustrated by a muscle in the bottom-left inner circle) and low-level sensorimotor coupling (illustrated by a spinal cord in the bottom-right inner circle). Computational complexity is associated to higher-level cognitive processes such as visual perception (top inner circle), planning (bottom-left inner circle) or decision making (bottom-right inner circle). Social complexity is associated to interactions between individuals of a population, such as a queen ant (top inner circle), a worker ant (bottom-left inner circle) and a soldier ant (bottom-right inner circle).

Using these definitions for the three types of complexities, we construct the following morphospace in figure 3.

Figure 3. Morphospace of consciousness. Autonomous, computational and social complexity constitute the three axes of the consciousness morphospace. Human consciousness is used as a reference in one corner of the space. Current AI implementations cluster together in the high computation, low autonomy and low social complexity regime, while multi-agent cognitive robotics cluster around low computational, but moderate autonomous and social complexities. Abbreviated legends: MADeep (multi-agent deep reinforcement system) [73]; TalkH (talking heads) [72]; DQN (deep Q-learning) [56]; DAC-X (distributed adaptive control) [52], CoBot (cockroach robot) [34], Kilobot (swarm robot) [67], Subsumption (mobile robot architecture) [19].

While figure 3 is only a first attempt at constructing the morphospace of prospective conscious systems, the precise coordinates of various systems within this morphospace are subject to change due to the rapid pace of new and developing technologies. However, one can justify the main clusters appearing in the figure as follows. The human is taken as the benchmark in this space. She/he can perform computational tasks across a variety of domains such as making logical inference, planning an optimal path in a complex environment or dealing with recursive problems and hence leads with respect to computational complexity due to these cross-domain capabilities. On the social axis, human social interactions have resulted in the emergence of language, music, art, culture, socio-political systems, etc. Other biological entities such as non-human primates [18], [82], bees and ants score lower on the social and computational axis than humans.

Current AI systems such as IBM Watson [36], AlphaGo [69], DQNs [56] and Siri [2] are powerful computing systems over a narrow set of domains, but in their current form they do not show general intelligence, that is, the capacity to independently interact with the world and successfully perform different tasks in different domains [48], or as proposed by Allen Newell, the capability with which anything can become a task [58]. These AI systems are still clustered high on the computational axis, but lower than humans (due to domain-specificity). Also they score low on autonomy and social complexity. Synthetic forms of life such as protocells show some levels of arousal, reacting to chemicals and stressors, but currently show minimal capabilities for computation or interactions with other agents [44]. Finally, interest in the field of collective robotics has led to the rise of machines where emergent macro-properties, e.g. coordination (KiloBot [67], Multi-Agent Deep Network [73]) or shared semantic conventions (Talking Heads [72]) self-organize out of multi-agent interactions. These systems are designed to display simple forms of navigation, object-detection, etc., while interacting with other agents performing the same task. However, they show lower social and autonomous complexity than most biological agents, and being embodied, they currently score lower on computational complexity than heavy-powered AI systems such as IBM Watson and AlphaGo. Notice also that a large region in the central zone of the morphospace in figure 3 is suspiciously vacant. A similar observation was made in [61] in the context of the morphospace of synthetic organs and organoids. In both cases, such an observation points towards new classes of future machines. In the following section, we discuss two possible manifestations of future conscious systems.

6. Other Embodiments of Consciousness

The three dimensional morphospace discussed above provides us with a framework to also identify other types of complex systems whose levels of computational, autonomous and social complexity might be sufficient to answer the H5W of consciousness? This suggests at least two other embodiments of future conscious systems (based on the same functional criteria as above).

6.1. Group Consciousness

In a sense biological consciousness, itself can be thought of as a collective phenomenon where individual cells making up an organism are themselves not considered to be conscious (with respect to the three complexity measures defining the morphospace), even though the organism as a whole is. But what happens when the system itself is not localized? We postulate group consciousness as an extension of the above idea to composite or distributed systems that display levels of computational, autonomous and social complexity that are sufficient to answer the H5W problem. Note that, as per this specification of group consciousness, the group itself is treated as one entity. Hence, social complexity now refers to the interactions of this group with other similar groups.

