Commentary on: Andy Clark, Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science.
(This is a **revised** version of the commentary to appear in Behavioral and Brain Science, circa 2013)

What else can brains do?

Aaron Sloman
School of Computer Science, University of Birmingham, UK
http://www.cs.bham.ac.uk/~axs

Abstract: The approach Clark labels "action-oriented predictive processing" treats all cognition as part of a system of on-line control. This ignores other important aspects of animal, human, and robot intelligence. He contrasts it with an alleged "mainstream" approach that also ignores the depth and variety of AI/Robotic research. I don't think the theory presented is worth taking seriously as a *complete* model, even if there is much that it explains.

Clark's paper deserves far more than 1,000 words, but I have to be brief and dogmatic. Characterizing brains as predicting machines ignores many abilities produced by evolution and development,[1] including mathematical discovery and reasoning, using evolved mechanisms (perhaps) shared by several species capable of the "representational redescription" postulated in Karmiloff-Smith (1992) and the meta-configured competences suggested in Chappell & Sloman (2007), including (largely unstudied) discoveries of "toddler theorems" (Sloman 2010). The "action-oriented predictive processing" approach treats everything as on-line control (Powers 1973), like "enactivist" theorists who usually ignore competences required to make predictions true and processes generating and choosing (sometimes unconsciously) between goals, plans, designs (for houses, machines, etc.), preferences, explanations, theories, arguments, story plots, forms of representation, ontologies, grammars, and proofs. Predictive processing doesn't explain termite cathedral building. (Compare Chittka & Skorupski 2011).

Simultaneous localisation and mapping (SLAM) robotic techniques, partly inspired by things animals do, create useful (topological, metrical, and possibly logical) representations of enduring extended environments. That's not learning about mappings between inputs and outputs. It's a special case of using actions, percepts, and implicit theories to derive useful information about the environment. Another is producing a theory of chemical valency.

Systematically varying how things are squeezed, stroked, sucked, lifted, rotated, and so forth, supports learning about kinds of matter, and different spatial configurations and processes involving matter (Gibson 1966). Predicting sensory signals is only one application. Others include creating future structures and processes in the environment, and understanding processes. Choosing future actions often ignores sensory and motor details, since a different ontology is used (e.g. choosing between a holiday spent practising French and a music-making holiday, or choosing insulation for a new house). For more on "off-line" aspects of intelligence ignored by many "enactivist" and "embodied cognition" enthusiasts, see Sloman (1996; 2006; 2009). Even for on-line control, the use of servo-control with qualitative modifications of behavior responding to changing percepts reduces the need for probabilistic prediction: Head for the center of the gap, then as you get close use vision or touch to control your heading. Choosing a heading may, but need not, involve prediction: it could be a reflex action.

Predicting environmental changes need not use Bayesian inference, for example when you predict that two more chairs will ensure seats for everyone, or that the gear wheel rotating clockwise will make the one meshed with it rotate counter-clockwise. And some predictions refer to what cannot be sensed, for example most deep scientific predictions, or a prediction that a particular way of trying to prove Fermat's last theorem will fail.

Many things humans use brains for do not involve on-line intelligence, for example mulling over a conversation you had a week ago, lying supine with eyes shut composing a piano piece, trying to understand the flaw in a philosophical argument, or just daydreaming about an inter-planetary journey.

I don't deny that many cognitive processes involve mixtures of top-down, bottom-up, middle-out (etc.) influence: I helped produce a simple model of such visual processing decades ago, Popeye (Sloman 1978, Ch. 9), and criticized over-simple theories of vision that ignored requirements for process perception and on-line control (Sloman 1982; 1989). David Hogg, then my student, used 3-D prediction to reduce visual search in tracking a human walker (Hogg 1983). Sloman (2008) suggests that rapid perception of complex visual scenes requires rapid activation and instantiation of many normally dormant, previously learnt model fragment types and relationships, using constraint propagation to rapidly assemble and instantiate multi-layered percepts of structures and processes: a process of *interpretation*, not *prediction* (compare parsing). Building working models to test the ideas will be difficult, but not impossible. Constraint propagation need not use Bayesian inference. Clarke writes:

> "Thus consider a black box taking inputs from a complex external world. The box has input and output channels along which signals flow. But all it 'knows' about, in any direct sense, are the ways its own states (e.g., spike trains) flow and alter….The brain is one such black box" (sect. 1.2).

