Conclusion

« 67 » The CALM mechanism can provide autonomous adaptive capabilities to an agent because it is able to construct knowledge incrementally to represent the deterministic regularities observed during its interaction with the environment, even in partially deterministic universes.

« 68 » CALM is able to deal with partially observable environments, detecting high-level regularities. The strategy is the induction and prediction of unobservable properties, represented by synthetic elements.

« 69 » Synthetic elements enable the agent to step beyond the limit of instantaneous and sensorimotor regularities. In the agent's mind, synthetic elements can represent three kinds of "unobservable things": (a) hidden properties in partially observed worlds, or sub-environment identifiers in discrete non-stationary worlds; (b) markers to necessary steps in a sequence of actions, or to different possible agent points of view; and (c), abstract properties, which do not exist properly, but which are powerful and useful tools for the agent, enabling it to organize the universe into higher levels.

« 70 » With these capabilities, CALM is able to step beyond sensorial perception, constructing more abstract terms to represent the universe and to "understand" its own reality in more complex levels. CALM can be very effective for constructing models in partially but highly deterministic ($1 > \partial \gg 0$) and partially but highly observable ($1 > \omega \gg 0$) environments, and when the transformation functions have well-structured causal dependencies ($0 < \varphi \ll n$).

« 71 » Currently, we are improving CALM to enable it to form action sequences by chaining schemas. It will allow the creation of composed actions and plans. The next research steps include: formally demonstrating the mechanism's robustness and correctness; making comparisons between CALM and related solutions proposed by other researchers; and analyzing the mechanism's performance when facing more complex problems. Future works could include the extension of CALM to deal with non-deterministic regularities, noisy environments and continuous domains.

Received: 12 July 2013
Accepted: 19 September 2013

Open Peer Commentaries

on Filipo Studzinski Perotto's "Computational Constructivist Model"

To Bridge the Gap between Sensorimotor and Higher Levels, AI Will Need Help from Psychology

Frank Guerin
University of Aberdeen, UK
f.guerin/at/abdn.ac.uk

> Upshot • Constructivist theory gives a nice high-level account of how knowledge can be autonomously developed by an agent interacting with an environment, but it fails to detail the mechanisms needed to bridge the gap between low levels of sensorimotor data and higher levels of cognition. AI workers are trying to bridge this gap, using task-specific engineering approaches, without any principled theory to guide them; they could use help from psychologists.

« 1 » The formulation of the problem as it appears in the abstract of Filipo Perotto's article packs in a lot of information that merits discussion:

> "The constructivist conception of intelligence is very powerful for explaining how cognitive development takes place. However, until now, no computational model has successfully demonstrated the underlying mechanisms necessary to realize it. In other words, the artificial intelligence (AI) community has not been able to give rise to a system that convincingly implements the principles of intelligence as postulated by constructivism, and that is also capable of dealing with complex environments."

« 2 » This suggests that the psychologists have succeeded in explaining how cognitive development takes place and that the AI community has failed in its job to implement these "principles of intelligence." However, I would throw the problem back at the psychologists. I think that significant work is still needed at the level of theoretical psychology before we have something close to a proper explanation of how cognitive development takes place. Psychological explanations are for the most part vague and woolly; they do not elucidate the mechanisms underlying development (Jean Piaget's theory being a good example). Furthermore, Piaget's theory is at times even at odds with experimental psychology. It may be many, many years before we have a suitably detailed theory from the psychologists that is consistent with the evidence from experiments. Until that time, one could argue that the "principles of intelligence as postulated by constructivism" are implemented very well by existing AI systems. Gary Drescher, for example, did implement the basic principles of constructivism, but "it could never be used to solve significant applied problems," because the techniques do not scale up to systems with large numbers of inputs and degrees of freedom. However, Piaget did not give us any idea of how to deal with these issues, so one could lay the blame on him.

« 3 » To quote from the abstract again, "there is a large distance between the descriptions of the intelligence made by constructivist theories and the mechanisms that currently exist." If we consider Piaget's theory, and Drescher's system or the CALM system, I am not sure that there is such a large distance. Piaget's descriptions of assimilation and accommodation are so all encompassing and so lacking in detail that it seems to me that Drescher's system or the CALM system constitute perfectly good implementations. Psychology tends to leave mechanisms very underspecified.