This bears some resemblance to the notion of collective intelligence, which is a widely studied phenomenon in complex systems ranging from ant colonies [25], to a swarm of robots (the Kilobot in [67] and the CoRobot in [34]), to social networks [30]. But these are generally not regarded as conscious systems. As a whole they are not considered to be life forms with survival drives that compete or cooperate with other similar agents. However, these considerations begin to get blurred at least during transient epochs when collective survival comes under threat. For example, when a bee colony comes under attack by hornets, collectively it demonstrates a prototypical survival drive, similar to lower-order organisms. Other examples of such behaviors have also been studied in the context of group interactions in humans, where social sensitivity, cooperation and diversity have been shown to correlate with the collective intelligence of the group [86]. Following this, the notion of collective intentionality has been discussed in [40]. More recently, [27] have applied integrated information Φ to group interactions, suggesting a new kind of group consciousness. While it is known that Φ in adapting agents increases with fitness [26], one can ask a similar question for an entire group: what processes (evolutionary games, learning, etc.) enable an increase in all three complexity types for an entire group such that it can solve the H5W problem while cooperating or competing with other groups?

6.2. Simulated Consciousness

Our discussions on complexity also suggest another type of consciousness, namely, simulated consciousness, wherein embodied virtual agents in a simulated reality interact with other virtual agents, while satisfying the complexity bounds that enable them to answer the H5W questions within the simulation. In this case, consciousness is strictly confined to the simulated environment. The agents cannot perceive or communicate with entities outside of the simulation but satisfy all the criteria we have discussed above within the simulation. How these embodied virtual agents could acquire consciousness is not yet known. Presumably by evolving across multiple generations of agents that adapt and learn to optimize fitness conditions. It is also not clear what precise traits or mechanisms would have to be coded into the simulation (as initializations or priors) in order to enable consciousness to evolve. The point here is simply that the same criteria that we have identified with consciousness in biological or synthetic agents in the physical world, could in principle be admitted by agents within a simulation and confined to their interactions within that simulation. This has parallels to the notion of "Machine Consciousness" discussed in [65], which proposes that neural processes leading to consciousness might be realizable as a machine simulation (it even goes further to claim that computer systems might someday be able to emulate consciousness). At the moment, these are all open challenges in AI and consciousness research. Examples of studies discussing embodied virtual agents can be found in the work of [22] and [21]. More recent implementations of embodied virtual agents have been using gaming technology, such as the Minecraft platform [1], [42].

7. Discussion

The objective of this article was to bring together diverse ideas; from neuroscience, AI, synthetic biology and robotics; that have recently been converging towards the science of consciousness. Following progress in these fields, we have attempted to generalize the applicability to current clinical scales of consciousness to synthetic systems. Combined with insights from social robotics, we have taken a functional perspective of consciousness, interpreting it as an evolutionary game-theoretic strategy. We make the case for at least three complexity types necessary for describing consciousness, namely, autonomous, computational and social complexity. These form the bases of a consciousness morphospace. Besides biological and synthetic consciousness, the above considerations also suggest other possible manifestations of consciousness, namely, group consciousness and simulated consciousness, each based on a distinct embodiment.

In our discussion, social complexity was important for constructing the morphospace. Social interactions play an important role in regulating many cognitive and adaptive behaviors in both, natural and artificial systems [80]. In [79] it has been suggested that complex social interactions may have evolutionary served as the trigger for consciousness. What is however not known is whether there are specific lower bounds on the scales of each of the stated complexity types, that an agent must cross in order to attain a given level of consciousness. Certainly, from developmental biology we know that both humans (and many higher-order animals) undergo extensive periods of cognitive and social learning from infancy to maturation. These phases of social and cognitive training are crucial for development of cognitive abilities leading to levels of consciousness attained by healthy adults.

Even though we may be far from understanding all the engineering principles required to build conscious machines, a complexity-based comparison between biological and artificial systems reveals interesting insights. For example, current AI systems using deep learning tend to cluster along the computational complexity axis of the morphospace, whereas synthetically engineered life forms group closer along the autonomous complexity axis. On the other hand, biologically conscious agents are distributed in regions of the morphospace corresponding to relatively high complexity along all the three axes (which suggests necessary, if not sufficient, conditions for consciousness). In terms of Newell's criteria, excluding those that refer exclusively to human-specific traits (language, symbolic reasoning), the remainder are completely satisfied by all agents located in the high complexity region (of all three axes) of the consciousness morphospace. In contrast, current AI or synthetic systems do not check-out on this list. Though in 1994 Newell was not explicitly referring to consciousness, it is remarkable to note how those ideas to formulate theories of cognition and intelligence seem to reconcile with current ideas of consciousness. One could summarize the crux of Newell's criteria as referring to agents displaying autonomous behaviors with cross-domain problem-solving capabilities, which can be decomposed to (at least) the three complexity classes discussed in this paper.

