way in which (in a genuinely constructivist spirit) the agent can impose some organization on the sensed environment. In continuation of §9 we can read

> "In a constructivist approach, cognitive development must be a process of gradual complexification of the intelligence, where more abstract structures (symbolic) are built from simpler sensorimotor interactions, in a way that harmonizes the lived experiences with the constructed model."

CALM does not provide a way to build "more abstract structures… from simpler sensorimotor interactions." At best, there are the synthetic elements that contain abstract properties in the sense that they do not correspond to any sensory inputs. Given that those abstract properties are added to schemas whose context and action vectors are equal, it is impossible to understand them as abstract/symbolic structures in the sense given in §9. In Piagetian parlance, the learning exhibited by CALM could be seen as model of the *empirical abstraction* but not of the *reflective abstraction* that is crucial for cognitive development and creative behavior. Briefly, via *empirical abstraction*, some quality (e.g., weight or color or contingency among actions and qualities) is abstracted from an object/situation. On the other hand, reflective abstraction is about reorganization of existing schemas and their projection onto a higher plane. (See Kitchener 1986: 61–65 for an informative discussion of empirical and reflective abstraction, as well as the discussion in Campbell & Bickhard 1993 on *the knowing levels*).

« 4 » In §13, one can read that the agent's body with its "internal states and metabolisms, elements that belong neither to the mind nor the environment… allow the agent to have intrinsic motivations…" I believe that the decision to introduce the two entities ("body" and "mind") is somewhat arbitrary, given that it is barely mentioned in the rest of the paper. It appears that the body is introduced only to have the above-mentioned possibility to have "intrinsic motivations." If this is the case, then the intrinsic motivations can be related to low-level physiological drives (hunger, pain-avoidance) with no possibility for development of more sophisticated forms of motivations such as curiosity. If, on the other hand, the intrinsic motivations can be placed in the "mind" of the agent, I see no reason to draw the arbitrary body-mind distinction.

Georgi Stojanov obtained his PhD degree in computer science and AI from UKIM University in Skopje, Macedonia in 1997 on "Anticipation theory and electro-expectograms in the context of biological and machine intelligence." His research interests include: interactivism and constructivism; cognitive development and creativity; analogy and metaphor. He was a visiting scholar at the University of Trieste, Italy, Les Archives Jean Piaget in Geneva, Université de Versailles-Saint-Quentin-en-Yvelines in Paris, and at the Institute for Non-linear Science, UCSD. As of 2005, he is an associate professor at the American University of Paris.

Received: 13 October 2013
Accepted: 18 October 2013

Action, Anticipation, and Construction: The Cognitive Core

Mark H. Bickhard
Lehigh University, USA
mhbo/at/lehigh.edu

> Upshot • Interaction-based models of cognition force anticipatory and constructivist models. The CALM model offers significant development of such models within a machine learning framework. It is suggested that moving to an entirely interactive-based model offers still further advantages.

« 1 » Charles Sanders Peirce introduced action and interaction as the proper loci for understanding the mind well over a century ago (Joas 1993). An interaction-based model of cognition, in turn, is intrinsically anticipatory – i.e., anticipations of potential actions and interactions (Bickhard 2009b; Buisson 2004; Piaget 1954). And an action and interaction-based model of cognition forces a constructivism: it is not feasible for the world to impress competent interactive system organization into a passive mind; it must be constructed. For yet another step, given that prescience does not exist, such a constructivism must be a variation and selection constructivism, an evolutionary epistemology (Campbell 1974). These characteristics, thus, form a coherent framework for understanding cognition, and, more generally, mind (Bickhard 2009b).

« 2 » Classic passive mind models, however, descending from the ancient Greeks, still dominate the scene, currently in their "recent" incarnations of symbolic computationalism and connectionism. Machine learning is an interesting combination of perspectives: learning about the environment requires checking what is tentatively learned against that environment, which requires action and anticipation and construction of what is checked. Most cleanly, what is checked are those anticipations per se. But there is still also a reliance on passive models of perception (generally based on sensations) and restricted models of action and construction.

« 3 » Filipo Perotto's CALM is a significant advance within this framework, especially in its ability to extract anticipatory information from an only partially observable and not fully deterministic world, and to use synthetic elements in doing so. It is important to demonstrate that these more realistic framework assumptions can be handled, and to show how they can be handled.

« 4 » But CALM, too, is built on sensation models of perceiving and on singleton actions. One of the current foci for development of the CALM model is to develop possibilities of chaining schemas – again, I would agree that this is exactly the right direction. I would like to comment, however, on an even more general approach that might be considered – a fully interactive approach.