« 4 » To quote again from the article's abstract: "…and that is also capable of dealing with complex environments." Here is perhaps the essence of the problem. When you start building an actual AI system that has to interact with the world, you face a daunting task of dealing with a complex environment. It seems that AI is being saddled with the burden of not only implementing the high-level theory, but also making sure it can deal with complex environments. The "complex environments" problem needs to be thrown back at the psychologists. The history of AI has shown that a theory of cognition that works at a high abstract level but cannot account for the interface to the sensorimotor level is not much of a theory of cognition at all. The devil is in the detail. There are many writers who convincingly show how high-level cognition is very much grounded in our sensorimotor intelligence (e.g., Barsalou 2008; Byrne 2005; Bril, Roux & Dietrich 2005). Psychological theories tend to overlook the need for complex mechanisms to bridge the gap between the sensorimotor level and high-level cognition. Psychologists may need to become computer scientists to some extent, so that they have an appreciation of the computational problems involved and the need for them to describe mechanisms to account for how humans successfully solve these.

« 5 » On the positive side, there are some works in cognitive science that are beginning to attempt to address the issue of providing some theoretical framework to account for how a sensorimotor level can connect with higher levels of cognition: for example, the multi-layered cognitive system of Bipin Indurkhya (1992, Chapter 5).

« 6 » For the CALM system itself, I feel the article has all the correct ideas from a philosophical and psychological point of view, e.g., about the agent constructing its own symbolic structures and not having access to the "ontological reality." However, if we are to evaluate it as a candidate for a "general artificial intelligence mechanism that learns like humans do" (first sentence of abstract), then it might suffer the same shortcomings as Drescher's work, i.e., "it could never be used to solve significant applied problems." For example, if the context were to be the visual input from two stereo cameras delivering a few million pixels in 24 bit colour at thirty frames per second and the system is trying to predict the consequences of actions, in the complexity of an everyday setting, in this visual stream, it might not be feasible to use each bit of input as a CALM variable. One could, of course, propose to hook the CALM system up to a higher-level abstracted version of the visual input, but then one runs into the issues of where to make the cut-off between what the core CALM system sees and what is the responsibility of other abstraction mechanisms. If the cut-off is at the wrong place, then one runs into classical AI problems of (a) having a core cognition that makes unreasonable assumptions about how accurately it can interface with the world or (b) having a prespecified worldview imposed by the provided abstractions (see Brooks 1991 or Stoytchev 2009 for problems with this). There does not seem to be any clear theory from psychology to guide us on how to connect the sensorimotor level with some higher levels. AI does have various different applied systems that successfully make a connection from high-level symbols to perception and action in complex settings: for example, robots that perform everyday tasks (Beetz et al. 2010). However, each applied AI system tends to be specialised and optimised for one particular task. None could claim to be a reasonable model of general human cognition, nor do they attempt to be. This is really a job for the psychologists.

Frank Guerin obtained his Ph.D. degree from Imperial College, London, in 2002. Since August 2003, he has been a Lecturer in Computing Science at the University of Aberdeen. He is interested in understanding the core of cognition in computational terms. He has focused on understanding infant cognitive development, as a first step to understanding later cognition.

Received: 11 October 2013
Accepted: 18 October 2013

Environments Are Typically Continuous and Noisy

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

> Upshot • The schema system presented in the target article suffers from problems that had been acknowledged more than ten years ago. The main point is that our world is neither deterministic nor symbolic. Sensory as well as motor noise is ubiquitous in our environment. Symbols do not exist a priori but need to be grounded within our continuous world. In conclusion, I suggest that research on schema-learning systems should tackle small but real-world, continuous, and noisy problem domains.

Heuristic learning principles are not enough

« 1 » About 15 years ago, I began working together with Wolfgang Stolzmann and Joachim Hoffmann on the development of anticipatory classifier systems (Stolzmann 2000). We attempted to tackle the fundamental problems of learning a cognitive model in well-structured environments, implementing contextual rule differentiation, rule adjustment, and rule integration mechanisms. With iterative improvements and additions, the ACS2 system was developed. ACS2 combines a heuristic rule differ-