This perspective on consciousness opens several possibilities for future work. For instance, it may be interesting to further refine the morphospace described here. In particular, computational complexity itself may involve several sub-types involving learning, adaptation, acquiring sensorimotor representations, etc, all of which are relevant for cognitive robotics [81]. Another question arising out of our discussion is whether the emergence of consciousness in a multi-agent social environment can be identified as a Nash equilibrium of a cooperation-competition game. In a game where say two species attain consciousness, the population pay-offs in cooperation and competition between them are likely to reach one of possible equilibria due to the recursive nature of intentional inferences, where an agent attempts to infer the inferences of other agents about its own intentions. Multi-agent models might offer a viable approach to test this idea.

7.1. Societal and Ethical Considerations

No discussion on conscious machines is complete without the very important issue of ethics. Both, the societal impact and ethical considerations of any form of advanced machine, especially conscious machines, for obvious reasons, constitutes a very serious issue. For example, the impact of medical nanobots for removing tumors, attacking viruses or non-surgical organ reconstruction has the potential to change medicine forever. Or AI systems to clear pollutants from the atmosphere or the rivers are absolutely essential for some of the biggest problems that humanity faces. However, as discussed above, purely increasing the performance of a machine along the computational axis will not constitute consciousness as long as these capabilities are not accessible by the system to autonomously regulate or enhance its survival drives. On the other hand, whenever the latter is indeed made possible, issues of societal interactions of machines with humans and the ecosystem, becomes an imminent ethical responsibility. It becomes important to understand the kind of cooperation-competition dynamics that a futuristic human society will face. Early stages of designing such machines are probably the best times to regulate their future impact on society. This analogy might not be surprising to any parent that has a child. Hence, a serious effort towards understanding the evolution of complex social traits is crucial alongside engineering advances required for the development of these systems.

Acknowledgments

We would like to thank Ricard Solé for correspondence on related work ([70]). We also thank Riccardo Zucca and Sytse Wierenga for help with graphics. This work has been supported by the European Research Council's CDAC project: "The Role of Consciousness in Adaptive Behavior: A Combined Empirical, Computational and Robot based Approach" (ERC-2013- ADG 341196).