This sounds like a variant of concept empiricism, defeated long ago by Kant (1781) and buried by philosophers of science.

Many things brains and minds do, including constructing interpretations and extending their own meta-cognitive mechanisms, are not concerned merely with predicting and controlling sensory and motor signals.

Evolutionary "trails", from very simple to much more complex systems, may provide clues for a deep theory of animal cognition explaining the many layers of mechanism in more complex organisms. We need to distinguish diverse *requirements* for information processing of various sorts, and also the different *behaviors* and *mechanisms*. A notable contribution is Karmiloff-Smith (1992). Other relevant work includes McCarthy (2008) and Trehub (1991), and research by biologists on the diversity of cognition, even in very simple organisms. I have been trying to do this this sort of exploration of "design space" and "niche space" for many years (Sloman 1971; 1978; 1979; 1987; 1993; 1996; 2002; 2011a; 2011b).

Where no intermediate evolutionary steps have been found, it may be possible to learn from alternative designs on branches derived from those missing cases. We can adopt the designer stance (McCarthy 2008) to speculate about testable mechanisms. (It is a mistake to disparage all "just so" stories, since some are based on deep experience of struggling to build working systems, and can be used to guide research rather than replace it.) This project requires studying many types of environment, including not only environments with increasingly complex and varied *physical* challenges and opportunities, but also increasingly rich and varied interactions with *other information processing systems:* predators, prey, and conspecifics (young and old). Generalizing Turing (1952), I call this the "Meta-morphogenesis project" (Sloman 2013).

Clark compares the prediction "story" with "mainstream computational accounts that posit a cascade of increasingly complex feature detection (perhaps with some top-down biasing)" (sect. 5.1). This fits some AI research, but labelling it as "mainstream" and treating it as the only alternative, ignores the diversity of approaches and techniques including constraint-processing, SLAM, theorem proving, planning, case-based reasoning, natural language processing, and many more. Much human motivation, especially in young children, seems to be concerned with extensions of competences, as opposed to predicting and acting, and similar learning by exploration and experiment is being investigated in robotics.

A minor point: Binocular rivalry doesn't always lead to alternating percepts. For example look at an object with one eye, with something moving slowly up and down blocking the view from the other eye. The remote object can appear as if behind a textured window moving up and down.

Clark claims (in his abstract) that the "hierarchical prediction machine" approach "offers the best clue yet to the shape of a unified science of mind and action". But it unifies only the phenomena its proponents attend to.

End of published version

(**Not in published version:** there are examples of mathematical reasoning about geometrical structures, such as appear to have led to the development of Euclid's *Elements*, which I suspect build on previously evolved abilities to detect and reason about possibilities for change in the environment, and restrictions on such possibilities, shared with other animals that perceive and reason about affordances. Some examples of the geometrical competences, which have nothing at all to do with probabilistic prediction, but are concerned with proofs concerning what is and is not possible, are presented in a draft document "Hidden Depths of Triangle Qualia", available at:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html

This briefly explains a distinction between online intelligence and offline intelligence that is ignored by the theory that brains are prediction machines.)

NOTE

1. For more details, see http://www.cs.bham.ac.uk/research/projects/cogaff/12.html#1203

References

Chappell, J. & Sloman, A. (2007) Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing* 3(3):211–39. http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609

Chittka, L. & Skorupski, P. (2011) Information processing in miniature brains. *Proceedings of the Royal Society B* (278):885–88. doi: 10.1098/rspb.2010.2699.

Gibson, J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin.

Hogg, D. (1983) Model-based vision: A program to see a walking person. *Image and Vision Computing* 1(1):5–20.

Kant, I. (1781) *Critique of pure reason,* tr. Kemp Smith 1929, Macmillan.

