International Journal of Philosophical Studies
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/riph20

Mind and Machine

Cathal O'Madagain[a]
[a] Institut Jean Nicod, Paris
Published online: 08 Apr 2014.

Routledge
Taylor & Francis Group

Book Review

Mind and Machine
By Joel Walmsley
Palgrave Macmillan (Palgrave Philosophy Today), 2012. Pp. xiv + 176. ISBN: 978-0-230-30294-5. $29.00 (pbk), $85.00 (hbk).

A centipede was happy quite
Until a toad in fun said
"Pray, which leg moves after which?"
This raised her doubts to such a pitch
She fell exhausted in a ditch
Not knowing how to run.

Katherine Craster, *The Centipede's Dilemma*, 1871

One insight about the mind that has been drawn from this rhyme is called 'Humphrey's Law': that for many of our learned skills – running, playing the banjo – we don't have to think in order to execute them, and that if we were to think about them, we'd probably mess them up (G. Humphrey, 1923, *The Story of Man's Mind*, Boston: Small, Maynard and Company). But the rhyme also highlights that thinking isn't something we are generally inclined to attribute to creatures whose skills include even such complex maneuvers as a centipede's walk. Such skills, we might expect, can be acquired without the centipede thinking at all. It might be surprising, then, to learn that some of the most recent developments in our attempts to understand the mind have focused on such mundane activities as an insect's ability to walk, or our ability to wag our fingers in synchrony. It is these developments that are the ultimate focus of Walmsley's *Mind and Machine*, which explores the subject of artificial intelligence, and defends the prospects of one recent approach to AI – the Dynamic Systems approach.

The first chapters of the book take us on a tour through the history of attempts to understand the mind in terms of principles of 'brute force' – the kind of principles that can be used to exhaustively describe the working of a mill, or an internal combustion engine. Pascal's *Pascaline*, a mechanical calculator he built in 1642 to help in his father's work as a tax collector, and Leibniz's *Stepped Reckoner* are both described in detail. The significance of these machines, which replicate the logical operations of addition, subtraction, multiplication and division by procedures of brute force, is made clear when considered in light of our current understanding of some higher-level cognitive achievements, for example the ability to speak a language. On Chomsky's view, language acquisition consists of the application of a set of logical procedures on the part of the language-learner to the stimulus input of a spoken language. If logical procedures can be mechanized, as these machines seem to show, then Descartes' famous contention that a machine could never acquire a language – the hallmark of human intelligence – is undermined.

The crowning achievement in the history of such inventions is of course Turing's machine – the 'Logical Computing Machine'. Turing introduced us to the notion of universal programmability – the possibility that a machine could be built that is not only capable of implementing certain logical procedures such as addition or subtraction, but that can be programmed to implement any algorithmic procedure whatsoever, where such a procedure is one that consists of a finite number of simple mechanical steps. Apart from laying the groundwork for the invention of computers as we now know them, Turing thereby gave us a precise way of asking the question whether a mind is a machine: once we have the definition of a machine that can carry out any algorithmic procedure, we can ask whether the mind might be a machine by asking if, for anything the mind can do, a Turing machine could do it too.

With this historical framework of AI in place, Walmsley takes us into the contemporary debate by exploring problems a Turing machine faces when it attempts to account for the operations that minds can accomplish, and the developments in our understanding of what kind of mechanisms may instantiate cognition, where these mechanisms seem to depart in various ways from the principles of a Turing machine. The first five chapters of the book are largely expository, being an introduction to the field. With respect to the problems facing Turing machines, this means exploring the problem raised by Gödel's incompleteness theorem, the possible shortcomings of Turing's test for intelligence, and Searle's Chinese Room argument, all with the most recent replies made on behalf of the advocate of artificial intelligence. With respect to alternative models of what kind of mechanisms might instantiate a mind, this means exploring connectionist or neural network models of the mind, along with the merits of such models and their shortfalls. Particularly helpful throughout is the balanced discussion of representationalism – the idea that states of minds can be understood as representations of the world outside the mind. Rather than understating the importance of representations, as proponents of connectionist (and particularly enactivist) theories of cognition sometimes seem inclined to do, or rejecting models of the mind that less obviously accommodate representations as commentators at the other end of the spectrum tend to do, Walmsley steers a careful path that highlights the possibility of retaining the insights of a representational theory of thought while embracing the advantages of less explicitly-representationalist models of cognition.

Where the book clearly moves beyond exposition is in chapter six, which is a defense of the claim that Dynamic Systems Theory (DST) offers an indispensable supplement to our account of the mind. While a Turing machine or a connectionist network models cognition as occurring in stages – an ordered series of events describable by an algorithm, for example – neither takes the time in which such processes occur (as opposed to the order of events) to form an essential component of the process itself. In contrast, DST takes the role of time seriously, including temporal variables in its models of cognitive processes. Following van Gelder and others, Walmsley takes this contrast to mark the first real departure from the principles of classical Turing-Machine-style AI. Points that the DST approach has scored include predicting the idiosyncracies of systems attempting to learn or execute various activities – such as finger wagging in humans, and walking in artificial insects. Conceptual tools such as state-space models that plot the trajectories of multiple systems interacting over time – a sensory feedback network and a motor-control network in an artificial insect attempting to walk – allow us to see differences in the outcome of different types of interaction between the networks over time, and the effects these differences have on skill acquisition.

