Exploration of the Functional Properties of Interaction: Computer Models and Pointers for Theory

Etienne B. Roesch • University of Reading, UK • e.b.roesch/at/reading.ac.uk
Matthew Spencer • University of Reading, UK • m.c.spencer/at/reading.ac.uk
Slawomir J. Nasuto • University of Reading, UK • s.j.nasuto/at/reading.ac.uk
Thomas Tanay • Goldsmiths, University of London, UK • thomas.tanay/at/gmail.com
J. Mark Bishop • Goldsmiths, University of London, UK • m.bishop/at/gold.ac.uk

> Context • Constructivist approaches to cognition have mostly been descriptive, and now face the challenge of specifying the mechanisms that may support the acquisition of knowledge. Departing from cognitivism, however, requires the development of a new functional framework that will support causal, powerful and goal-directed behavior in the context of the interaction between the organism and the environment. **> Problem** • The properties affecting the computational power of this interaction are, however, unclear, and may include partial information from the environment, exploration, distributed processing and aggregation of information, emergence of knowledge and directedness towards relevant information. **> Method** • We posit that one path towards such a framework may be grounded in these properties, supported by dynamical systems. To assess this hypothesis, we describe computational models inspired from swarm intelligence, which we use as a metaphor to explore the practical implications of the properties highlighted. **> Results** • Our results demonstrate that these properties may serve as the basis for complex operations, yielding the elaboration of knowledge and goal-directed behavior. **> Implications** • This work highlights aspects of interaction that we believe ought to be taken into account when characterizing the possible mechanisms underlying cognition. The scope of the models we describe cannot go beyond that of a metaphor, however, and future work, theoretical and experimental, is required for further insight into the functional role of interaction with the environment for the elaboration of complex behavior. **> Constructivist content** • Inspiration for this work stems from the constructivist impetus to account for knowledge acquisition based on interaction. **> Key words** • **Stochastic diffusion search, clustering, multi-agent system, cognitivism.**

Introduction

Constructivism and the cognitive science method

« 1 » Constructivist views stemmed from the difficulty of traditional accounts, behaviorism and cognitivism in particular, to explain how knowing arises as individuals develop (Piaget 1937; Foerster 1972; Smock & Glasersfeld 1974; Glasersfeld 1984). A central tenet in constructivist theories indeed lies in the notion that knowledge of the world, for an individual, is created from the interaction with the environment, rather than existing in an ontic reality, supposedly pre-existing or available for registration from the physical world, as implied by traditional accounts. In particular, radical constructivism describes knowing as the "result of self-organizing construction" (Glasersfeld 2005: 11) based on the individual's experience. In other words, the entirety of the individual is proposed to be engaged in the acquisition of knowledge, which is brought forth at many levels in the embodied and situated interaction of the individual with the environment (Roesch, Nasuto & Bishop 2012).

« 2 » Despite decades of development, however, constructivist perspectives have mostly remained applied (e.g., in education) or descriptive. They now face the challenge of specifying the mechanisms at play in the processes of both acquiring knowledge and combining it into new knowledge. Taking root in the computational theory of mind (computationalism), most proposals from cognitivism describe a framework of information-theoretical processes operating on informational content representing the world. These frameworks further detail the causal manipulations that these representations undergo in support of goal-directed behavior (Dretske 2003; Churchland 2005). To offer a viable alternative to cognitivism, constructivist theories must, therefore, grow in precision and attempt to answer the same fundamental questions, such as: What constitutes this knowledge that is constructed through interaction with the environment? Where is it constructed? How is it acted upon, leading to the construction of new knowledge?

« 3 » The present paper is an attempt at going beyond descriptive accounts by exploring some of the properties that are intrinsic to the interaction of biological organisms with the environment and that may support causal, powerful, goal-directed behavior. We introduce one family of mechanisms based on swarm intelligence, which we posit exhibit some of these properties. The mechanisms may be used as a metaphor to explore their practical implications with a view to accounting for goal-directed behavior in a constructivist and enactive way. We describe results obtained from three computer models implementing swarms of agents whose behavior is modeled after these properties. These three models grow in complexity, and our ambition is not to demonstrate biologically plausible mechanisms, but to explore the functional abilities of dynamical systems implementing fundamental properties of interaction. This approach is analogous to that of Hofstadter and colleagues, who explored some of the mechanisms underlying general intelligence using idealized computer models (Hofstadter et al. 1995).

Interactive foundations of biological organisms

« 4 » Biological organisms function owing to the interaction of deeply coupled dynamical systems, which support features of experience such as embodiment and embeddedness. These features are paramount to the organism's identity in its lived world, quite literally constituting the interface between the organism and the environment. Through this interface, the organism finds itself sensing changes in the environment as it moves and evolves in the world.

« 5 » Cognitivism, on the one hand, describes this interface in terms of modules specialized in particular aspects of this interaction. Perception, for instance, is often described as one such module, whose purpose is to sense information from the world and transmit it to higher levels of computation. The modules composing these levels are often depicted in terms of the functions they are believed to implement, in an information processing framework grounded in the manipulation of representations as internal models of the outside world (e.g., Fodor 1983; Fodor & Pylyshin 1988; Scheutz 2002; Dretske 2003). While this approach might have provided a fruitful theoretical handle on the computation that might be performed in the brain, we posit that this stance lacks momentum in accounting for the scalability of knowledge acquisition and the creation of new knowledge (Roesch, Nasuto & Bishop 2012).

« 6 » Constructivist theories, on the other hand, emphasize the functional role of the organism's adaptability to changes in the environment to explain how the organism settles in the environment in ways optimal for adaptation and survival. Sensorimotor contingency theory, for instance, accounts for perception and experience jointly (O'Regan & Noë 2001). For Kevin O'Regan and Alva Noë, perception "is a mode of exploration of the world that is mediated by knowledge" of sensorimotor contingencies (ibid: 940). This knowledge is created as the organism experiences the world, and its biological substrate becomes attuned to the laws governing the changes in the environment that it can expect to encounter. They emphasize that this knowledge is practical in nature, and does not relate to propositional dispositions that would be gathered and manipulated – as would be the case in cognitivism.

« 7 » In this theoretical perspective, as in other (more) enactive perspectives, perception is an integral part of the process from which knowledge of the world arises, yielding awareness of the environment and the ability to combine information to create new knowledge. More so, the boundary between perception and higher-level cognitive modules in fact disappears. Perception is a way of acting in the world, mediated by the body. This mediation occurs through the many interleaved systems that compose the body, thus providing a flexible way to respond to the environment. These considerations have important consequences for the definition of cognition, which can no longer be ascribed solely to the brain, but must include the many levels that compose the body in the world (Anderson, Richardson & Chemero 2012).

« 8 » Importantly, this realization also highlights the foundational character of the organism's interaction with the environment, which is no longer depicted as a mere by-product of the fact that the organism is situated within the environment, but quite literally constitutes the basis of more elaborate behavior. Exploration of the environment provides the organism with the ability to sense and become attuned to the laws governing change. Therefore it is critical to understand the properties of this interaction that may affect further elaboration, and that should be taken into account by constructivist theories.

« 9 » One assertion that is typically made by cognitivists, for instance, is that the environment is perceived in full by the cognizing system. A typical scenario would thus see the perception module capturing the instantaneous and static information present at the interface between the organism and the environment; this signal is transduced into a representation of the current state of the environment, in as much detail as possible, such that it can be manipulated and serve in higher-level modules. One extreme example may be the "fear module" thesis, which asserts that one area of the brain, the amygdala, assumes the function underlying the detection and reaction to objects from the environment that could be harmful to the organism (e.g., Öhman et al. 2007). The theory goes as far as hypothesizing low-level, hard-coded, rapid, automatic, pre-attentional and unconsciousness mechanisms in mammals' brains, which will extract and detect representations of snakes and spiders from instantaneous visual perceptions.

« 10 » There are many reasons to believe perception may not work that way, both at the anatomical and the functional levels. Alternative proposals point to the intrinsic limitations of cognition, related to resource consumption and energy conservation mainly, and explore the compensatory strategies that render perception seamless. These often yield surprising illusory effects when observed in particular conditions – see O'Regan & Noë (2001) for a review of such effects in the visual domain. This body of work emphasizes that biological organisms evolve in the environment in the light of partial and often noisy information. A critical aspect of interaction therefore concerns the exploration of the environment and the aggregation of information about the regularities that are present in the environment that may support further elaboration. This information is intrinsically multimodal, and processed in distributed ways. Consequently, unitary knowledge representing the entirety of the environment does not exist at the level of the organism per se, but emerges though communication between the systems that compose the organism. Additionally, to support goal-directed behavior, exploration of the environment must follow adequate rules that will permit the selection of information and direct resources towards relevant parts of the environment.

« 11 » In the next section, with a view to taking a pragmatic perspective on the properties discussed here, we introduce a number of idealized computer models, which we use as a metaphor for interaction in the context of goal-directed tasks.

Computer models of interaction based on Stochastic Diffusion Search

« 12 » To use computer models in cognitive science, as in many other academic fields, is to commit oneself to a circumscribed experimental situation, which makes theoretical assumptions explicit and forces the rules governing the processes under scrutiny to be spelled out (Korsten et al. 2011). These self-imposed constraints, leading to the characterization of a situation that only resembles that of the real world, may seem unable to provide much insight, especially for a theoretical domain as broad as that envisaged by constructivism. The scope of constructivist theory, spanning developmental learning and educational constructivism, indeed encompasses a very broad group of processes that contribute to the organism constructing itself through experience, which may be difficult to capture in a single model. Additionally, the current state of description of these theoretical processes does not currently lend itself well to algorithmic investigation and modelers must attempt to be creative whilst keeping within the boundary of the domain. Invariably, computer models of constructivist processes will thus seem incomplete.

« 13 » Nevertheless, we consider it a healthy exercise to subject certain aspects of constructivism to the scrutiny of computer models. In doing so, we hope to create pointers to aspects of the theory that turned out to be important for a working computer model, and which may benefit from further theoretical characterization. In the present work, we followed a number of assumptions taken from constructivist theory, which we mapped to selected modeling techniques as best as possible.

« 14 » Firstly, we posit that the organism is composed of a structured collection of systems that communicate and influence each other, that evolve over time and that are critical to the rise of experience. Secondly, we fare away from the notion that the organism holds a model (representation) of the environment that is full, used for further computation and constantly kept up-to-date. Instead, we assume that every system that composes the organism accesses partial and noisy information from the environment, asynchronously and more or less directly, that is used to perform a very particular and simple task. Finally, we ascribe the overall computational power to the higher level of the organism, which adapts itself to suit the environment, interact with it and modify it. Importantly, we further assume the organism to be endowed with a given goal, the success of which we use as measure of this computational power. It is tempting for constructivist theorists to want the organism to define its own goals but, for simplicity, we relate goals to the physiological drive that an organism must pursue unconditionally. Broadly speaking, we assume this goal to be the result of processes akin to autopoiesis (Maturana & Varela 1980) or representational normativity (Bickhard 2006), for instance. The processes leading to the emergence of this goal, although of great interest indeed, do not enter the scope of our models.

« 15 » Given this set of assumptions, we developed computer models based on multi-agent systems to represent the processes that an idealized organism might go through when interacting with the environment. Our goal in doing so was to assess the computational ground that such a set of assumptions would lay in the context of multi-agent systems. We based this work on *Stochastic Diffusion Search* (SDS), which is a family of algorithms that originates from the field of swarm intelligence, exploring the collective intelligence of populations of agents (Bishop 1989). Inspiration for these algorithms comes from nature, including flocks, colonies or societies of individuals from all species (Berdahl et al. 2013), as well as bacteria or cells of the immune system (Banchereau & Steinman 1998). These biological systems have indeed been shown to be able to solve very complex tasks that require the integration, computation and exchange of information amongst the entire population of agents (Afek et al. 2011). Such tasks broadly include search in a problem space and optimization of behavior in regard to a particular set of constraints, requirements and the very simple abilities intrinsic to each agent; for instance, determining the best location for settling a nest (desert ants), the tuning of sensory networks in the nervous system (Afek et al. 2011) or the modulation of the immune response to a pathogen (Banchereau & Steinman 1998).

