Review:
The Computer Revolution in Philosophy:
Philosophy, Science and Models of Mind
1978
by Aaron Sloman

Reviewed by: Stephen P. Stich
The Philosophical Review,
Vol. 90, No. 2 (Apr., 1981), pp. 300-307

This file is at
http://www.cs.bham.ac.uk/research/projects/cogaff/crp/stich-review-crp.html
PDF version:
http://www.cs.bham.ac.uk/research/projects/cogaff/crp/stich-review-crp.pdf

---

REVIEW OF
THE COMPUTER REVOLUTION IN PHILOSOPHY: PHILOSOPHY,
SCIENCE AND MODELS OF MIND.
AARON SLOMAN.
Humanities Press, 1978. Pp. xii, 304
[Now out of print.
Slightly revised version, with additional notes added after 2001,
freely available online at
http://www.cs.bham.ac.uk/research/projects/cogaff/crp/]

The central theme of Aaron Sloman's book is that developments in computer science and the art of programming "can change our thinking about ourselves: giving us new models, metaphors and other thinking tools to aid our efforts to fathom the mysteries of the human mind-and heart" (p. x). Thus Sloman sets out to "re-interpret some age-old philosophical problems, in the light of developments in computing" (p. 5). To make his case for the revolutionary potential of computing ideas in philosophy, Sloman offers several extended examples of the ways in which computational concepts and models can clarify philosophical issues and enable us to ask deeper questions. He also illustrates the way in which a philosophical project can be recast to make it continuous with ongoing research in artificial intelligence (AI). Sloman does not pretend to offer completed accounts or theories. All of his proposals are avowedly tentative, fragmentary, and oversimplified. But this is no defect. What Sloman is proposing is an ongoing exploration of the fruits that may emerge from the hybridization of philosophy and AI. The embryonic results of this cross-fertilization are often exciting; Sloman makes a plausible case for a rich harvest ahead.

The book is notable for its abundance of intriguing asides on the possible implications of AI models for education, psychotherapy, social policy, and the arts of communication. Though these asides are sometimes a bit wild, they are often thought provoking and insightful. On the darker side, Sloman is also prone to gratuitous nastiness, unsubstantiated allegations, and simplistic "solutions" to subtle and difficult problems. Thus, for example, he rails against "academic colleagues" (unnamed) who are convinced by "fine prose, impressive looking diagrams or jargon" (p. 13). He warns us that "most psychologists never even think of the important questions, and those who do usually lack the techniques of conceptual analysis required for tackling them" (p. 37). No wonder, then,

that "so much of philosophy, psychology, and social science is vapid, or simply false" (p. 15). But fear not. In setting out the questions that have exercised philosophers, Sloman will "ignore the many pseudo-questions posed by incompetent philosophers who cannot tell the difference between profundity and obscurity" (p. 65).

In that latter category, I must surmise, are those benighted souls who puzzle over the questions of personal identity and survival after death. The issue can be dealt with definitively in a single paragraph.

> The computational metaphor, paradoxically, provides support for a claim that human decisions are not physically or physiologically determined, since, as explained above, if the mind is a computational process using the brain as a computer then it follows that the brain does not constrain the range of mental processes, any more than a computer constrains the set of algorithms that can run on it. Moreover, since the state of a computation can be frozen, and stored in some non-material medium such as a radio signal transmitted to a distant planet, and then restarted on a different computer, we see that the hitherto non-scientific hypothesis that people can survive bodily death, and be resurrected later on, acquires a new lease of life. [1]

The refutation of reductionism (p. 9) is only a few sentences longer.

Sloman’s cantankerousness is most pronounced in the first three chapters, where his principal concern is to set out his account of the nature and aims of science. That account is, by far, the weakest part of the book. The unhappy combination of abrasive tone and dubious substance may lead many readers to leave the remainder of the book unread. That would be a pity. For despite its faults, there is much here that is valuable and important. The chapters on conceptual analysis, intelligent mechanisms, analogical reasoning, arithmetic knowledge, and perception are each useful contributions to the literature. Taken together these chapters constitute an impressive defense of Sloman’s central thesis: the ideas and techniques developed in AI provide powerful new tools for tackling philosophical problems. In the paragraphs that follow I will elaborate a bit on this theme, then sketch my misgivings about Sloman’s account of science.

In Chapter Four, "What Is a Conceptual Analysis?" Sloman endorses the orthodox Oxbridge line on the nature and importance of conceptual analysis.