This takes us back, however, to the question that I opened with. Even if models that include time as a variable give us a better handle on predicting the kind of learning processes that allow us to engage in such activities as walking and finger-wagging, should we expect that this is giving us a grip on the mind, the seat of thought, when we feel little inclination to attribute thought to creatures like a centipede? The answer given by the DST theorist comes under the label 'incrementalism', with the neat tagline 'today the earwig, tomorrow, man' (D. Kirsch, 1991, 'Today the Earwig, Tomorrow Man?', *Artificial Intelligence* 47: 161-184). Walmsley defends the incrementalist line, offering three reasons for thinking that full-blown thought may have developed out of the tools of simpler skills. First, the ability to deftly react in real time to a changing environment took far longer to evolve than the ability to speak or play chess, so we have a good reason to think that the former is the more difficult to accomplish, and that the latter might fall into place quickly once the tools for the former have been established (cf. R. Brooks, 1991, 'Intelligence without Representation', *Artificial Intelligence* 47: 139-159). Second, in light of the first, we might be wrong to think that it is skills like playing chess rather than skills like navigating a changing environment that are the real achievements of cognition, and hence the ultimate explananda of cognitive science. And finally, although the landmark achievements of DST to date have focused on skills such as walking, DST has also been effectively employed to model higher processes such as decision-making and perceptual categorization, so that the assumption that it need only tell us about 'low-level' skills is false.

Allowing, then, that DST may be of use in telling us about the full range of cognitive achievements, a further question remains. The strong claim of this part of the book is not just that DST is a helpful tool in allowing us to understand cognition, after all, but that it's essential – because cognitive processes are essentially temporal processes. At a first pass, this claim seems worryingly close to the claim made by the theorist who identifies cognitive processes with the material processes they are instantiated in. Such identifications, such as the identification of pain with C-fibres firing, have faced a serious difficulty since the functionalist pointed out that we have good reason to expect that creatures without C-fibres could experience pain, and that identifying mental types with the meaty types they happen to be instantiated in for us humans invites the charge of a kind of meat-chauvinism. The functionalist alternative to such identity theories holds that it is not necessary to be a particular kind of material thing to instantiate a particular cognitive type. And this introduces the logical possibility that the same cognitive types that are realized in material states in us might in principle be realized in non-material entities too (even if we don't think there are any such entities). If this is a logical possibility, then material instantiation cannot be essential to the identity of a cognitive type, and cannot be included in its individuation conditions.

Similarly, although we humans have cognitive processes that are drawn out over time, it may be problematic to suppose that these time-stretches are essential to the processes being the cognitive types they are. Walmsley carefully discusses the functionalist status of the DST model of cognition, but I'm not sure that his account avoids this worry. He suggests that if the equations that characterize cognitive and neural dynamics had the same form (as described by DST), then:

> We could say that the dynamics of cognition had been (functionally) reduced to the dynamics of the brain, whilst allowing that the dynamics of cognition could be implemented by some other (perhaps non-biological) kind of concrete system, so long as its components were arranged in such a way as to fill the requisite functional role (p. 150).

But what this degree of multiple-realizability doesn't allow for is the possibility that the same cognitive type – such as representing the proposition '1 + 1 = 2' – could be realized in a dynamic or a non-dynamic system. Although the focus of DST is on cognitive types that respond to changing environments, as is repeatedly stressed, it is far from obvious that all cognitive types are concerned with dealing with a changing environment. Representing the proposition just considered – a necessary or so-called 'timeless' truth – certainly doesn't seem to be. What Walmsley's position seems to rule out is the possibility that some cognitive types might be at least in principle instantiated in non-durational entities (even if we don't think there are any) or instantaneously within a durational entity (which seems at least easier to conceive) – either way entailing that it is not essential to those cognitive types that they unfold over time. Even the claim that the same cognitive type could be realized over substantially different time-scales seems to be potentially ruled out when we take the time-scale of a cognitive type to be essential to its individuation. For example, although highly contentious, there is a growing interest in the possibility that plants might be capable of functionally realizing cognitive processes – even though these processes would take place on a vastly different time-scale than their type-identical analogues in humans (M. Pollan, 'The Intelligent Plant', *The New Yorker*, December 23, 2013). Taking DST at its word seems to rule such possibilities out a priori, which may be too quick. While the mind-brain identity theorist invites the charge of matter-chauvinism, then, the DST theorist seems to invite the charge of time-chauvinism.

Overall, Walmsley's book is challenging and illuminating, whether or not we agree with its final position. The book contributes substantially to the current debate on the way forward for AI, and will also work well as an introduction to the area for advanced undergraduate or graduate classes. Finally, the detailed accounts of the theoretical principles invoked in contemporary AI will be particularly welcomed by those who hope to keep their philosophy of mind up to date with the latest work in cognitive science.

Institut Jean Nicod, Paris *Cathal O'Madagain*

http://dx.doi.org/10.1080/09672559.2014.896616