« 16 » In effect, SDS algorithms are probabilistic, agent-based global search and optimization algorithms that rely on the exploration of the search space and diffusion of information between the agents that constitute the colony or the organism. Information is sparse, noisy and incomplete, and agents explore the search space concurrently, exercising two modes of functioning sequentially: each agent formulates a hypothesis about the environment, which is evaluated against a minimal source of information (test phase); the result of these evaluations is then shared with the other agents (diffusion phase). This two-phase process carries on over time until a termination criteria is met. The individual agents' computational powers are minimal and success is determined at the level of the population, which has been shown to converge towards a global optimum (Nasuto & Bishop 1999). At the level of the population, a final step may include the consolidation of behavior.

« 17 » A simple example may be as follows. Consider the task of determining whether a string of symbols, the target, is contained within a (much) longer string of symbols, the environment or search space. For simplicity, assume a perfect match to the target exists in the environment. A population of agents is devised such that each agent knows only a small part of the target in the form of a series of adjacent symbols. The process evolves over time. At each time step, each agent is assigned a random position in the environment and makes the hypothesis that the small part of the target it knows can be found in the vicinity of this position. This hypothesis is evaluated against the environment, and every agent whose hypothesis has been verified takes active status, and inactive if it has not. In the diffusion phase, inactive agents take the hypotheses of active agents. It can be shown that, as test and diffusion phases alternate, the population of agents self-organizes and converges towards the most promising hypothesis, the location of the target in the environment.

« 18 » We posit that SDS may be a fruitful metaphor for some of the mechanisms at play in the acquisition of knowledge and the creation of new knowledge, in the context of radical constructivist theories. In particular, we show that SDS algorithms can be used to perform complex computation, and explore two paradigms that are representative of the mechanical processes put forth in more traditional views: summation and sorting of series of digits. These paradigms have been used to illustrate the computational power of foundational perspectives, such as the Turing Machine (Turing 1936; Sieg & Byrnes 1996), and have formed the basis of theoretical conceptions based on the mechanistic manipulation of symbols. In this paper, we show that, even though agents' abilities are minimal, the population of agents as a whole, interacting with the environment, is able to organize and consolidate a solution to such complex problems without referring to a representation of the external world.

Digit summation

« 19 » The simplest example of goal-directed behavior may be that of digit summation, which is characteristic of the mechanical operations that inspired traditional computer science (Turing 1936; Sieg & Byrnes 1996). In this first implementation, we endow a population of agents with limited knowledge about the regularities that may be found in the environment, and the task of the swarm is simply to play out this knowledge against the environment. We use this framework as point of departure for more complex examples, and to ease the reader's comprehension of subsequent models.

Implementation

« 20 » *Initialization*: A search space is created at random, in the form of a series of digits, which is then shuffled.

« 21 » *Exploration*: A population of agents is created, each holding a hypothesis whereby a pair of digits defined at random will be found at a particular location in the search space. At each step, every agent explores the hypothesis it is holding against the search space. If the hypothesis is validated, agents take active status, and inactive if it is not.

« 22 » *Diffusion*: Hypotheses for all active agents are gathered, and the search space is modified accordingly. When the sum of two digits leads to a number greater than 10, a particular flag signals a carry-over.

Results

« 23 » In this very simple example, a population of agents is tasked to sum a search space composed of 10 digits (Figure 1). At each time step, a fraction of the population sees its hypotheses validated, and the search reduces until it contains only one number, the result of the sum.

« 24 » The scope of this model may seem remote from constructivist considerations, compared to our other models. With this example, we hope to demonstrate the mechanics at play when a collection of systems acts out a particular strategy. Of particular importance is the kind of information available at each level of the model. At the level of the agents, information from the environment is partial and distributed. From the interaction of the agents with the environment arises a particular experience that grounds the population of agents in the environment, thereby permitting a meaningful behavior.

Sorting: Communication

« 25 » In this second example, we extend our swarm of agents by giving them the ability to communicate between each other. Members of the swarm possesses the same kind of limited knowledge about the environment and, with a view to sorting the series of digits that composes the environment, the swarm alternates between exploration and diffusion of information, before collating the intelligence gathered and manipulating the environment.

Step 0: [6, 3, 5, 4, 7, 3, 5, 3, 9, 5]
Step 1: [6, 3, 5, 11, 3, 8, 9, 5]
Step 2: [6, 8, 14, 8, 9, 5]
Step 3: [6, 8, 22, 14]
Step 4: [6, 8, 36]
Step 5: [6, 44]
Step 8: [50]

Figure 1: Representation of the evolution of the search space over time

Implementation

« 26 » *Initialization*: A search space is created at random, in the form of a series of digits, which is then shuffled.

« 27 » *Exploration*: A population of agents is created, each holding a hypothesis whereby a pair of ordered digits will be found at a particular location in the search space. At each step, every agent explores the hypothesis it is holding against the search space. If the hypothesis is validated, it takes active status, and inactive if it is not.

« 28 » *Diffusion*: The diffusion phase follows an active recruitment strategy, whereby active agents select an inactive agent at random, which is then provided with their hypothesis. To avoid testing a location that has already been tested, these newly created hypotheses are modified to spread the exploration of the search space to either side of the search space, around the initial location. Locations that yielded a validated hypothesis are swapped. Termination occurs when the search space reaches the expected state, or when the population of agents has been inactive for a particular lapse of time.

Results

« 29 » Communication leads to bursts of activation in the swarm of agents, and these progressively support the sorting of the search space. Bursts of activity indicate that clustering of agents is called upon when unsorted locations of the search space have been identified by the swarm (Figure 2). This mode of organization, akin to the synchronization of resources, is critical for the organism to construct an experience of the environment and direct behavior (De Meyer, Bishop & Nasuto 2000).

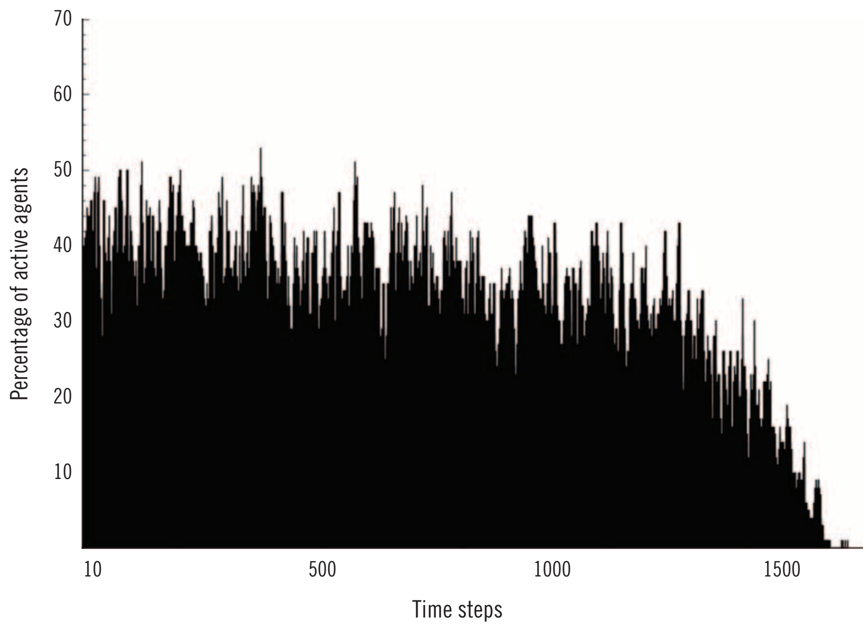

Figure 2: Effect of the communication between 100 agents sorting a search space composed of 200 digits.

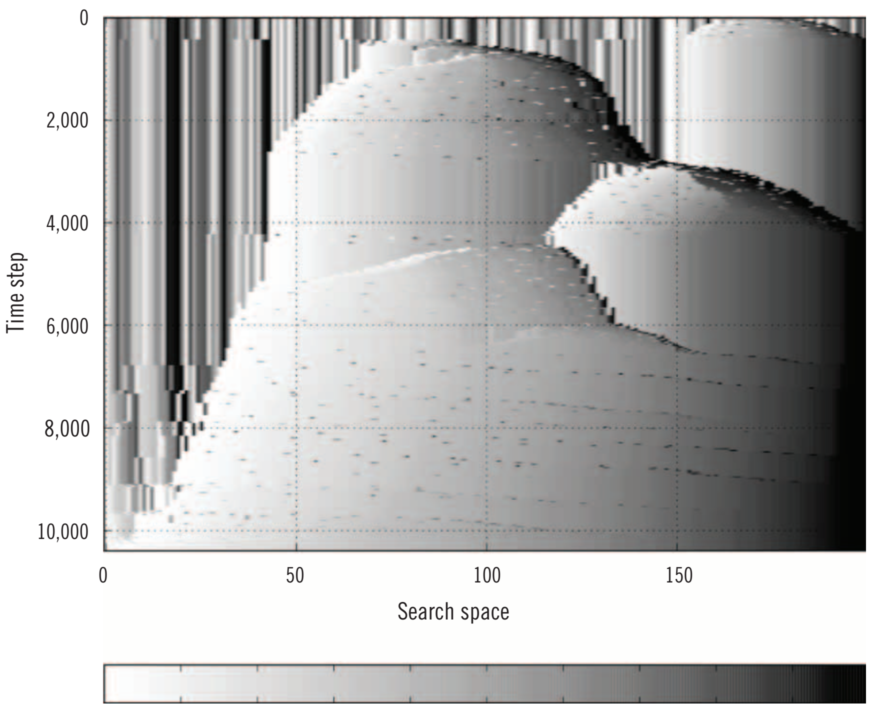

Figure 3: Representation of the evolution of the search space over time. In this simulation, it took c.a. 10 000 time steps for 10 agents to sort a search space containing 200 digits. The state of the search space is represented on the x-axis, where each pixel is color-coded, from white to black, to indicate the ordering of the digits. A representation of the ordered search space is shown below the graph. Communication between the clusters of agents can be identified when one digit is moved away to its rightful position.

Sorting: Communication and Exploration

« 30 » In this final example, we enhance our agents with the ability to explore their surroundings. This active perception generally provides the swarm with increased acuity over the environment in which they are evolving. Additionally, we introduce a pairwise recruitment mechanism with the aim of fostering collaboration between agents.

Implementation

« 31 » *Initialization*: A search space is created at random, in the form of a series of digits, which is then shuffled.

« 32 » *Exploration*: Agents are given random indices (hypotheses), which they use to explore the search space in either direction, carrying on for as long as they find sorted digits or until they hit a boundary – thereby constructing partial knowledge of their immediate surroundings. Upon hitting two consecutive digits that are unsorted, the agent takes active status, and inactive status if it does not, i.e., it hits a boundary. Communication ensues when all agents are assigned a status.

« 33 » *Diffusion*: Active agents that have hit a violation perform a swap at that location and enter into recruitment mode. Every active agent contacts another agent at random, one that is active and not otherwise already engaged in communication; the pair compare the length of their explorations; and the winner assigns a new hypothesis to the other agent, within the patch of sorted digits it knows, before continuing to explore its initial patch of digits (variants as to what hypothesis they themselves follow within this patch make a big difference in terms of performance).

Results

« 34 » A population of agents endowed with communication and exploration abilities is shown to be able to sort arrays of any size (Figure 3). As groups of agents cluster in the search space in response to perceived regularities in the environment (Figure 4), several clusters of agents may co-exist at the same time. Our algorithm introduces a bias for exploration, and agents seek to expand the cluster they belong to, resulting in information transfer between clusters so that digits are transported between clusters, across the search space.