> We have a rich and subtle collection of concepts for talking about mental states and processes and social interactions.... These have evolved over thousands of years, and they are learnt and tested by individuals in the course of putting them into practical use.... All concepts are theory-laden, and the same is true of these concepts. In using them we are unwittingly making use of elaborate theories about language, mind and society. The concepts could not be used so successfully in intricate interpersonal processes if they were not based on substantially true theories. So by analyzing the concepts, we may hope to learn a great deal about the human mind and about our own society. [84-85]

What is unique, and delightful, in Sloman’s chapter is that he proceeds to give a step by step account of how to embark on a conceptual analysis -- a how-to-do-it guide compiled with a programmer’s eye to detail. We start out by collecting varied instances and noninstances of the concept in question. Dictionaries and Roget’s Thesaurus are then consulted for tentative definitions and for lists of related words and phrases. This is followed by a variety of probes designed to illuminate features of the concept being explored or the commonsense theory in which it functions: "Ask what the role of the concept is in our culture (p. 91).... Ask what sort of things can be explained by instances

of the concept (p. 91).... Often some question about the analysis of a concept can be investigated by telling elaborate stories about imaginary situations" (p. 95). For the student, or the newcomer to the art, Sloman's manual is an invaluable guide. And even the practiced hand is bound to find useful hints here. The least familiar suggestion in Sloman's recipe for conceptual analysis is the one he saves for last.

> Try to test your theories by expressing them in some sort of computer program or at least in a sketch for a design of a working program.... Test your analysis by designing a program whose behavior is intended to instantiate the concept, then see whether the actual behavior is aptly described using the concepts in question. You will usually find that you have failed to capture some of the richness of the concept.... The methods of A. I. provide a useful extension to previous techniques of conceptual analysis, by exposing unnoticed gaps in a theory and by permitting thorough and rapid testing of very complex analyses. [97]

In Chapter Six, "Sketch of an Intelligent Mechanism," Sloman takes up his own suggestion. He sketches the design of a program aimed at simulating intelligence, "purposiveness, flexibility, and creativity" (p. 116). What, Sloman asks, would a computer-driven robot have to do for us comfortably to describe it as intelligent, flexible, creative, etc.? And what sort of program could possibly generate such activity? The answers, needless to say, are exceptionally complex. To count as intelligent and purposeful a robot must be capable of doing many different sorts of things. It must form plans and subplans; it must learn from experience how to produce better plans; it must have a reasonable store of information about itself and its environment; it must monitor its environment and update its information store efficiently; it must decide quickly among varying courses of action, often on the basis of incomplete information, etc. What is more, all of these various capacities must integrate properly with each other. Because of the richness and complexity of the intuitive theory in which our concepts are embedded, it is almost unavoidable that a detailed account of our commonsense theory be cast as a computer program, since programming languages provide the only available formalism for representing complex interacting processes. Sloman's sketch of an intelligent mechanism can be viewed with equal justice as an unusually detailed effort at conceptual analysis, or as a rather sketchy outline for an ambitious AI research project. The ambiguity is a persuasive argument for Sloman's contention that the philosopher's project and the AI programmer's project are continuous with each other.

In the chapter on "Perception as a Computational Process," Sloman sets out to recast in computational terms the Kantian claim that perception presupposes prior knowledge and abilities. The project described is billed as an "attempt to design a machine which can see" (p. 217). The idea is to program a computer to recognize various objects or patterns scanned by a TV camera (or presented in some other preceded, but essentially equivalent way). Humans are remarkably good at recognizing shapes, letters, and objects even when the visual input is ambiguous, unclear, distorted, or staggeringly complex. And simulating the human -achievement turns out to be one of the more difficult tasks tackled by the artificial intelligencia. The most successful of currently available programs take an essentially Kantian view of the project. "The program has to work up the raw material by comparing representations, combining them, separating them, classifying them, describing their relationships, and so on. What Kant failed to do was describe such processes in detail" (p. 230). It is Sloman's contention that the best way to elaborate on the workings of Kantian schemata is to construct programs capable of simulating human perceptual capacities. The brief sketch he gives of his own POPEYE project is enough to make it plausible that AI offers a promising new technique for exploring the ways in which preexisting knowledge, theories, and concepts interact in the process of perception.

One of the lessons to emerge from work on computer vision is how very complex

perception is, and how little of the complexity is introspectively available. It turns out "that very complex computational processes are required for what appeared previously to be very simple abilities, like seeing a block, or even seeing a straight line" (pp. 219-20). The case for complex cognitive processes underlying phenomenologically simple perception is, i think, quite undeniable. What is more dubious is Sloman's contention that unconscious cognitive processes are "essentially similar in character to intellectual processes of which we are sometimes conscious" (p. 224). It is probably true that much early work in AI began with a bias in favor of postulating unconscious processes modeled on rational conscious processes. But surely there is no reason a priori to assume that all unconscious cognitive processes are "essentially similar in character" to familiar conscious processes.[1]

---

[1] For an elaboration on this theme, cf. my "Between Chomskian Rationalism and Popperian Empiricism," British Journal for the Philosophy of Science, 30, 4.