Karmiloff-Smith, A. (1992) *Beyond modularity: A developmental perspective on cognitive science*. MIT Press

McCarthy, J. (2008) The well-designed child. *Artificial Intelligence* 172(18):2003–14

Powers, W. T. (1973) *Behavior, the control of perception*. Aldine de Gruyter.

Sloman, A. (1971) Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In: *Proceedings of the 2nd IJCAI,* ed D.C.Cooper*,* pp. 209–26. William Kaufmann. http://www.cs.bham.ac.uk/research/cogaff/04.html#200407.

Sloman, A. (1978) *The computer revolution in philosophy*. Harvester Press (and Humanities Press).

Sloman, A. (1979) The primacy of non-communicative language. In: *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference*, Oxford, March 1979, ed. M. MacCafferty & K. Gray, pp. 1–15. Aslib. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43.

Sloman, A. (1982) Image interpretation: The way ahead? In: *Physical and biological processing of images (Proceedings of an International Symposium* organised by The Rank Prize Funds, London, 1982, ed. O. Braddick & A. Sleigh, pp. 380–401. Springer-Verlag. http://www.cs.bham.ac.uk/research/projects/cogaff/06.html#0604

Sloman, A. (1987) Reference without causal links. In: *Advances in artificial intelligence – ii*, ed. J.B.H. du Boulay, D. Hogg, & L. Steels, pp. 369–81. North-Holland. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#5

Sloman, A. (1989) On designing a visual system (towards a Gibsonian computational model of vision). *Journal of Experimental and Theoretical AI* 1(4):289–37. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#7

Sloman, A. (1993) The mind as a control system. In: *Philosophy and the cognitive sciences*,ed. C. Hookway & D. Peterson, pp. 69–110. Cambridge University Press. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18

Sloman, A. (1996) Actual possibilities. In: *Principles of knowledge representation and reasoning: Proceedings of the 5th International Conference (KR '96)*, ed. L. Aiello & S. Shapiro, pp. 627–38. Morgan Kaufmann Publishers.

Sloman, A. (2002) Diagrams in the mind. In: *Diagrammatic representation and reasoning*, ed. M. Anderson, B. Meyer & P. Olivier, pp. 7–28. Springer-Verlag.

Sloman, A. (2006, May) *Requirements for a fully deliberative architecture (or component of an architecture)* (Research Note No. COSY-DP-0604). School of Computer Science, University of Birmingham

Sloman, A. (2008) *A multi-picture challenge for theories of vision* (Research Note No. COSY-PR-0801). School of Computer Science, University of Birmingham.

Sloman, A. (2009) Some requirements for human-like robots: Why the recent over-emphasis on embodiment has held up progress. In: *Creating brain-like intelligence*, ed. B. Sendhoff, E. Koerner, O. Sporns, H. Ritter & K. Doya, pp. 248–77. Springer-Verlag.

Sloman, A. (2010) If learning maths requires a teacher, where did the first teachers come from? In: *Proceedings of the International Symposium on Mathematical Practice and Cognition, AISB 2010 Convention*, ed. A. Pease, M. Guhe & A. Smaill, pp. 30–39. De Montfort University.

Sloman, A. (2011a) Varieties of meta-cognition in natural and artificial systems. In: *Metareasoning: Thinking about thinking*, ed. M. T. Cox & A. Raja, pp. 307–23. MIT Press.

Sloman, A. (2011b, Sep) *What's vision for, and how does it work? From Marr (and earlier) to Gibson and beyond.* (Online tutorial presentation, also at http://www.slideshare.net/asloman/).

Sloman, A. (2013) Paper 4: Virtual machinery and evolution of mind (part 3) meta-morphogenesis: Evolution of information-processing machinery. In: *Alan Turing – his work and impact*, ed. S. B. Cooper & J. van Leeuwen. Elsevier

Trehub, A. (1991) *The cognitive brain*. MIT Press. http://www.people.umass.edu/trehub/.

Turing, A. M. (1952) The chemical basis of morphogenesis. *Philosophical Transactions of Royal Society of London B* 237:37–72