Discussion and limitations

« 35 » We described multi-agent models as a metaphor for the processes at play when an organism interacts with its environment. We sought to challenge constructivist considerations by taking an operational stance and creating empirical situations that would test their computational power. In particular, we based our models on a number of assumptions. We assumed that (a) an organism is composed of distributed systems, which (b) perceive, react to and exchange partial information, thereby (c) elaborating an experience of the environment over time, and supporting a goal-directed behavior.

« 36 » We modeled this idealized organism as a multi-agent system based on Stochastic Diffusion Search that explored a given environment in a structured manner. Computational power ensued at the level of the organism from the consolidation of the intelligence gathered about the environment. In three models, we manipulated the abilities of agents to communicate with each other and interact with the environment. In light of the empirical situations we created, our results show that a certain level of computational power is attainable given this set of constraints. Against this result, it is important to raise some of the limitations to our work.

« 37 » The issue of scalability is paramount for constructivist theory, which aims to account for the construction of the individual as a result of their experience of the world. In this context, the question of whether our models do in fact generalize to a full-blown constructivist perspective may rightfully be asked. An answer to this question would concern both the computational capabilities of the models, and the scope of the theoretical interpretations that one would formulate from the results. SDS is a very powerful family of algorithms based on distributed processing with a view to exploring and identifying specified patterns in a noisy search space. Mathematical analyses of the behavior exhibited by this dynamical system prove convergence towards statistical equilibrium, and further analyses indicate good robustness and scalability with problem size (Nasuto & Bishop 1999). A more theoretical criterion of scalability, however, may relate to intentionality and the emergence of meaning in the experiencing individual, which goes beyond the capabilities of our models. One of the central hypotheses found in constructivism indeed asserts that an individual gradually constructs a sense of self-identity in their lived world. Future work may tackle this particular issue, which is non-trivial for modelers.

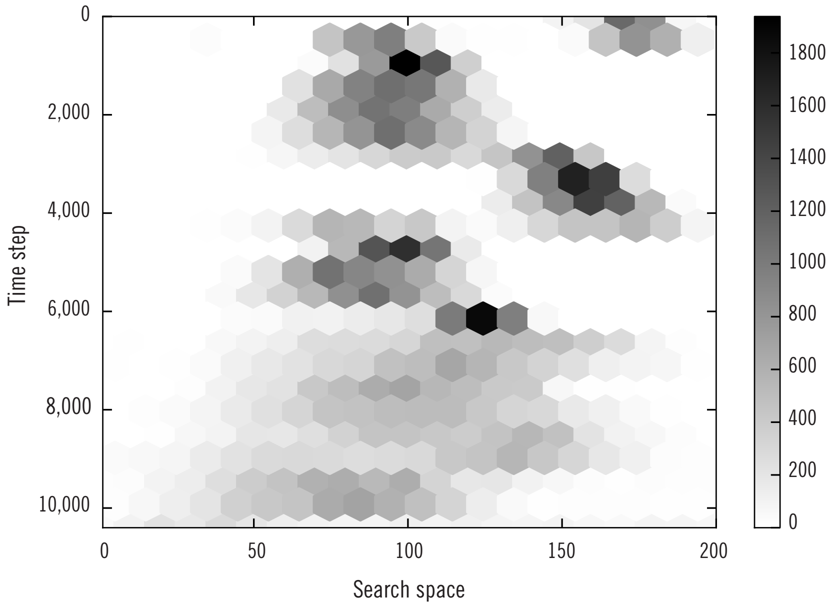

Figure 4: Representation of the evolution of the clustering of agents in the search space, corresponding to the simulation depicted in Figure 3. The total number of agents present at any point in time in the search space is color-coded. Each hexagon aggregates the activity of agents over about 500 time steps.

« 38 » Another important aspect to constructivist theory, and the scalability of our models, concerns the goal-directedness that emerges within an individual as it develops. As noted earlier, different proposals have been made to describe the processes as well as the physical pressure at play in the shaping of these goals (e.g., Maturana & Varela 1980; Bickhard 2006). Because our initial interest was in the mechanisms underlying interaction between the organism and the environment, and the assessment of the ensuing computational power in light of particular examples from more mechanistic manipulation of symbols (Turing 1936; Sieg & Byrnes 1996), we chose to endow our models with tailored tasks that resembled this previous work and that they could achieve more or less successfully. This choice certainly constrained what can be inferred from the results of our models.

« 39 » Finally, our work did not aim to propose biologically plausible mechanisms that would support knowledge acquisition. Rather, we used computer models to identify aspects of interaction between an organism and the environment that may support complex behavior, and act as pointers to theorists. Future work may focus on specific mechanisms and compare models against data collected in vivo.

Conclusion

« 40 » With a view to contributing to the theoretical debate surrounding the proposal of alternatives to cognitivism, we reviewed properties of the interaction between the organism and the environment that, we posit, are paramount for the acquisition of information from the environment, and for causal, powerful and goal-directed behavior. We identified several of such properties,

ETIENNE B. ROESCH

is a Lecturer in Cognitive Science at the University of Reading, UK. His research interests include general brain functioning, emotion, perception, attention, consciousness, and cognitive science methods, i.e., theoretical, (biologically plausible) computational, and experimental work (ERG, coupled EEG-fMRI, and psychophysics). His current efforts involve the development of a novel framework for investigating the role of interaction with the environment in giving rise to our experience of the world.

MATTHEW C. SPENCER

completed his PhD in 2012 on data-driven and modelling techniques for the study of evolving complex networks describing neuronal functional connectivity dynamics, in both cellular cultures and whole-brain EEG. Applications of his research pertain to animats and brain-computer interfaces as well as to the classification of complex system dynamics in general. He is currently working as a post-doctoral research assistant in the Brain Embodiment Lab, investigating the computational power of bio-realistic machine intelligence paradigms and continuing his work on stochastic network models of brain connectivity.

SLAWOMIR J. NASUTO

received an MSc in Mathematics from the University of Marie Curie-Sklodowska in Lublin, Poland in 1993 and a PhD in Cybernetics from the University of Reading in 1999. He has been a member of the academic staff in the School of Systems Engineering at Reading since 2000. His research interests span computational neuroscience and neuroanatomy, analysis of signals generated by the nervous system including EEG, and single neuron and multivariate spike trains or EMG and their applications for brain-computer interfaces and animats (robots controlled by neural cultures). He also has research interests in Stochastic Diffusion Search, an algorithm belonging to a family of swarm intelligence methods, and has been involved in applications of such techniques in aerospace problems.

THOMAS TANAY

is currently enrolled in a MRes/PhD programme at UCL in the Centre for Mathematics and Physics in the Life Sciences and Experimental Biology (CoMPLEX). In December 2012, he graduated from Goldsmiths, University of London with a MSc in Cognitive Computing, before taking a short term position of research assistant under the supervision of Prof. Mark Bishop. During this period, his work has essentially consisted in adapting the swarm intelligence method called Stochastic Diffusion Search (SDS) to solve complex forward planning problems such as game-playing.

MARK BISHOP

is Professor of Cognitive Computing at Goldsmiths, University of London and Chair (Elect) of the AISB (the UK society for the study of artificial intelligence and the simulation of behaviour). His research is in the field of cognitive computing: its theory – where his interests centre on the foundations of stochastic diffusion processes; its application – where he has worked on industrial problems in autonomous robotics, neural networks, and colour; and its philosophical foundations – where he has developed a novel argument against the possibility of machine consciousness. Together with John Preston, Mark has co-edited a critique of John Searle's arguments against machine intelligence, *Views into the Chinese Room* (2002).

such as partial information, exploration, distributed processing and aggregation of information, emergence of knowledge and directedness towards relevant information, that may support such operation, and used computer models to evaluate these properties in situ. In particular, we used Stochastic Diffusion Search algorithms as a metaphor for exploring dynamical aspects of the interaction between a population of agents and the environment. We showed that a computer model created with the explicit intention of embodying these properties achieves tasks such as digit summation and sorting successfully. This work contrasts sharply with the more descriptive accounts from constructivism in that we were forced to make practical moves to fit both engineering and theoretical constraints. The dynamical systems we described are grounded on both an internal structure and specific mechanisms steering an interaction with the environment. We hope that future theoretical work will include detailed description of these mechanisms.

Acknowledgements

The authors wish to thank the editors of this special issue on computational constructivism, Alexander Riegler and Tom Ziemke, as well as two anonymous reviewers for their comments on previous drafts of this manuscript. This work was partly funded by an award to JMB and SN by the Sir John Templeton Foundation.

Received: 15 June 2013
Accepted: 24 September 2013

Open Peer Commentaries

on Etienne Roesch et al.'s "Exploration of the Functional Properties of Interaction"

Systems Sciences and the Limitations of Computer Models of Constructivist Processes

Manfred Füllsack
ISIS/University of Graz, Austria
manfred.fuellsack/at/uni-graz.at

> Upshot • Why computer models of constructivist processes can enhance constructivist matters even though the models will always "seem incomplete."

« 1 » Etienne Roesch and his colleagues suggest a convincing model on stochastic diffusion search as an example for how computer models can simulate constructivist processes. When asked to write this commentary, I was prompted by a somehow incidental remark on an issue that I was recently musing about myself. The remark is about "computer models of constructivist processes" that will always "seem incomplete" (§12).

« 2 » Constructivism, in its philosophical sense, is an all-embracing theory. One might say it is always radical. Delimitations of its scope (as could be conjectured from sub-theories such as *social* constructivism) do not withstand its epistemological rigor. There is no transition from constructed to less constructed, or from real to less real. It was, in particular, Ernst von Glasersfeld (e.g., Glasersfeld 1995) who argued in favor of a "radical" (i.e., thorough) application of constructivism: if you go for constructivism you have to apply it all the way, however bumpy (since counter-intuitive) the road may become.

« 3 » This awkward feature associates constructivism with the concept of systems, a concept that hardly ever aims at less than the whole. Though at times deployed just to delimit objects for investigation, the term by definition[1] obliges one to consider interrelations, and interrelations tend to reach beyond any limits. They are never purely internal. The concept hence pushes its own boundaries. It forces a system's context to be considered and thereby calls to mind that any analytical containment called "system" will always be a system *in systems*. No final delimitation will stand. So, as with constructivism, if you decide to go for systems you have to go all the way.

« 4 » Analysis, however, would be lost in a limitless space. Where should scientific observation begin and where should it end? Hence, like other scientific objects, systems too must be – at least temporarily – limited in order to function as analytical means. But if they are delimited anyway, you may object, what then, apart from their name, distinguishes systems from conventional objects as they are delimited in sciences all along?

« 5 » The difference that makes a difference in this case – and this is the topic of this issue of *Constructivist Foundation* – is the digital computer. Systems sciences deploy digital machines in their own right as analytical tools. These tools, as we all know, allow for an unprecedented synopsis of interrelations even though they are limited in the sense of having to start at an externally

1 | The term "system "originates from the Greek term "σύστεμα" for "being put together."

predefined level of order, at an *ontological* base, so to speak, a base that constructivism rejects. In other words, computers enable investigation into a complexity of relations and interactions that was not conceivable before the digital age, and they do so in spite of their technical limitations. They allow "crawling the micro-causal web," as Mark Bedau (1997) called it, a web of illustrative, albeit incomplete, interactions of lower-level components that by themselves do not show the properties of the higher-level phenomena to which their aggregations emerge.