« 5 » Consider that passive sensations, insofar as they exist at all, functionally serve to detect properties of the environment, and that such detection – as a strictly factual matter – is all that is functionally relevant to the system. In particular, such detections need not be understood to *represent* that which is detected in order to account for their influences on system processing. Still further, such detections can also be realized by fully interactive processes, not just by passive receptive processes (Bickhard & Richie 1983). On the other hand, anticipations concerning possible interactions with the environment also can, and arguably do, occur with respect to *whole patterns* of interaction, not just singular actions. Chaining of schemas is precisely a step in this direction, but it requires more than "chains" to be able to model general interactive patterns.

« 6 » So, I would suggest that patterns of interaction can serve:

1 | detection functions, rather than sensations and perceptions interpreted as representational (with all of the classic problems that that interpretation entails: Bickhard 2009b),[2] *and*
2 | as patterns of interaction that are anticipated as possible in the future, *and*
3 | as patterns that can be tentatively constructed in learning more about the environment – learning more about what patterns of interaction can be anticipated as possible, given what prior interactive detections have already occurred.

« 7 » Such shifts generate a dynamic systems model, more than a classic computational model, but one in which representation is not absent:. truth value emerges in anticipations that are capable of being true or false, and cognitive representation more generally can be built from organizations of such anticipations (Piaget 1954; Bickhard 2009b). In such a model, representation is not built on or out of presumed sensations as representations.

« 8 » Modeling the dynamics of such dynamic systems is difficult. For one class of problems, there are no topological or metric spaces built in to serve as spaces for generalization. On the other hand, if the construction of such topological spaces can itself be constructed, then we can model the cognitive development and organization and re-organization of such spaces in children and adults – a higher level of learning and development than is usually addressed (Bickhard & Campbell 1996). For another class of problems, cognitive representations of, for example, objects or numbers, cannot be simply presupposed, but must themselves be constructed. However, that was among the basic insights of Piaget some time ago (Piaget 1954; Allen & Bickhard 2011, 2013a, 2013b, 2013c).

« 9 » Overall, then, moving to a fully interactive dynamic systems framework makes a number of modeling problems much more difficult. But it offers advantages of avoiding classic problems concerning, for example, the nature of representation (Bickhard 2009b), and offers direct approaches to modeling phenomena that are very difficult to approach within standard frameworks (e.g., re-organizing the topology of representational spaces in response to understanding an analogy; Bickhard & Campbell 1996).

2 | Note that this also frees the model from being able to generalize only along the dimensional variables that are built into the sensation apparatus, and from the related built-in metric spaces for error, etc. Of course, it also makes the dynamics of such generalization more difficult to model.

Mark Bickhard is the Henry R. Luce Professor in Cognitive Robotics and the Philosophy of Knowledge at Lehigh University. He is affiliated with the Departments of Philosophy and Psychology, and is Director of the Institute for Interactivist Studies. His work focuses on the nature and development of persons, as biological, psychological, and social beings.

Received: 11 October 2013
Accepted: 21 October 2013

Representing Knowledge in a Computational Constructivist Agent

Thomas Degris
ENSTA ParisTech-INRIA, France
thomas.degris/at/inria.fr

> Upshot • The aim of this commentary is to relate the target article to recent work about how to represent the knowledge acquired from experience by a constructivist agent.

« 1 » Constructivist agents acquire new knowledge and maintain existing knowledge by experimenting with their environment. A key question is then how to represent knowledge for such an agent.[3] In the target paper, knowledge that can be created and updated from data is emphasized, but a different mathematical framework and a different architecture, namely the Horde architecture (Sutton et al. 2011), is used. This commentary presents the similarities and differences between the target paper and the Horde architecture.

« 2 » Both papers focus on a situated agent embedded in its environment. The agent does not have access to the full state of the environment. To be able to understand better its interaction with the environment, the agent needs to construct abstract internal structures from a low level sensorimotor loop. In both papers, the internal representation built by the agent comes from its own experience and does not need to match an arbitrary absolute representation of its environment.

« 3 » The target paper has chosen to represent the knowledge of the agent with tree-structured representations. While trees can, in principle, take advantage of specific structures in the data, they also have issues that can make them impractical to use as a lifelong constructivist agent in the actual world. More specifically, as mentioned in the target paper, the main idea of structured representations is that the system dynamics can be factored to save memory and computational time. But such structure just may not be in the data. For instance, if one would like to predict the next value of a bump sensor on a small mobile robot, it is likely that information from all the sensors on the robot, as well as many abstract representations, may help in one way or another to make a better prediction. Thus, a prediction as simple as the value of a bump sensor may simply not be factorable. Moreover, even when some of the system dynamic may be factorable, there is no guarantee that other representations, such as value functions or policies, will be factorable (Boutilier, Dearden & Goldszmidt 2000). For an agent to take complex decision or to understand a complex environment, perhaps it is unavoidable to consider a large number of variables or signals. In comparison, the Horde architecture focuses on algorithms with a complexity in computation

3 | Joseph Modayil and I proposed an answer to that question in their presentation "Scaling-up knowledge for a cognizant robot" at the AAAI Spring Symposium on Designing Intelligent Robots: Reintegrating AI, Stanford University, 2012.