A neurocomputational jeremiad

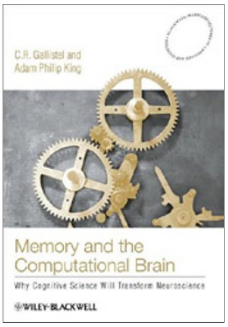

Memory and the Computational Brain: Why Cognitive Science will Transform Neuroscience

by C R Gallistel & Adam Philip King

Wiley-Blackwell, 2009
336 pp, hardcover, $99.95
ISBN 1405122889

Reviewed by Peter Dayan

Along with a light complement of fascinating psychological case studies of representations of space and time, and a heavy set of polemical sideswipes at neuroscientists and their hapless computational fellow travelers, this book has the simple goal of persuading us of the importance of a particular information processing mechanism that it claims does not currently occupy center stage. The authors maintain that "there must be an addressable read/write memory mechanism in brains that encodes information received by the brain into symbols (writes), locates the information when needed (addresses) and transports it to computational machinery that makes productive use of the information (reads)". Most of the chapters are devoted to unpacking this statement, describing conventional computer science notions of representations, symbols, information processing and Turing machines. The authors stress that neuroscientists completely ignore the issue of addressable read/write memory and/or propose preposterously inadequate solutions. The book coyly forswears a solution of its own, bar the Conradian possibility that something close to the molecular heart of neurons might be involved.

Some issues the book brings to center stage are spot on, notably representation: how indeed can the brain realize complex cognitive entities, sentences or even just visual scenes. The book contains good discussions about important notions such as productivity, compositionality and systematicity that were the focus of previous debates about connectionist representations, thus providing insights into how computers can be engineered and programmed to cope. However, such concerns nearly get buried by a reluctance to seriously consider the possibility that neural realizations of representation and computation might look nothing like engineered solutions. Indeed, there is altogether somewhat little regard for the structural and physiological facts of the brain.

Central billing in the book's conception of what is missing in contemporary neuroscience goes to two rather different computer architectures: the conceptually important Turing machine and something like a conventional computer, which is of more practical concern. Turing machines have two key components, a finite state machine, instantiating if-then-rules associated with transitions between states, and an infinitely long tape storing symbols (for example, just 0 and 1), which can be read from or written to on the finite state machine's orders. Among the hardware differences between a universal Turing machine and my desktop computer are an addressable memory and finiteness. The former makes it very easy to construct powerful data structures such as lists and trees, which in turn make it simple to write programs to solve the representational and processing puzzles posed by the book's psychological examples. The book claims that neuroscience is ignoring symbols and addressing, for instance, maintaining that information located in synapses is somehow too inaccessible and implicit. It also suggests that Turing tape is missing, deriding the two obvious neural possibilities of plastic synapses and reverbatory activity.

For the issue of addressing, the book really hoists itself on its own petard of Turing universalism. After all, a perfectly good Turing machine can also lack addressing; the tape is just read step by step without explicit (for example, numbered) locations. However, it can nevertheless simulate my computer's addressing schemes. It seems no less plausible to interrogate a synapse via presynaptic activity than to traverse reels of tape step by step, in both cases to get to read their states. Equally, the state of an elevated or depressed synaptic conductance seems no more or less able to realize a symbol when embedded in an appropriate computational milieu than would be the state of a tape entry being 1 or 0. When, at the end of the book, the authors finally accept something like this functionalist point, the overall thesis somewhat unravels.

The book suffers similar architectural blinders in suggesting that there is a fundamental difference between a Turing machine with a finite tape and a finite state machine, with the latter facing nasty combinatorial explosions in representing information. This is not true; my computer can actually be considered as a huge finite state machine with a very particular structure, in the form of stringent restrictions on the possible transitions. It is facile to demand a transparent mapping of "silicone" (sic) concepts onto biological ones; just because it doesn't look or walk like a duck doesn't mean that it can't realize the quack. Finally, we can come back to the odd parts of the phrase in the quote above "locates ... and transports [information]"; this is wedded to a deeply conventional notion that there is 'dumb' peripheral memory (or tape) and a 'smart' centralized computational device. In comparison, from the relative uniformity of cortical architecture, the brain looks as if computational power and storage are generally colocated and broadly distributed; there is no reason to expect any transportation in any conventional sense.

How about the mechanistic realization of something equivalent to tape? Some of the specific systems the book describes (for example, desert ants doing dead reckoning across featureless desert rocks) present compelling biophysical puzzles, but are not persuasive about the sort of general computations we expect to be enabled by tape. Outside the temporal window of reverbatory memory, we would mostly require the realization of one-shot or snapshot storage and recall. This point is not lost on neuroscientists, but is rather the intense focus of an impressive and vocal array of them.

Students of Gallistel's influential previous books, *The Organization of Action* and *The Organization of Learning*, had been eagerly awaiting an *Organization of Computation*. It's not clear that this one is quite yet ready for writing to tape. ■

Peter Dayan is in the Gatsby Computational Neuroscience Unit, University College London, London, UK.
e-mail: dayan@gatsby.ucl.ac.uk