« 6 » With this, the computer allows for the verification of insights, of which previously the possibility could be suspected at best but could never be proven. The computational power of contemporary digital machines indeed provides ground for "a new kind of science" (Wolfram 2002). And in this kind of science – and I consider the paper of Roesch et al. to be part of it – it does not matter that the machine so far is not (yet) able to simulate emergences that span much more than two levels of order, or as Roesch et al. (§12) state it, that "computer models of constructivist processes will thus seem incomplete." As in all sciences, it is the *extrapolation* that matters. It is the insight (never the proof) that a local interaction of autonomous agents with noisy and incomplete information can cause higher order phenomena such as solving complex tasks that require the integration, computation and exchange of information amongst an entire population of agents. It is the model that in its abstractness and incompleteness suffices to illustrate how highly complex phenomena can be caused by simple mechanisms that by themselves would not be suspected of having far-reaching effects. Or, in still other words, it is the maybe abstract, incomplete and noisy model that in interaction with other such models might give rise to insights that none of them could have triggered on their own.

Manfred Füllsack is Professor of Systems Sciences at the University of Graz. His research includes: systems, complexity, networks, games and computational theory, work – its history, its sociology, its economy and its philosophy and computer-based simulations.

Received: 11 October 2013
Accepted: 21 October 2013

Reactive Rules Alone Cannot Construct Cognition

Martin V. Butz
University of Tübingen, Germany
martin.butz/at/uni-tuebingen.de

> Upshot • Although the authors investigate a form of distributed swarm intelligence and solve some problems with it – including sorting and summing – the major goal, which is constructing cognition, cannot be achieved by this approach alone. I propose that anticipatory mechanisms have the potential to construct cognition and may very well be combined with swarm intelligence principles.

Why cognition cannot be achieved by reactive rules

« 1 » It is out of the question that reactive rules can produce intelligent behavior – even more so when endowed with communication and interaction. Many organisms, starting from the simplest bacteria, survive and reproduce by executing mostly reactive rules. Models of such behaviors exist and have even generated robots that exhibit very clever interactions with their environments. Valentino Braitenberg's "vehicles" (Braitenberg 1984) may have illustrated such potential in the most brilliant way. Along the lines of Marvin Minsky's *Society of Mind* (Minsky 1988), Rodney Brooks (1991) has shown similar capabilities with various robots using "subsumption architectures," in which a collection of small, mostly reactive intelligent agents produces seemingly intelligent behavior.

« 2 » This is also the case for the swarm intelligence system presented in the target article: endowed with clever, local rules, the swarm manages to sort numbers or sum up a bunch of numbers. However, the reader should note where the intelligence resides in the presented architectures. The provided rules generated by the designers (i.e., the authors) yield the task-specific intelligent behavior. The systems may also be compared with the "Swiss Robots," which are able to clean up an area of blocks using extremely simple reactive rules (Pfeifer & Bongard 2008). Intelligence *emerges* via swarm intelligence.

« 3 » What is missing in these reactive approaches? Maybe the most fundamental lack is the fact that these systems and robots do not *understand* what is going on. They are not able to predict future events, they cannot imagine other states-of-affairs than those currently perceived, and they are able to act neither in anticipation of the future nor goal-directedly. Nonetheless, they may achieve goals (without them actually "knowing" it), and for an observer they may seem "intelligent" in what they are doing.

« 4 » While putting forward that cognition cannot be achieved by reactive rules, I do not want to imply that swarm intelligence is not worth further pursuit. In fact, I wholeheartedly agree that distributed, modularized systems are the key to developing higher-level cognitive systems. The authors indirectly mention the homunculus problem in §9. Representation without purpose is certainly one of the major errors made by traditional cognitivist approaches. And, as the authors elaborate on in §10, intrinsically multimodal systems with distributed processing capabilities need to be investigated further. I also agree with the proposition that modularized components should process information asynchronously and concurrently (cf. §14), as does the brain. However, the authors do not investigate principles of autopoiesis, normativity, or the emergence of goals, which may be the key components for the emergence of cognition.

How anticipations can help to construct cognition

« 5 » The authors themselves ask the question: "What constitutes the knowledge that is constructed through interaction with the environment?" (§2). However, they do not provide an answer to this question. I propose that the answer just lies in the principle of anticipation with a focus on *purposeful anticipation*.

« 6 » In a previous target article in *Constructivist Foundations* (Butz 2008), I suggested that an anticipatory drive may reside in cognitive organisms, such as us humans, with which we strive to structure incoming (sensory) information into anticipatory structures for interacting with the environment purposefully. I described how a conscious self can emerge due to the anticipatory drive in a highly complex social and cultural world.

« 7 » Similar, more mathematically-oriented theories can be found in the vision literature (Rao & Ballard 1999) and the neuroscience literature (Friston 2009). Indeed, Karl Friston proposes that a *free energy principle* may be one fundamental processing principle in the brain, and in particular in the neo-cortex. The principle is essentially based on generative, predictive models and aims at minimizing the differences between predictions and actual perceptions, where these differences are quantified by an information theoretic measure of "free energy." Due to its generality, the principle effectively subsumes the Bayesian brain hypothesis, which suggests that the brain develops a probabilistic, generative model of the world, as well as principles of efficient encoding, value-learning mechanisms, mechanisms of attention and decision-making, and optimal control principles (Friston 2010). The principle results in the compression of information and thus the minimization of undesired circumstances, such as the lack of energy, either by re-representing states of the environment in a compact form or by acting upon the environment to minimize the undesiredness of aspects of internal bodily or external states in the environment.

« 8 » What does this mean for the presented article and its focus on a stochastic diffusion search mechanism simulating swarm intelligence? I believe it suggests that the individual agents need to be endowed with predictive capabilities. A swarm intelligence system with individual predictors that compete with each other may subsume the agents put forward in the target article. When swarm intelligence systems are then allowed to interact with each other by exchanging information, passing on state estimations, and suggesting manipulative actions, the path may be paved towards an actual constructivist approach to cognition. When internal needs are added to the system, it may become purposeful by also effectively structuring the information processed in other modules for pursuing the satiation of these needs. In this way, information will be selectively processed further and the environment will not be re-represented in the sense of the good old-fashioned cognitivist approaches, but it may actually be pro-represented with the system's purpose in mind.

« 9 » While this is certainly hardly trivial to realize and should be tackled first in small settings with restricted complexity, I would like to see the constructivist community move forward in this direction. Along the way, though, it should be kept in mind that environments exhibit fundamental structural properties that can be exploited for this purpose. Our brains are not number-processing systems or even symbol-processing apparatuses, nor are they general purpose learning systems without any learning biases (remember the *no free lunch theorem*; Wolpert & Macready 1997). Anticipatory mechanisms – most likely along the lines of the free energy principle – must be designed to search for particular structural properties found in our experiences. To reveal these structural properties and to exploit them with maximum effect for developing cognitive systems in the constructivist sense is the grand challenge, yet to be solved convincingly by the community and, beyond, by researchers in cognitive science, computational neuroscience, artificial intelligence, and cognitive robotics.

Martin V. Butz is Full Professor at the Eberhard Karls University of Tübingen, in the Faculty of Science, Department of Computer Science and Department of Psychology. His research group works on cognitive modeling, focusing on how the brain develops representations of the body and the surrounding space and how these representations are used to manipulate the environment goal-directedly.

Received: 14 October 2013
Accepted 21 October 2013

Self-organization in Brains

Peter Cariani
Boston University, USA
cariani/at/bu.edu

> Upshot • Artificial life computer simulations hold the potential for demonstrating the kinds of bottom-up, cooperative, self-organizing processes that underlie the self-construction of observer-actors. This is a worthwhile, if limited, attempt to use such simulations to address this set of core constructivist concerns. Although we concur with much of the philosophical perspective in the target article, we take issue with some of the implied positions related to dynamical systems, sensorimotor contingency theory, and neural information processing. Ideally, we would like to see computational approaches more directly address adaptive, constructive processes and mechanisms operant in minds and brains. This would entail using tasks that are more relevant to the psychology of human and animal learning than performing digit sums or sorts. It also could involve relating the dynamics of agents more explicitly to ensembles of communicating neural assemblies.

History and philosophy of constructivism

« 1 » In our understanding, constructivism is, and has always been, more of a theoretical and philosophical alternative (rather than a reaction, as the introduction implies) to behaviorism, realism, and purely computational theories of mind (computationalist cognitivism, Chomskian linguistics, symbolic AI). Constructivist attitudes and ideas reach far back into history, even to the pre-Socratics (Glasersfeld 1995).

« 2 » In contrast to behaviorism, Piagetian constructivism deals with the structure and internal self-organization of minds. In contrast to realism, in which knowledge is conceived, following Plato, as "justified true belief," constructivism takes "knowledge" and "information" as purposively acquired, useful experience for understanding and acting within an ill-defined environment. Realism goes hand-in-hand with correspondence theories of truth. Mental "representations" (beliefs), for the realist, can either be true or incorrect, whereas for the constructivist, they are effective or ineffective guides for understanding, planning, and action in pursuit of some particular internal goal. In constructivist theory, pragmatist criteria of efficacy vis-à-vis some purpose replaces realist truth value criteria that depend on the veridicality of representations.

« 3 » Computational theories of mind (see Boden 2006 for a comprehensive intellectual history), focused almost completely on internal deliberation processes based on discrete symbols and rules and all but ignored sensorimotor interactions and transactions with bodies and environments.

« 4 » What makes a theory *computationalist* in this sense is not whether it is realized on a digital computer per se, but whether its aim is to demonstrate the consequences of computational rule systems apart from their embeddings in self-organizing observers and their interactions with their environments. Without such pragmatic and epistemic context, demonstrations of the power of symbol-and-rule-based systems are mathematics, not epistemology.

« 5 » In contrast to most computationalist models of mind, constructivist approaches are concerned with *all* of the processes that are involved in how the mind becomes organized through the structure of its interactions with external environments.[2] These processes encompass all internal psychological operations, i.e., perception, cognition, conation, affect, hedonics, memory, anticipation, deliberation, and action and include both analog-subsymbolic-dynamical and digital-symbolic modes of neural information processing.

« 6 » There are some inherent strengths and weaknesses of computationalist theories of mind. On the one hand, they are able to give relatively straightforward accounts of the ease with which different conceptual attributes can be combined together (Fodor & Pylyshyn 1988; Marcus 2001). This leads to the impressive combinatorics of human conceptual and linguistic productions, which are extremely difficult to realize using connectionist nets or dynamical systems. But on the other hand, models based on unitary symbol-primitives are incapable of dealing with the processes that constitute the symbols themselves and that can create entirely new ones (Carello et al. 1984; Cariani 2012). One cannot expand an alphabet by simply concatenating its existing symbols. Such closed-world symbol systems, even with their combinatoric generativity, have difficulty explaining the apparent open-ended nature of human concept formation (Piaget 1980).

« 7 » Perhaps the most critical weakness of the traditional computationalist paradigm involved an almost complete neglect of interactions between symbols and non-symbolic, environmental contexts. Sensing was reduced to streams of input symbols and actions to streams of symbolic decisions. Cycles of goal-driven, sensory-motor transactions involving learning through experience (pragmatically-relevant empirical evidence) were not part of the computationalist picture. At the height of the computationalist philosophical wave in the mid-1980s, it was even asserted that science required no measurements or observables, i.e., that pure computation could suffice (Churchland 1985; Thagard 1988). At that time, it was a radical position to propose a general model of epistemic systems in which adaptively constructed sensors and effectors played a pivotal role (Cariani 1989, 1992).

« 8 » Interaction with the environment, through sensing and effecting, is central to any epistemology, and for this reason, it should not be mystified. The authors attribute embodiment and embeddedness to "deeply coupled dynamical systems" without any explanation (§4). What does this mean exactly? One wants to think that this means that there is a rich and open-ended set of ill-defined, physical interactions available to sensory and motor interfaces. But one gets the feeling from the wording that dynamical systems are some privileged mode of description and that any approach that trades in symbolic computations, information processing, or neural mechanisms is inherently at odds with constructivism. In a similar vein, the authors appear to subscribe to radical enactivism (§6), which posits that *all* perception is bound up with sensorimotor contingencies. Although action can weakly modulate perception and vice-versa, the extreme version of the theory (no perception or consciousness without action, conscious experience not determined by organization of neural activity) is easily falsified. We perceive many sensory distinctions that have no immediate relevance in terms of goal-directed actions, and it is remarkable how invariant perception can be in different action contexts. We listen to music to modulate our internal psychological states, not to act on an external world. Yes, perception and action can interact, but the two are very different and distinct processes (Cariani 1989), and sometimes a percept is just a percept. A subtle 1% change in the concentration of an inhaled anesthetic agent, such as isoflurane, leaves all else but the organization of neural activity intact, but completely abolishes conscious awareness.