References

- [1] Aluru, K., Tellex, S., Oberlin, J., MacGlashan, J.: Minecraft as an experimental world for ai in robotics. In: AAAI Fall Symposium (2015)
- [2] Aron, J.: How innovative is apple’s new voice assistant, siri? *New Scientist* 212(2836), 24 (2011)
- [3] Arsiwalla, X.D., Verschure, P.F.M.J.: Integrated information for large complex networks. In: The 2013 International Joint Conference on Neural Networks (IJCNN). pp. 1–7 (Aug 2013)
- [4] Arsiwalla, X.D., Betella, A., Bueno, E.M., Omedas, P., Zucca, R., Verschure, P.F.: The dynamic connectome: A tool for large-scale 3d reconstruction of brain activity in real-time. In: ECMS. pp. 865–869 (2013)
- [5] Arsiwalla, X.D., Dalmazzo, D., Zucca, R., Betella, A., Brandi, S., Martinez, E., Omedas, P., Verschure, P.: Connectomics to semantomics: Addressing the brain’s big data challenge. *Procedia Computer Science* 53, 48–55 (2015)
- [6] Arsiwalla, X.D., Herreros, I., Moulin-Frier, C., Sanchez, M., Verschure, P.F.: Is Consciousness a Control Process?, pp. 233–238. IOS Press, Amsterdam (2016)
- [7] Arsiwalla, X.D., Herreros, I., Verschure, P.: On Three Categories of Conscious Machines, pp. 389–392. Springer International Publishing, Cham, Switzerland (2016)
- [8] Arsiwalla, X.D., Verschure, P.: Computing Information Integration in Brain Networks, pp. 136–146. Springer International Publishing, Cham, Switzerland (2016)
- [9] Arsiwalla, X.D., Verschure, P.F.M.J.: High Integrated Information in Complex Networks Near Criticality, pp. 184–191. Springer International Publishing, Cham, Switzerland (2016)
- [10] Arsiwalla, X.D., Verschure, P.F.: The global dynamical complexity of the human brain network. *Applied Network Science* 1(1), 16 (2016)
- [11] Arsiwalla, X.D., Zucca, R., Betella, A., Martinez, E., Dalmazzo, D., Omedas, P., Deco, G., Verschure, P.: Network dynamics with brainx3: A large-scale simulation of the human brain network with real-time interaction. *Frontiers in Neuroinformatics* 9(2) (2015)
- [12] Avena-Koenigsberger, A., Goñi, J., Solé, R., Sporns, O.: Network morphospace. *Journal of The Royal Society Interface* 12(103), 20140881 (2015)
- [13] Ay, N.: Information geometry on complexity and stochastic interaction. *Entropy* 17(4), 2432–2458 (2015)
- [14] Baars, B.J.: Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in brain research* 150, 45–53 (2005)
- [15] Balduzzi, D., Tononi, G.: Integrated information in discrete dynamical systems: motivation and theoretical framework. *PLoS Comput Biol* 4(6), e1000091 (2008)
- [16] Barrett, A.B., Seth, A.K.: Practical measures of integrated information for time-series data. *PLoS Comput Biol* 7(1), e1001052 (2011)
- [17] Blumenfeld, H.: Impaired consciousness in epilepsy. *The Lancet Neurology* 11(9), 814–826 (2012)
- [18] Borjon, J.I., Takahashi, D.Y., Cervantes, D.C., Ghazanfar, A.A.: Arousal dynamics drive vocal production in marmoset monkeys. *Journal of neurophysiology* 116(2), 753–764 (2016)
- [19] Brooks, R.: A Robust Layered Control System for a Mobile Robot. *IEEE Journal on Robotics and Automation* 2(1), 14–23 (1986)
- [20] Bruno, M.A., Vanhaudenhuyse, A., Thibaut, A., Moonen, G., Laureys, S.: From unresponsive wakefulness to minimally conscious plus and functional locked-in syndromes: recent advances in our understanding of disorders of consciousness. *Journal of neurology* 258(7), 1373–1384 (2011)
- [21] Burden, D.J.: Deploying embodied ai into virtual worlds. *Knowledge-Based Systems* 22(7), 540–544 (2009)
- [22] Cassell, J.: Embodied conversational agents. MIT press (2000)
- [23] Ciresan, D.C., Meier, U., Masci, J., Maria Gambardella, L., Schmidhuber, J.: Flexible, high performance convolutional neural networks for image classification. In: *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*. vol. 22, p. 1237. Barcelona, Spain (2011)
- [24] Deci, E.L., Ryan, R.M.: The” what” and” why” of goal pursuits: Human needs and the self-