Let me turn, now, to Sloman's account of the aims and methods of science, which occupies the bulk of Chapters Two and Three. It is a bit anomalous that a discussion of scientific methodology should occupy so prominent a place in a volume whose main focus is elsewhere. Sloman is obviously aware of the anomaly, but makes little effort to explain it away. It is just that "the issues are generally misunderstood, and I felt something needed to be done about that" (p. xi). But I think a bit of reading between the lines reveals a more compelling motive. AI, after all, is a puzzling discipline. Unlike an empirical science, it does not seem to aim at explaining the workings of some part of nature. Consider, for example, those paradigms of AI virtuosity, the chess playing computer programs. It is possible to inquire into how people actually go about playing chess; how they formulate strategies, select among them, recognize impending attacks, etc. The answers, no doubt, would involve cognitive mechanisms of which the players are at best dimly aware. And there would be different answers for different players with different levels of skill. But none of this story need be of much interest to the AI researcher whose goal is to produce a winning program. The strategies of chess masters may, of course, provide the AI researcher with some useful ideas. But he is free to adopt them or pass them up. It is not the point of an AI chess program to explain how people play chess. If, however, AI is not the empirical study of human cognitive processes, what is it?

A number of writers have noted that AI is not concerned with how people accomplish a task, but rather with the question of how it is possible for any physical mechanism to execute some task requiring intelligence. It is surely of considerable philosophical importance to show that a task or activity hitherto accomplished only by creatures with minds can be accomplished also by artifacts, physical through and through. And of course, it may be of some technological importance to build machines which can recognize patterns, play chess, transcribe the spoken word, etc. Should we conclude, then, that AI is not a science at all, but rather a curious hybrid of philosophy and technology?

Sloman's discussion of science is, I think, largely motivated by his wish to avoid such a conclusion and to argue instead that AI falls squarely within the realm of science. His strategy is a heroic one. Rather than show that AI really does share what are generally taken to be the aims and methods of empirical science, he maintains that empirical science, contrary to common misconception, shares the aims and methods of AI. According to Sloman, a principal aim of science is to extend "knowledge of what sort of things are possible and impossible in the world, and how or why they are . . . " (p. 24). One of the subgoals into which this broad aim is divided is the "constructing [of] theories to explain known possibilities" (p. 27, emphasis his). A theory "of the

constituents of atoms" is offered as an example of such a theory explaining possibilities. Generative grammars are a second example. And "artificial intelligence models provide a major new species of explanations of possibilities" (p. 27). However, Sloman does not pause to subject his notion of possibility to the sort of conceptual analysis he advocates elsewhere. This is unfortunate since it seems that Sloman is using the term 'possibility' in a bewildering variety of senses; what is worse, he often uses the term in contexts where it is hard to believe it means anything at all. He tells us that " 'pure' science first discovers instances of possibilities then creates explanations of those possibilities" (p. 32). He then goes on to list examples. But in many of the examples the term 'possibility' seems to be idling-doing no work at all. Consider:

> Newton's gravitational theory explained how it was possible for the moon to produce tides on earth. His theory of the relation between force and acceleration explained how it was possible for water to remain in a bucket swung overhead. [46]

Now I should have thought that Newton's gravitational theory explained how the moon actually does produce tides on earth, and that his theory of the relation between force and acceleration explained why water remained in a bucket swung overhead. If, as I suspect, Sloman intends his claims to be equivalent to these, then the talk about possibilities is quite empty. If he does not intend the two to be taken as equivalent, then explaining how it is possible for the moon to produce tides on earth must be something different from explaining how the moon actually does produce tides on earth. But, though I can conjure a variety of nonvacuous readings for 'explaining how it is possible for the moon to produce tides on earth', none of them make it at all plausible that Newton's theory explained any such thing. Here are some of Sloman's other examples:

> The kinetic theory of heat explained, among other things, how it was possible for heating to produce expansion, and how heat energy and mechanical energy could be interconvertible. [46]

> The theory of genes explained how it was possible for offspring to inherit some but not all of the characteristics of each parent, and for different siblings to inherit different combinations. [46]

In these cases too, the only reading on which the claims are plausible is one that takes the modal locutions as idiosyncratic paraphrases of more familiar indicatives: the kinetic theory of heat explained why heating produces expansion, and the theory of genes explained why offspring inherit some but not all of the characteristics of each parent. I conclude that Sloman's attempt to show that natural science aims at explaining possibilities, and thus that AI is of a piece with the rest of natural science, does not succeed. Whatever problems there are about the status of AI and its relations to other disciplines are problems that remain to be solved.

Let me add a final complaint. Sloman's book must surely mark a low point in the book manufacturer's art. The right margins are not justified, and the book abounds with misprints. The typeface frequently changes to smaller print, and occasionally to boldface, for no evident reason. Emphasis is sometimes indicated by italics, sometimes by boldface, and at least once by underlining.

All of this grousing should not be misconstrued, however. Despite its faults of form and content, Sloman's book is a useful and important one. The vision he offers of a merger between philosophy and AI is exciting, and I would predict that a growing number of philosophers will follow the path Sloman has helped to forge.

STEPHEN P. STICH

Installed here 19 Nov 2014
With the permission of Stephen Stich
by Aaron Sloman
School of Computer Science
The University of Birmingham

I accept all of the criticisms of the style of the book, but have a partial response to the comments about explanations of possibility, here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html