« 9 » In my opinion, constructivism does not entail a particular theory of mind and brain beyond the two stipulations that (1) knowledge arises from interaction with an environment, through perceiving and acting, and (2) that there are internal, adaptive self-modification processes that realize new epistemic functions that in turn modify behavior. In terms of constructivist theory, there is nothing inherently preferable about dynamical systems over either symbolic or analog neural signal processing descriptions. These are complementary descriptive modes – the attractor basins in the dynamical systems description correspond to the symbol states of computational models (Cariani 2001; Pattee 2012).

« 10 » One needs to choose the type of description that best addresses the particular issues at hand. If one is interested in the neural basis for construction of new mental categories, then neural information processing models are most appropriate because they operate using those dimensions of the dynamics that are relevant to particular functions. If one is interested in the neuronal substrates that realize informational processes, then detailed biophysical dynamical systems and/or realistic neural network models are the most appropriate modes of description. Artificial life simulations of swarm dynamics are useful as heuristic devices for exploring general classes of bottom-up cooperative mechanisms that might underlie self-organization in brains.

« 11 » Some of the rejection of neural information processing approaches seems to stem from realist interpretations of concepts related to "information" and "representation" that are grounded in truth values, e.g., Dretske (1981). The difficulties with realist and referentialist understand-

2| Some might find any talk of "external" environments inherently dualistic. However, the division here is between observer-actor and environment, and not a metaphysical Cartesian substance dualism between a spiritual mind and a material body. One must speak of inside and outside, internal and external in order to be able to discuss the structure of observer-actors and their transactions with their environments.

ings of mental representations have been extensively critiqued (Bickhard & Terveen 1995). But strong rejection of referentialist conceptions need not mean that all functional concepts related to signals, information, symbols, and even computation should therefore be avoided in favor of function-free descriptions cast in the language of dynamics.

« 12 » There are many senses of "information," and a number of them have meanings that are completely compatible with non-referentialist, constructivist perspectives. For example, Ross Ashby's (1956) concept of information as uncertainty reduction in the limited observer is operationalist, pragmatic, and makes no assumptions about truth or correspondence with "reality." The core idea of "genetic information" in DNA involves a pattern (the nucleotide base-pair sequence) that ultimately determines the construction of a protein amino acid polymer. Such an informational process, which switches behavior (resulting protein sequence) by different configurations (DNA coding sequence), involves a causal chain (orderly mapping), but no representation in the referentialist sense. A given DNA sequence, a specific piece of genetic information, is neither true nor false, but the sequence can produce proteins that are useful (or not) in a given cellular context.

« 13 » Similarly, other concepts that the authors seem to reject, such as neural codes, neural representations, and neural information processing, can be used effectively, without realist and computationalist baggage, in a manner that comports well with constructivism. The less ontologically-laden senses in which these terms are used in the neurosciences almost never carry the realist, referential connotations that they bear in more philosophical and computationalist discourses.

« 14 » For many years, I have worked on the problem of "neural coding of pitch" (Cariani 1995, 1999; Cariani & Micheyl 2012). What this means is not that there is some objective attribute "pitch" that is out there in the world somewhere in the acoustics that impinge on our ears. Rather, pitch is a subjective quality that is usually associated with sounds having particular periodic structure, and with particular (temporal) patterns of neural activity (as observed via sound recording and neurophysiological measuring devices). We seek to identify those neural activity patterns that covary with our subjective experiences and overt perceptual judgments of pitch. The neural correlates of pitch in the early stages of auditory processing are now well-understood, such that the pitches that are heard at any given time can be accurately predicted from the neural firing patterns that are present, all other conditions being the same (Cariani 1999).

« 15 » The problem of neural coding thus involves identifying which aspects of neural activity convey the distinctions that subserve our psychological functions and that determine the contents of our conscious experience. This is arguably the most fundamental problem in neuroscience: to understand the specific nature of the "signals of the system" that subserve mental functions (neuropsychology) and conscious awareness (neurophenomenology). It is a reverse-engineering process analogous to elucidating the genetic code. One must consider the nature of neural coding in order to begin to envision how new concepts might be created in minds and brains (Cariani 2012). If we, as constructivists, abandon the quest to understand how brains work as coherent, functional systems, then we shall never attain deep understanding of the nature of construction processes in natural minds.[3] We should therefore be very selective in what we reject and why. What makes a given account relevant to constructivism lies not in the particular description (e.g., dynamical systems good, information-processing bad), but whether and how well its interpretation addresses core constructivist questions. What is the essential functional organization of observer-actors? How do reciprocal interactions between observer-actor and their environments bootstrap construction and complexification? What are the specific processes and mechanisms of mind and brain that subserve this construction? How is open-ended creation of new concepts possible?

3 | However, it might well be possible to understand construction processes in artificial systems before we understand how they operate in biological minds/brains.

« 16 » In our search for answers to these questions, we should not be afraid to draw from psychology, the neurosciences, cybernetics, and other relevant disciplines, as well as from our own direct experiences, There are reasonable psychological models of learning and concept-formation, neural network theories that posit the kinds of brain mechanisms that could potentially be involved, and cybernetic frameworks that illuminate the functional organization of adaptive, goal-directed percept-action systems (Cariani 2011).

Swarm intelligence as a model for constructive epistemology

« 17 » Swarm intelligence has some interesting features, but it is somewhat odd a choice for virtual epistemology. It demonstrates how a collection of active, semi-independent, interacting and/or communicating elements, each using simple decision strategies based on limited, local sensory inputs can collectively realize behavioral functions that would not be possible by individuals acting completely independently of one another. Stochastic diffusion search (SDS) algorithms can be interpreted in neural terms, where the agents are neural assemblies, or in social terms, where the agents are individual human or robotic actors. Such interpretations help relate these simulations more directly to concrete human, animal, and adaptive robotic systems that we would commonly think of as examples of self-constructing epistemic systems.

« 18 » I very much appreciate the authors' candidness about the limitations of their particular simulations examples. There are always difficult choices to make regarding what to simulate in terms of level of detail and what exactly is to be explained by the whole exercise.

« 19 » One problem I do have with the swarm intelligence examples (digit summation and sorting) vis-à-vis constructivism is the relative difficulty in understanding how these collective behaviors are epistemic functions and how the agent operations involved are constructions. It is hard to see exactly what is being constructed here. Although the complex dynamics of the swarm do unfold over time, the individual agents themselves do not change. It is therefore not clear what has been learned or what new processes or structures have been constructed. The swarm simply does what it does, irrespective of its history. It has its collective dynamics, but it does not learn to do the right thing. Neither its constituent agents nor the swarm has any memory, so no internal processes change with experience. Swarm intelligence without adaptive agents looks like yet another demonstration of how to get complex global behavior from many simple interacting components and processes. That is fine, and interesting in itself, but it becomes a mathematics demonstration, not epistemology. Complex behavior, by itself, is not new function.

« 20 » This suggests that we would be better off using simulated agents that have some capacity for memory and learning, such that their percept-action functions can change with experience. In the last decade there have been some excellent demonstrations of individual and social constructions of meanings that have come out of artificial life evolutionary robotics simulations (Steels 2003) and the developmental psychology of language acquisition (Tomasello 2003). Now that plausible processes for the evolution of social cooperation and shared meanings through iterated action have been demonstrated, the next stage is to connect these processes to plausible models of minds and brains that themselves have open-ended capacities for adaptive self-organization (Steels 2008; Cariani 2012).

Peter Cariani has worked in theoretical biology, biological cybernetics, and neuroscience. His doctoral work developed a semiotics of percept–action systems, formulated a taxonomy of self-constructing adaptive systems, and explored epistemic implications of evolutionary robotics. He currently teaches courses on the neuropsychology of music and conducts research related to temporal codes in the brain.

Received: 11 October 2013
Accepted: 24 October 2013

Towards Constructive Foundations of Cognitivism: Breaking in Open Doors?

Jiří Wiedermann
Academy of Sciences
of the Czech Republic
jiri.wiedermann/at/cs.cas.cz

> Upshot • We challenge the authors' claim in the target article that "departing from cognitivism requires the development of a new functional framework that will support causal, powerful and goal-directed behavior in the context of the interaction between the organism and the environment." We argue that rather than a departure from cognitivism, the indicated goal is a natural complement or extension of the classical understanding of cognitivism. In order to reach such a goal, no new functional framework has to be developed right from scratch: there are many insights in related areas of research that can serve such a purpose well and can become an integral part of constructive cognitivism. We welcome the idea to build constructive foundations of cognitivism.

Extending the scope of cognitivism

« 1 » In §1 and §2, the authors of the article briefly recapitulate the development of constructivist views since Piaget and claim that despite decades of development, constructivist perspectives have mostly remained applied (e.g., in education) or descriptive. This is partially correct since traditionally, and until recently, less attention has been paid to the underlying mechanisms participating in the knowledge generation process, quite understandably. However, especially in the past few decades, advances in the related research areas have brought new insights into the processes underlying cognition. Inevitably, this new knowledge is also ready to find its way into cognitivist theory, not as a departure from it or an alternative to it (as the authors claim, e.g., in the structured abstract of their work and in §2), but as a natural complement and extension of a narrow understanding of cognitivism. In this new perspective, cognitivism is seen as the theory describing methods of how humans (and in general, cognitive systems) generate knowledge leading to the development of an individual's cognition abilities, and the mechanisms (or algorithms) causing such effects. Effectively, this means a shift from the descriptive approach to a procedural approach, a shift from describing *what* a real cognitive system is doing to *how* it is doing what it is doing. Extending cognitivism by a theory dealing with the mechanisms supporting knowledge acquisition and development of human cognitive abilities will enable understanding of cognitive processes through their implementation, too. At the same time, such an approach will provide a framework for testing the classical constructivist theories of cognition as applied, e.g., in education. For the cognitivistic methods, it will offer a means of their evaluation and tailoring.

The theoretical sources of constructive cognitivism

« 2 » Naturally, the authors are aware of new developments in the area of computationalism and of the impacts of these developments on cognitivism. In fact, it is exactly this impact and the methods and mechanisms of achieving it that are the focal point of current efforts in all areas of the related research (some works are mentioned in the paper). However, as regards the inventory of the associated research, I would opt for a more systematic and broad approach. The sources of inspiration and, indeed, of ready-made knowledge to be used in this context are really quite rich. These sources are so numerous and diverse that it is not possible to mention all of them in this commentary. Nevertheless, I will try to list at least those sources that seem to me to be most relevant to the paper at hand and that I expected to see somewhere in the introductory part of the paper. In doing so, I abstain from referencing the basic or survey works in the widely-known research fields

« 3 » First of all, there is epistemology dealing with the philosophical study of the nature and grounds of knowledge. Along these lines, Wiedermann & van Leewuen's (2013) recent effort to see computations as knowledge generating processes deserves special attention. In fact, this definition of computation enlarges its domain to the entire cognitive processes, answering thus the question "What is cognition." Then there is philosophy of mind, investigating the nature of mind, whose main attributes are skills in acquiring and creating new knowledge. Not only are descriptions of such skills and their development being investigated but increasingly more often, the underlying mechanisms are, too. In more concrete terms, this eventually leads to the computational models of the mind (or of the brain, for that matter). Dealing (so far) with a lower "intellectual" level of cognition, developmental robotics has achieved significant progress in the field of embodied cognition. All the previous research efforts are encompassed in the general field of artificial intelligence (AI) and, recently, in artificial general intelligence (AGI). Further, we have general computational models as a part of the computational complexity theory, focusing on the computability and efficiency of various means of computation that can be ascribed to cognitive systems. Such models, especially the so-called non-standard models, deviate substantially from the classical Turing machine model and the computational paradigm represented by it, and deserve special attention in cognitivism. This is because they aim at the investigation of the resources specific to cognition, such as interactivity, non-uniform evolution, learning or non-terminating behavior – the properties characteristic of cognitive systems and thus important from the view point of the extended definition of cognitivism. I will concentrate upon such properties in the next section. Finally, there is a vast body of life sciences (especially psychology, animal psychology, neurosciences, biology, ethology, linguistics, etc.) that have their say in cognitivism, too.