- determination of behavior. *Psychological inquiry* 11(4), 227–268 (2000)
- [25] Dorigo, M., Birattari, M., et al.: Swarm intelligence. *Scholarpedia* 2(9), 1462 (2007)
- [26] Edlund, J.A., Chaumont, N., Hintze, A., Koch, C., Tononi, G., Adami, C.: Integrated information increases with fitness in the evolution of animats. *PLoS Comput Biol* 7(10), e1002236 (2011)
- [27] Engel, D., Malone, T.W.: Integrated information as a metric for group interaction: Analyzing human and computer groups using a technique developed to measure consciousness. *arXiv preprint arXiv:1702.02462* (2017)
- [28] Giacino, J.T.: The vegetative and minimally conscious states: consensus-based criteria for establishing diagnosis and prognosis. *NeuroRehabilitation* 19(4), 293–298 (2004)
- [29] Giacino, J.T., Ashwal, S., Childs, N., Cranford, R., Jennett, B., Katz, D.I., Kelly, J.P., Rosenberg, J.H., Whyte, J., Zafonte, R., et al.: The minimally conscious state definition and diagnostic criteria. *Neurology* 58(3), 349–353 (2002)
- [30] Goleman, D.: Social intelligence. Random house (2007)
- [31] Griffith, V.: A principled infotheoretic ϕ -like measure. *arXiv preprint arXiv:1401.0978* (2014)
- [32] Haggard, P.: Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience* 9(12), 934–946 (2008)
- [33] Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C.J., Wedeen, V.J., Sporns, O.: Mapping the Structural Core of Human Cerebral Cortex. *PLoS Biology* 6(7), 15 (2008)
- [34] Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., Said, I., Durier, V., Canonge, S., Amé, J.M., et al.: Social integration of robots into groups of cockroaches to control self-organized choices. *Science* 318(5853), 1155–1158 (2007)
- [35] Herreros, I., Arsiwalla, X., Verschure, P.: A forward model at purkinje cell synapses facilitates cerebellar anticipatory control. In: *Advances in Neural Information Processing Systems*. pp. 3828–3836 (2016)
- [36] High, R.: The era of cognitive systems: An inside look at ibm watson and how it works. IBM Corporation, Redbooks (2012)
- [37] Hofbauer, J., Huttegger, S.M.: Feasibility of communication in binary signaling games. *Journal of theoretical biology* 254(4), 843–849 (2008)
- [38] Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P.: Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences* 106(6), 2035–2040 (2009)
- [39] Hubel, D.H., Wiesel, T.N.: Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology* 160(1), 106–154 (1962)
- [40] Huebner, B.: *Macrocognition: A theory of distributed minds and collective intentionality*. Oxford University Press (2013)
- [41] Hutchison, C.A., Chuang, R.Y., Noskov, V.N., Assad-Garcia, N., Deerinck, T.J., Ellisman, M.H., Gill, J., Kannan, K., Karas, B.J., Ma, L., et al.: Design and synthesis of a minimal bacterial genome. *Science* 351(6280), aad6253 (2016)
- [42] Johnson, M., Hofmann, K., Hutton, T., Bignell, D.: The malmo platform for artificial intelligence experimentation. In: *International joint conference on artificial intelligence (IJCAI)*. p. 4246 (2016)
- [43] Koch, C., Massimini, M., Boly, M., Tononi, G.: Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience* 17(5), 307–321 (2016)
- [44] Kurihara, K., Okura, Y., Matsuo, M., Toyota, T., Suzuki, K., Sugawara, T.: A recursive vesicle-based model protocell with a primitive model cell cycle. *Nature communications* 6 (2015)
- [45] Laureys, S.: The neural correlate of (un) awareness: lessons from the vegetative state. *Trends in cognitive sciences* 9(12), 556–559 (2005)
- [46] Laureys, S., Celesia, G.G., Cohadon, F., Lavrijsen, J., León-Carrión, J., Sannita, W.G., Szabon, L., Schmutzhard, E., von Wild, K.R., Zeman, A., et al.: Unresponsive wakefulness syndrome: a new name for the vegetative state or apallic syndrome. *BMC medicine* 8(1), 68 (2010)
- [47] Laureys, S., Owen, A.M., Schiff, N.D.: Brain function in coma, vegetative state, and related