The computational basis of constructive cognitivism

« 4 » When modeling a cognitive system, it is of utmost importance to choose the right computational model, especially if one wants to model a certain cognitive feature. Since the set of computational models available for modeling cognitive tasks is too broad to be even briefly mentioned in this commentary, I will only consider the models suitable for the task considered by the authors.

« 5 » In their paper, the authors have concentrated on the interaction between the organism and the environment, aiming to identify the "computational power of … interaction," as they call it in the structured abstract. The computational power of interaction has been studied in the field of complexity theory for years (cf. van Leeuwen & Wiedermann 2006, and the respective edited volume) and it is relatively well characterized. A less-known result in this field that might be of interest in the context of cognitivism is probably the one related to the super-Turing computational power of evolving interactive computations (cf. Wiedermann & van Leeuwen 2002). However, it appears that the computational power of interaction (at least in the standard sense as used, e.g., in computational complexity theory) is not the property of computation the authors are after. Rather, it is *the knowledge acquisition ability of a computation* that interests the authors. This is seen from the properties of computation affecting this ability that the authors consider. Among these properties, the authors have included "partial information from the environment, exploration, distributed processing and aggregation of information, emergence of knowledge and directedness towards relevant information" (structured abstract).

« 6 » To study the aforementioned properties of a computation, the authors have chosen the less-known Stochastic Diffusion Search algorithm (SDS, cf. §16), introduced by one of the co-authors, Mark Bishop (1989). This is a family of algorithms that originates from the field of swarm intelligence, exploring the collective intelligence of populations of agents. This is a reasonable choice for the purpose at hand, although in the meantime other suitable and more formal models with a complete characterization of their computational power have emerged (cf. Angluin's 2007 population protocols or Wiedermann's 2013 amorphous computing systems).

« 7 » The lack of formal description of the SDS model makes it difficult to assess the parts of the paper that give examples of computational abilities of SDS algorithms. For instance, in §17 an SDS algorithm is described for locating a substring in a string of symbols. Apparently, each agent must be able to access a random location in the string but it remains unclear what the mechanism is behind the task. It must then remember a series of adjacent symbols. For that, a finite memory will do. Can agents be finite automata? In the diffusion phase, the agents must communicate. What is the communication mechanism? Is it one- or two-way communication, is it synchronous or asynchronous? What happens if the target strings contained repeated substrings? This would mislead the agents unless they have remembered the positions of symbols they have observed. Then they cannot be finite automata. Etc.

« 8 » Similar problems arise in other examples. In §19, digit summation is considered. This does not appear to be a typical cognitive task (the references in §18 to other, mostly universal models of computations solving the same tasks is of no help in this context), but let it be. It appears that the agents must be able to add two numbers and store the result in their internal memory. That is, they cannot be finite automata. They also possess the ability to modify the search space and to communicate with each other. Again, the details of the model and of the computational scenario are unclear. This same holds for the next example of sorting, in §25, which should illustrate communication among the agents. Did not the agents communicate already in the previous examples? Must not the agents have a sense of direction in order to be able to explore the search space on either side of the search space, around the initial location (§28)?

« 9 » In general, all three examples of SDS computation given above are not typical examples of cognitive processing. This is because although they use interaction, they do not capture the remaining two main properties of cognitive computations – viz. evolution (learning) and potentially unbounded operation (cf. Wiedermann & van Leeuwen 2002).

« 10 » Nevertheless, in spite of this criticism, I agree with the conclusion at the end of §18 that

> "even though agents' abilities are minimal, the population of agents as a whole, interacting with the environment, is able to organize and consolidate a solution to such complex problems, without referring to a representation of the external world."

« 11 » I see the previous examples as the first trials to explore the knowledge-acquiring and processing abilities of families of simple embodied computational devices in a bottom-up manner, starting with the simplest gadgets and the simplest cognitive tasks. To me, it appears that the first thing to do is to specify the set of simple cognitive tasks and scenarios out of which solutions of more complex cognitive tasks can be composed, possibly under more complex scenarios. Among such tasks, I would include mirroring of observed signals, learning of sequences of signals, reinforced learning, imitation learning, communication via simple actions, etc. The literature is full of similar approaches but I am not aware of work following the evolutionary sequence of cognitive tasks in parallel with the evolutionary sequence of increasingly more complex (embodied) computational devices. This could be a rewarding task for constructive cognitivism.

Conclusion

« 12 » Although I welcome the idea to build the constructive foundations of cognitivism, whose very first and timid germs I see in the commented paper, and although I have indicated that there is a tremendous amount of results and insights to begin with, it is not going to be an easy enterprise. The doors mentioned in the title of this commentary are wide open, but an immense amount of stuff has to be categorized, ordered, critically evaluated and sorted so that only the minimal body necessary to build, maintain and develop a coherent constructive theory of cognitivism remains. If properly done, cognitivism and its applications will benefit much from such an endeavor.

Jiří Wiedermann is a computer scientist specialized in computability, complexity theory and non-standard models of computations, with interests in cognitive computations and epistemology. He is a member of the Czech Learned Society and of Academia Europaea. From 2000 to 2012, he served as the director of the Institute of Computer Science of the Academy of Sciences of the Czech Republic in Prague.

Received: 8 October 2013
Accepted: 22 October 2013

Single Agents Can Be Constructivist too

Olivier L. Georgeon
Université de Lyon, France
olivier.georgeon/at/liris.cnrs.fr

Salima Hassas
Université de Lyon, France
salima.hassas/at/univ-lyon1.fr

> Upshot • We support Roesch and his co-authors' theoretical stance on constructivist artificial agents, and wish to enrich their "exploration of the functional properties of interaction" with complementary results. By revisiting their experiments with an agent that we developed previously, we explore two issues that they deliberately left aside: autonomous intentionality and dynamic reutilization of knowledge by the agent. Our results reveal an alternative pathway to constructivism that addresses the central question of intentionality in a single agent from the very beginning of its design, suggesting that the property of distributed processing proposed by Roesch et al. is not essential to constructivism.

« 1 » In their paper, Etienne Roesch and his coauthors formulate a constructivist approach to artificial learning in which

> "knowledge of the world, for an individual, is created from the interaction with the environment, rather than existing in an ontic reality, supposedly pre-existing or available to registration from the physical world." (§1)

They propose three models to illustrate this idea, in which a swarm of agents performs different tasks in an environment made of a string of digits. We fully agree with this theoretical stance but we feel that these models do not illustrate it as well as possible. In particular, one might argue that the swarm's knowledge does, in fact, "exist in an ontic reality" since the agents directly "perceive" the digits and apply predefined rules to process the digits for the purpose intended by the designer.

« 2 » Here, we present an alternative model that does not make the knowledge of the environment directly available to registration by the agent. The environment is the string of digits presented in §23, and the agent was designed to produce similar results as proposed in §25: sorting the string. Yet the agent's observations are reduced to a single bit whose significance depends on the dynamics of the agent's interactions rather than directly reflecting the state of the environment. The agent remains "unaware" that it "exists" at a particular position in a string of digits, and its own goal is not to sort this string. For the agent, the construction of knowledge consists of learning to organize its behavior to fulfill a form of intentionality defined independently of the environment.

Implementation

« 3 » *Initialization*: The environment is a string of 10 digits $E_0 = [6,3,5,4,7,3,5,3,9,5]$ plus an integer p in the interval $[0,9]$ that represents the agent's position. $E_t[p_t]$ denotes the digit at the agent's position at time t. At time 0, $p_0 = 0$, thus the current digit $E_0[p_0] = 6$.

« 4 » *Behaviors*: At time t, the agent chooses an action from amongst the set of three possible actions $A = \{step, feel, swap\}$, and then receives a binary observation from amongst the set of two possible observations $O = \{true, false\}$. The set $A \times O$ thus contains 6 possible *interactions*. The agent initially ignores the meaning of actions and observations, i.e., it implements no rule to process them specifically. However, each interaction has a predefined *valence* that plays a role in defining the agent's intentionality, as explained below. Unbeknownst to the agent, *step* consists of stepping to the next digit. If this action takes the agent to a greater or equal digit then it produces the observation *true* and has a positive valence, otherwise, it produces the observation *false* and has a strongly negative valence. *Feel* consists of testing whether the next digit is greater than or equal to the current one; if yes, it produces *true*, otherwise *false*. *Feel* interactions have a mildly negative valence. *Swap* consists of trying to swap the current digit with the next. It succeeds only if the current digit is greater than the next, producing the observation *true* and a positive valence; otherwise it does nothing and produces the observation *false* and a strongly negative valence. When the agent is at position 9, *step* returns the agent to position 0, *swap* does nothing, and the three actions produce the observation *false*. Table 1 summarizes the implementation of these possibilities of interaction.

« 5 » *Agent*: We used an agent presented previously (Georgeon & Ritter 2012), which was programmed to exhibit two forms of intentionality: the tendency to select sequences of actions that produce well-predicted observations, and the tendency to enact positive interactions while avoiding strongly negative interactions. The former type of intentionality relates to Steels's (2004) autotelic principle (the enjoyment of being *in control* of one's activity), and was implemented as a tendency to record, hierarchically organize, and appropriately re-enact sequences of interactions that capture regularities in the coupling between the agent and the environment. The latter is called interactional motivation (Georgeon, Marshall & Gay 2012), and was implemented through preferentially engaging in sequences of interactions that have the highest total valence.

Action	Condition	Effect	Observation	Interaction	Valence
step	$p_t<9$ and $E_t[p_t] \le E_t[p_t+1]$	$p_{t+1}=p_t+1$	true	step_up	4
	$p_t<9$ and $E_t[p_t] > E_t[p_t+1]$	$p_{t+1}=p_t+1$	false	step_down	–10
	$p_t=9$	$p_{t+1}=0$	false	step_down	–10
feel	$p_t<9$ and $E_t[p_t] \le E_t[p_t+1]$	–	true	feel_up	–4
	$p_t<9$ and $E_t[p_t] > E_t[p_t+1]$	–	false	feel_down	–4
	$p_t=9$	–	false	feel_down	–4
swap	$p_t<9$ and $E_t[p_t] \le E_t[p_t+1]$	–	false	not_swap	–10
	$p_t<9$ and $E_t[p_t] > E_t[p_t+1]$	$E_{t+1}[p_t+1]=E_t[p_t]$ $E_{t+1}[p_t]=E_t[p_t+1]$	true	swap	4
	$p_t=9$	–	false	not_swap	–10

Table 1: Possibilities of interaction available to the agent.

Results

« 6 » Table 2 reports selected strips of behaviors, with the current digit marked in bold. The agent started by randomly picking the *step* action at times 1 and 2. Over time, the agent organized its behavior as if it had discovered that the *feel* action could be used to test the next digit. If this action resulted in the *feel_up* interaction, then the *step_up* interaction could subsequently be enacted, otherwise, the *swap – step_up* sequence could subsequently be enacted. This dynamics resulted in the behavior of "carrying digits to the right." This behavior is illustrated in Table 2 from time 106 to 113: the agent "carried" the "5" digit from position 2 to 4, by repeating the *feel_down – swap – step_up* sequence until the "5" digit got "blocked" by a greater or equal digit (another "5" digit at position 5). This behavior resulted in the string being entirely sorted at time 130.

Time	Interaction	Environment
0	–	**6** 3 5 4 7 3 5 3 9 5
1	step_down	6 **3** 5 4 7 3 5 3 9 5
2	step_up	6 3 **5** 4 7 3 5 3 9 5
…		
106	feel_down	3 4 **5** 3 3 5 5 6 7 9
107	swap	3 4 **3** 5 3 5 5 6 7 9
108	step_up	3 4 3 **5** 3 5 5 6 7 9
109	feel_down	3 4 3 **5** 3 5 5 6 7 9
110	swap	3 4 3 **3** 5 5 5 6 7 9
111	step_up	3 4 3 3 **5** 5 5 6 7 9
112	feel_up	3 4 3 3 **5** 5 5 6 7 9
113	step_up	3 4 3 3 5 **5** 5 6 7 9
…		
130	swap	3 3 **3** 4 5 5 5 6 7 9

Table 2: Behavior strips.

« 7 » Figure 1 reports the agent's behavior until time 200 in terms of what matters to the agent: the enacted interactions, their valence, and the level of control that the agent has over its activity manifested by the length of the sequences intentionally enacted.

« 8 » In summary, this experiment helps clarify the distinction between the designer's goal (sorting the string, illustrated in Table 2) and the agent's intentionality (being in control and enacting interactions that have positive valence, illustrated in Figure 1). While the agent remained unaware of the underlying structure of the environment, it learned to master sensorimotor contingencies as if it enjoyed being able to predict its activity and to "step up," and disliked "stepping down" and failing to swap digits. The agent learned to use the *feel* action – in spite of its negative valence and the ignorance of its meaning – as an active perception of the environment to inform subsequent behaviors. This activity illustrates the property pointed out by Roesch et al. that "perception is an integral part of the process from which knowledge of the world arises" (§7) and that "Exploration of the environment provides the organism with the ability to sense and become attuned to the laws governing change" (§8). We believe that these properties, associated with the capacity of the agent to engage in incremental learning, qualify the agent as a candidate to illustrate key aspects of constructivism.

« 9 » Roesch et al. conclude by listing the properties of interactions that they judge to be paramount to the constructivist approach: "partial information, exploration, distributed processing and aggregation of information, emergence of knowledge and directedness towards relevant information" (§40). Our results support all of these except the *distributed processing* property (insofar as it applies to a swarm of agents), and suggest the additional property of *intrinsic intentionality*.

Olivier L. Georgeon is currently an associate researcher at the LIRIS Lab with a fellowship from the French government (ANR-RPDOC program). He received a master's in computer engineering from Ecole Centrale de Marseille in 1988, and a PhD in psychology from the Université de Lyon in 2008.

Salima Hassas leads a research activity in the area of multi-agent systems with self-*(self-adaptation, self-organization, etc.) abilities. Her research is related to complex systems, bio-inspired computing, constructivist approaches to cognition and enaction. She is an associate editor of ACM TAAS, and serves on scientific committees for major conferences in her research domain.

Received: 11 October 2013
Accepted: 18 October 2013

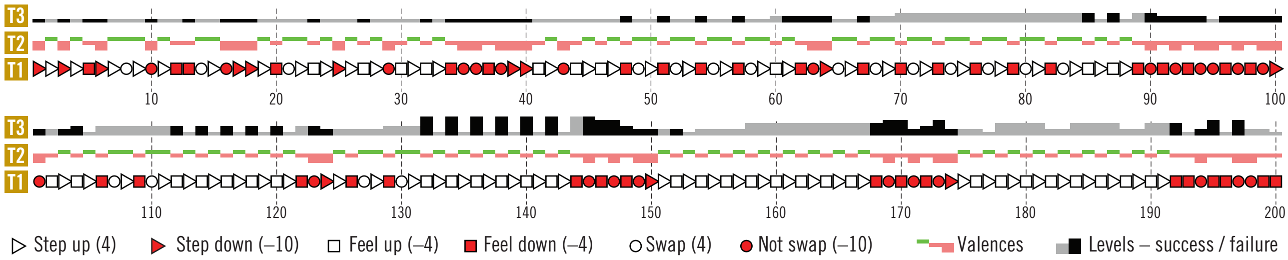

Figure 1: Analysis of the first 200 interactions enacted by the agent. Tape T1: the enacted interactions (the shape represents the action and the color the resulting observation). Tape T2: the valence of the enacted interactions displayed as a bar graph (green when positive, red when negative). Tape T3: The length of the sequences intentionally enacted, displayed as a bar graph. Higher levels of gray indicate better control over the activity; black segments indicate that an intended sequence was interrupted due to the failure to predict correctly the resulting observation. This trace shows that the behavior was unorganized approximately until time 40 (no regularities in the symbols in T1 and the presence of step_down and not_swap interactions that have strong negative valence represented by high red bars in T2). The agent intentionally enacted the second order sequence swap – step_up for the first time during time 68–69 (second level in T3), then the third-order sequence feel_down – swap – step_up during time 70–72 (third level in T3), repeating this sequence until time 85. After time 130, the digits were entirely sorted, and the agent engaged in repeating the sequence feel_up – step_up, except when reaching the end of the string, in which case it continued experimenting with other behaviors (episodes 144–150, 168–174, and 192–200). After time 310 (not shown), the agent resigned itself to merely enacting the step_down interaction when reaching the end of the string, acknowledging that it had no better possibilities.

Authors' Response: Learning, Anticipation and the Brain

Etienne B. Roesch et al.

> **Upshot** • Albeit mostly supportive of our work, the commentaries we received highlighted a few points that deserve additional explanation, with regard to the notion of learning in our model, the relationship between our model and the brain, as well as the notion of anticipation. This open discussion emphasizes the need for toy computer models, to fuel theoretical discussion and prevent business-as-usual from getting in the way of new ideas.

« 1 » Computer simulations hold great potential for charting unexplored territories. This applies especially to constructivism, which now faces the challenge of explaining the mechanisms that support knowledge acquisition. However, this potential comes at a price, i.e., that of forcing oneself to remain humble as to what can reasonably be expected of the tool. Constructivism is an "all-embracing theory," as pointed out by **Manfred Füllsack** (§2), that demands a delineation of scope that goes beyond the mere description of the organism's internal states. In this target article, we aimed to explore aspects of the interaction between the organism and the environment that support what might be called a constructivist process; aspects such as partial information, exploration, distributed processing and aggregation of information, emergence of knowledge and directedness towards relevant information.

« 2 » We do not mystify interaction – and our single-minded focus can indeed be misinterpreted (**Peter Cariani** §8) – for that would be casting a blind eye over an impressive body of data; nor do we subscribe to the mystification of internal mechanisms often found in mainstream cognitive science. In placing such a focus on aspects of interaction, we wish to highlight the contribution of these aspects to causal, powerful and goal-directed behavior, alongside more internal mechanisms, which we voluntarily did not include in our models. These passive, dynamic effects, also called "nontrivial causal spread" (Wheeler 2005: 250), contribute to cognition and our interaction with the world in a significant way (Pfeiffer & Bongard 2007), and are all too often ignored. If there is one family of theories that ought to bring credit to the interaction between the organism and the environment, it is constructivism.

« 3 » To put it differently, we concur with the notion that constructivism does not imply a particular theory of mind and brain beyond two simple axioms:

> "(1) knowledge arises from interaction with an environment, through perceiving and acting, and (2) there are internal, adaptive, self-modification processes that realize new epistemic functions that in turn modify behavior" (**Cariani** §9).

But, in the work presented, we simply wished to ascertain criteria of sufficiency and necessity for these axioms. Arguably, in the limited context we chose, our results seem to imply that the first axiom is sufficient to achieve our set goals.

« 4 » The commentaries on our target article emphasise the need for such toy models to fuel theoretical discussion, and made our work, and that of the editors and reviewers of this special issue, worthwhile. Computer simulations demand practical moves to fit both engineering and theoretical constraints. This practice aspires to, and is bound to make some of us feel uneasy, for it is only through the world of our experiences, good and bad, that we construct ourselves – and our theories about the mind.

« 5 » Hereafter, we go over some of the points that arose, and provide some clarification where needed.

Learning

« 6 » Most, if not all, of the commentators rightfully pointed out that our agents do not exhibit special abilities, such as that of memorizing a history of interaction with the environment, which one might see as a constructivist way of acquiring knowledge. It was our ambition, however, not to endow our systems with too much sophistication, which would have prevented our analysis of the very low-level properties of interaction. Our models nonetheless embody a form of learning in two ways. First, at the level of the population of agents, the state of the swarm remains stable for a significant amount of time, and influences the unfolding of future interactions. The swarm/population can be said to hold a representation – for lack of a better word – of its interaction with the environment, which serves as a foundation to acting. This transient state of activation can be observed in Figure 4 in our target article, for instance, which depicts the clustering of agents in the organism's perceptive field. As the swarm reorganizes itself around the hypotheses it holds, and acts upon the environment, it remains stable for a significant amount of time (time steps). A more formal characterisation of SDS's ability to converge towards global optima has been described by Nasuto & Bishop (1999). Second, as the swarm interacts with the environment and quite literally modifies it, it may not be unreasonable to depict the environment as a memory for the swarm, albeit external to the organism. Here, one must be careful not to read into this statement explicit support for the often oversimplified extended mind thesis (Clark & Chalmers 1998). In the work presented, our goal has not been to provide a definite and conclusive definition of the mind, cognition and the mechanisms at play; rather, we point out aspects of interaction that can yield further theoretical development. The footprint of the observer-actor's interaction in the environment must not be discarded on the basis that it outgrows traditional boundaries.

SDS and neural assemblies

« 7 » A second point of discussion concerns the use of SDS as a metaphor, and swarm intelligence as an appropriate level of description for the processes under study. In particular, due to its nature, it is easy to misconstrue swarm intelligence as a plain rejection of neural representations and neural information processing (Cariani §13). This is particularly the case here, given the philosophical discussion introducing our article, which opposed cognitivism to constructivism and turned more readily to sensorimotor contingency theory or radical enactivism.

« 8 » This misunderstanding spawns from the qualitative difference between swarm intelligence and neural networks. At least in the case of SDS, which is a particular case of swarm intelligence, previous work from our group showed that a connectionist framework based on spiking neurons was able to implement a robust and efficient search akin to the operations of its swarm counterpart (Nasuto, Bishop & De Meyer 2009). Consequently, one can imagine reproducing the results presented in the target article within a network of neural assemblies. In this new empirical situation, the currency would thus become spatio-temporal neural activity patterns that covary with the network's experience of the environment. Whether these neural activations qualify as symbolic or subsymbolic representations of the external world is another issue, which goes beyond the scope of the present discussion. In other words, if we used SDS as an abstract representation to investigate aspects of interaction, this pragmatic move was primarily dictated by technical and empirical considerations, and should not be seen as a claim favouring swarm intelligence over other models of the mind and neural processing. Other models may, of course, have been as appropriate (Jiri Wiedermann §6), and may have provided fruitful insight into the contribution of learning and unbounded operation to knowledge acquisition.

Anticipation and emergence of "purposeful anticipation"

« 9 » Both Olivier Georgeon & Salima Hassas, and Martin Butz present ideas supportive to the notion that some core intelligence can emerge from the swarm. They explore the extent to which individual agents may contribute to the organism's behaviour, by limiting the agents' awareness of the environment and by supplementing the swarm with anticipatory mechanisms, respectively. Albeit radically different, both these approaches are spot on and extremely interesting for the present discussion. Our models, based on SDS, do not represent specific biological mechanisms; they do, however, highlight aspects of interaction that may be shown to play crucial roles in the way biological organisms evolve in their environment. Similarly, it is very likely that some part of the organism (any organism) is more or less attuned to the environment whereas some other part may rely on anticipatory mechanisms. A comprehensive constructivist model may thus need to account for both types of mechanisms to account for different types of interactions.