- disorders. *The Lancet Neurology* 3(9), 537–546 (2004)
- [48] Legg, S., Hutter, M., Others: A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and applications* 157, 17 (2007)
- [49] Lewis, D.: *Convention: a philosophical study* (1969)
- [50] Lewis, D.: *Convention: A philosophical study*. John Wiley & Sons (2008)
- [51] Liao, Q., Poggio, T.: Bridging the gaps between residual learning, recurrent neural networks and visual cortex. *arXiv preprint arXiv:1604.03640* (2016)
- [52] Maffei, G., Santos-Pata, D., Marcos, E., Sánchez-Fibla, M., Verschure, P.F.: An embodied biologically constrained model of foraging: from classical and operant conditioning to adaptive real-world behavior in *dac-x*. *Neural Networks* 72, 88–108 (2015)
- [53] Malyshev, D.A., Dhimi, K., Lavergne, T., Chen, T., Dai, N., Foster, J.M., Corrêa, I.R., Romesberg, F.E.: A semi-synthetic organism with an expanded genetic alphabet. *Nature* 509(7500), 385–388 (2014)
- [54] McCulloch, W.S., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5(4), 115–133 (1943)
- [55] Mischiati, M., Lin, H.T., Herold, P., Imler, E., Olberg, R., Leonardo, A.: Internal models direct dragonfly interception steering. *Nature* 517(7534), 333–338 (2015)
- [56] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. *Nature* 518(7540), 529–533 (2015)
- [57] Moulin-Frier, C., Puigbò, J.Y., Arsiwalla, X.D., Sanchez-Fibla, M., Verschure, P.F.: Embodied artificial intelligence through distributed adaptive control: An integrated framework. *arXiv preprint arXiv:1704.01407* (2017)
- [58] Newell, A.: You Can’t Play 20 Questions with Nature and Win: Projective Comments on the Papers of this Symposium. *Visual Information Processing* pp. 283–308 (1973)
- [59] Newell, A.: *Unified theories of cognition*. Harvard University Press (1994)
- [60] Oizumi, M., Albantakis, L., Tononi, G.: From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS Comput Biol* 10(5), e1003588 (2014)
- [61] Ollé-Vila, A., Duran-Nebreda, S., Conde-Pueyo, N., Montañez, R., Solé, R.: A morphospace for synthetic organs and organoids: the possible and the actual. *Integrative Biology* 8(4), 485–503 (2016)
- [62] Parton, A., Malhotra, P., Husain, M.: Hemispatial neglect. *Journal of Neurology, Neurosurgery & Psychiatry* 75(1), 13–21 (2004)
- [63] Petersen, K., Wilson, B.: Dynamical intricacy and average sample complexity. *arXiv preprint arXiv:1512.01143* (2015)
- [64] Rashevsky, N.: Outline of a physico-mathematical theory of excitation and inhibition. *Protoplasma* 20(1), 42–56 (1933)
- [65] Reggia, J.A.: The rise of machine consciousness: Studying consciousness with computational models. *Neural Networks* 44, 112–131 (2013)
- [66] Rosenblatt, F.: The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological review* 65(6), 386 (1958)
- [67] Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. *Science* 345(6198), 795–799 (2014)
- [68] Sak, H., Senior, A.W., Beaufays, F.: Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In: *Interspeech*. pp. 338–342 (2014)
- [69] Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.: Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587), 484–489 (2016)
- [70] Sole, R.: Rise of the humanbot. *arXiv preprint arXiv:1705.05935* (2017)
- [71] Steels, L.: Evolving grounded communication for robots. *Trends in cognitive sciences* 7(7), 308–312 (2003)

- [72] Steels, L., Hild, M.: Language grounding in robots. Springer Science & Business Media (2012)
- [73] Tampuu, A., Matiisen, T., Kodelja, D., Kuzovkin, I., Korjus, K., Aru, J., Aru, J., Vicente, R.: Multiagent cooperation and competition with deep reinforcement learning. *PloS one* 12(4), e0172395 (2017)
- [74] Tegmark, M.: Improved measures of integrated information. *arXiv preprint arXiv:1601.02626* (2016)
- [75] Tononi, G.: An information integration theory of consciousness. *BMC neuroscience* 5(1), 42 (2004)
- [76] Tononi, G., Boly, M., Massimini, M., Koch, C.: Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* 17(7), 450–461 (2016)
- [77] Tononi, G., Sporns, O.: Measuring information integration. *BMC neuroscience* 4(1), 31 (2003)
- [78] Tononi, G., Sporns, O., Edelman, G.M.: A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences* 91(11), 5033–5037 (1994)
- [79] Verschure, P.F.: Synthetic consciousness: the distributed adaptive control perspective. *Phil. Trans. R. Soc. B* 371(1701), 20150448 (2016)
- [80] Verschure, P.F., Pennartz, C.M., Pezzulo, G.: The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Phil. Trans. R. Soc. B* 369(1655), 20130483 (2014)
- [81] Verschure, P.F., Voegtlin, T., Douglas, R.J.: Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature* 425(6958), 620–624 (2003)
- [82] de Waal, F.B.: Apes know what others believe. *Science* 354(6308), 39–40 (2016)
- [83] Weinstein, N., Przybylski, A.K., Ryan, R.M.: The index of autonomous functioning: Development of a scale of human autonomy. *Journal of Research in Personality* 46(4), 397–413 (2012)
- [84] Wennekers, T., Ay, N.: Stochastic interaction in associative nets. *Neurocomputing* 65, 387–392 (2005)
- [85] Wibral, M., Vicente, R., Lindner, M.: *Transfer Entropy in Neuroscience*, pp. 3–36. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
- [86] Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N., Malone, T.W.: Evidence for a collective intelligence factor in the performance of human groups. *Science* 330(6004), 686–688 (2010)
- [87] Wyss, R., König, P., Verschure, P.F.J.: A model of the ventral visual system based on temporal stability and local memory. *PLoS Biol* 4(5), e120 (2006)