« 10 » Interestingly, subsequent work on the models we presented builds on anticipatory mechanisms. In particular, we operationalized a strong form of anticipation that arises from the natural coupling of both the observer-actor and the environment, based on the delayed feedback between the physical elements of this system, the properties of their synchronization and the strength of the coupling between them (Stepp & Turvey 2010; Roesch, Nasuto & Bishop 2012; Spencer et al. in press).

Acknowledgements

We wish to thank our reviewers and all commentators; most were supportive of our efforts, and all provided us with constructive ideas, which will certainly benefit future work.

Received: 6 November 2013
Accepted: 9 November 2013

Combined references

Afek Y., Alon N., Barad O., Hornstein E., Barkai N. & Bar-Joseph Z. (2011) A biological solution to a fundamental distributed computing problem. Science 331: 183–185.

Anderson M. L., Richardson M. J & Chemero A. (2012) Eroding the boundaries of cognition: Implications of embodiment. Topics in Cognitive Science 4(4): 717–730.

Angluin D., Aspnes J., Eisenstat D.& Ruppert E. (2007) The computational power of population protocols. Distributed Computing 20(4): 279–304.

Ashby W. R. (1956) An introduction to cybernetics. Chapman and Hall, London.

Banchereau J. & Steinman R. M. (1998) Dendritic cells and the control of immunity. Nature 392: 245–252.

Bedau M. A. (1997) Weak emergence. Philosophical Perspectives 11: 375–399.

Berdahl A., Torney C. J., Ioannou C. C., Faria J. J & Couzin I. D. (2013) Emergent sensing of complex environments by mobile animal groups. Science 339: 574–576.

Bickhard M. H. (2006) Developmental normativity and normative development. In: Smith L. & Voneche J. (eds.) Norms in human development. Cambridge University Press, Cambridge: 57–76.

Bickhard M. H. & Terveen L. (1995) Foundational issues in artificial intelligence and cognitive science: Impasse and solution. Elsevier, New York.

Bishop J. M. (1989) Stochastic searching networks. In: Proceedings of the First IEE conference on Artificial Neural Networks. IEE Conference Publications, London: 329–331.

Boden M. A. (2006) Mind as machine: A history of cognitive science. Oxford University Press, Oxford.

Braitenberg V. (1984) Vehicles: Experiments in synthetic psychology. MIT Press, Cambridge MA.

Brooks R. A. (1991) Intelligence without representation. Artificial intelligence 47: 139–159.

Butz M. V. (2008) How and why the brain lays the foundations for a conscious self. Constructivist Foundations 4(1): 1–14 & 32–37. Available at http://www.univie.ac.at/constructivism/journal/4/1/001.butz

Carello C., Turvey M., Kugler P. N. & Shaw R. E. (1984) Inadequacies of the computer metaphor. In: Gazzaniga M. S. (ed.) Handbook of cognitive neuroscience. Plenum, New York: 229–248.

Cariani P. (1989) On the design of devices with emergent semantic functions. Unpublished Ph.D. thesis at the State University of New York at Binghamton.

Cariani P. (1992) Emergence and artificial life. In: Langton C. G., Taylor C., Farmer J. D. & Rasmussen S. (eds.) Artificial life II. Addison-Wesley, Redwood City CA: 775–798.

Cariani P. (1995) As if time really mattered: Temporal strategies for neural coding of sensory information. Communication and Cognition – Artificial Intelligence 12(1–2): 161–229. Reprinted in: Pribram K. (ed.) (1994) Origins: Brain and self-organization. Lawrence Erlbaum, Hillsdale NJ: 208–252.

Cariani P. (1999) Temporal coding of periodicity pitch in the auditory system: An overview. Neural Plasticity 6(4): 147–172.

Cariani P. (2001) Symbols and dynamics in the brain. Biosystems 60(1–3): 59–83.

Cariani P. (2011) The semiotics of cybernetic percept–action systems. International Journal of Signs and Semiotic Systems 1(1): 1–17.

Cariani P. (2012) Creating new primitives in minds and machines. In: McCormack J. & D'Inverno M. (eds.) Computers and creativity. Springer, New York: 395–430.

Cariani P. & Micheyl C. (2012) Towards a theory of infomation processing in the auditory cortex. In: Poeppel D., Overath T. & Popper A. (eds.) Human auditory cortex: Springer handbook of auditory research. Springer, New York: 351–390.

Churchland P. M. (1985) The ontological status of observables. In: Churchland P. M. & Hooker C. A. (eds.) Images of Science (Chicago: University of Chicago).

Churchland P. M. (2005) Functionalism at forty: A critical retrospective. The Journal of Philosophy 102(1): 33–50.

Clark A. & Chalmers D. (1998) The extended mind. Analysis 58(1): 7–19.

De Meyer K., Bishop J. M & Nasuto S. J. (2000) Attention through self-synchronisation in the spiking neuron stochastic diffusion network. Consciousness and Cognition 9(2): S81.

Dretske F. I. (1981) Knowledge & the flow of information. MIT Press, Cambridge MA.

Dretske F. I. (2003) Experience as representation. Philosophical Issues 13(1): 67–82.

Fodor J. A. (1983) The modularity of mind. An essay on faculty psychology. MIT Press, Cambridge MA.

Fodor J. A. & Pylyshin Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. Cognition 28(1–2): 3–71.

Foerster H. von (1972) Notes on an epistemology for living things. BCL Report. No. 9.3 (BCL Fiche No. 104/1), Biological Computer Laboratory, Department of Electrical Engineering, University of Illinois, Urbana IL. Reprinted in: Foerster H. von (1981) Observing systems. Intersystems Publications, Seaside CA: 258–265.

Friston K. (2009) The free-energy principle: A rough guide to the brain? Trends in Cognitive Sciences 13: 293–301.

Friston K. (2010) The free-energy principle: A unified brain theory? Nature Reviews Neuroscience 11: 127–138.

Georgeon O. & Ritter F. (2012) An intrinsically-motivated schema mechanism to model and simulate emergent cognition. Cognitive Systems Research 15–16: 73–92.

Georgeon O., Marshall J., & Gay S. (2012) Interactional motivation in artificial systems: Between extrinsic and intrinsic motivation. In: Proceedings of the Second International Conference on Development and Learning, and on Epigenetic Robotics (EPIROB2012). San Diego CA: 1–2.

Glasersfeld E. von (1984) An introduction to radical constructivism. In: Watzlawick P. (ed.) The invented reality. Norton, New York: 17–40. Available at http://www.vonglasersfeld.com/070.1

Glasersfeld E. von (1995) Radical constructivism: A way of knowing and learning. Falmer Press, London.

Glasersfeld E. von (2005) Thirty years radical constructivism. Constructivist Foundations 1(1): 9–12. Available at http://www.univie.ac.at/constructivism/journal/1/1/009.glasersfeld

Hofstadter D. R. and the Fluid Analogies Research Group (1995) Fluid concepts and creative analogies. Basic Books, New York.

Korsten N., Roesch E. B., Fragopanagos N., Taylor J. G. & Sander D. (2011) Biological computational constraints to the psychological modelling of emotion. In: Petta P., Pelachaud C. & Cowie R. (eds.) Handbook for research on emotions and human-machine interactions – HUMAINE. Springer, Berlin: 132–145.

Marcus G. F. (2001) The algebraic mind: Integrating connectionism and cognitive science. MIT Press, Cambridge MA.

Maturana H. R. & Varela F. J. (1980) Autopoiesis and cognition: The realization of the living. Reidel: Dordrecht.

Minsky M. (1988) The society of mind. Simon & Schuster, New York.

Nasuto S. & Bishop M. (1999) Convergence analysis of stochastic diffusion search. Parallel Algorithms and Applications 14(2): 89–107.

Nasuto S. J., Bishop J. M. and De Meyer K. (2009) Communicating neurons: A connectionist spiking neuron implementation of stochastic diffusion search. Neurocomputing 72(4–6): 704–712.

O'Regan J. K. & Noë A. (2001) A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24(5): 939–1031.

Öhman A., Carlsson K., Lundqvist D. & Ingvar M. (2007) On the unconscious subcortical origin of human fear. Physiology & Behavior 92(1–2): 180–185.

Pattee H. H. (2012) Discrete and continuous processes in computers and brains. In: Pattee H. H. & Raczaszek-Leonardi J. (eds.) Laws, language and life. Springer, Dordrecht. Originally published in: Conrad W. & Güttinger W. (eds.) (1974) Physics and mathematics of the nervous system. Springer, Berlin: 125–142.

Pfeifer R. & Bongard J. C. (2006) How the body shapes the way we think: A new view of intelligence. MIT Press, Cambridge MA.

Piaget J. (1937) La construction du réel chez l'enfant. Delachaux & Niestlé, Neuchâtel. English translation: Piaget J. (1957) The construction of reality in the child. Routledge & Kegan Paul, London.

Piaget J. (1980) The psychogenesis of knowledge and its epistemological significance. In: Piatelli-Palmarini M. (ed) Language and learning. The debate between Jean Piaget and Noam Chomsky. Harvard University Press, Cambridge MA: 23–34.

Rao R. P. N. & Ballard D. H. (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience 2: 79–87.

Roesch E. B., Nasuto S. J & Bishop J. M. (2012) Emotion and anticipation in an enactive framework for cognition (Response to Andy Clark). Frontiers in Psychology 3(398): 1–2.

Scheutz M. (2002) Computationalism. New Directions. MIT Press, Cambridge MA.

Sieg W. & Byrnes J. (1996) K-graph machines: Generalizing Turing's machines and arguments. In: Hajek P. (ed.) Gödel '96. Springer, New York: 98–119.

Smock C. D. & Glasersfeld E. von (1974) Epistemology and education. Follow Through Publications, Athens GA.

Spencer M., Roesch E. B., Bishop J. M. & Nasuto S. J. (in press) Emergent representations from distributed interactive dynamics. In: Müller V. C. (ed.) Fundamental issues of artificial intelligence. Springer, Berlin.

Steels L. (2003) Evolving grounded communication for robots. Trends in Cognitive Sciences 7(7): 308–312

Steels L. (2004) The autotelic principle. In: Fumiya I., Pfeifer R., Steels L., & Kunyoshi K. (eds.) Embodied artificial intelligence. Springer, New York: 231–242.

Steels L. (2008) The symbol grounding problem has been solved. So What's next? In: de Vega M. (ed.) Symbols and embodiment: Debates on meaning and cognition. Oxford University Press, Oxford: 223–244.

Thagard P. (1988) Computational philosophy of science. MIT Press, Cambridge MA.

Tomasello M. (2003) Constructing a language. A usage-based theory of language acquisition. Harvard University Press, Cambridge MA.

Turing A. M. (1936) On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society (Series 2) 42: 230–265.

van Leeuwen J. & Wiedermann J. (2006) A theory of interactive computation. In: Goldin D., Smolka S. & Wegner P. (eds.) Interactive computation: The new paradigm. Springer, New York: 119–142.

Wheeler M. (2005) Reconstructing the cognitive world. MIT Press, Cambridge MA.

Wiedermann J. (2013) The many forms of amorphous computational systems. In: Zenil H. (ed.) A computable universe: Understanding and exploring nature as computation. World Scientific: Singapore: 243–256.

Wiedermann J. & van Leeuwen J. (2002) The emergent computational potential of evolving artificial living systems. AI Communications 15(4): 205–216.

Wiedermann J. & van Leeuwen J. (2013) Rethinking computation. In: Proceedings of the Sixth AISB Symposium on Computing and Philosophy: "The scandal of computing: What is computation?" University of Exeter, Exeter UK: 6–10.

Wolfram S. (2002) A new kind of science. Wolfram Media, Champaign IL.

Wolpert D. H. & Macready W. G. (1997) No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation 1: 67–82.