Pinker, S. (1997). *How the Mind Works*, New York: Norton.
Turing, A. (1950). 'Computing Machinery and Intelligence', *Mind* 59, pp. 433–460.
Turing, A. (1936). 'On computable numbers', *Proc. Lond. Math. Soc.* 2nd Series, 42, pp. 230–265.

YORICK WILKS
Department of Computer Science
University of Sheffield
Sheffield, S11 4DP, UK
E-mail: yorick@dcs.shef.ac.uk

John von Neumann, *The Computer and the Brain*, 2nd edition, Mrs. Hepsa Ely Silliman Memorial Lectures, New Haven: Yale University Press, 2000, xxviii + 82 pp., $9.95 (paper), ISBN 0-300-084373-0.

When John von Neumann turned his interest to computers, he was one of the leading mathematicians of his time. In the 1940s, he helped design two of the first stored-program digital electronic computers. He authored reports explaining the functional organization of modern computers for the first time, thereby influencing their construction worldwide (von Neumann, 1945; Burks et al., 1946). In the first of these reports, von Neumann described the computer as analogous to a brain, with an input "organ" (analogous to sensory neurons), a memory, an arithmetical and a logical "organ" (analogous to associative neurons), and an output "organ" (analogous to motor neurons).

His experience with computers convinced him that brains and computers, both having to do with the processing of information, should be studied by a new discipline–automata theory. In fact, according to von Neumann, automata theory would cover not only computers and brains, but also any biological or artificial systems that dealt with information and control, including robots and genes. Von Neumann never formulated a full-blown mathematical theory of automata, but he wrote several important exploratory papers (von Neumann, 1951, 1956, 1966). Meanwhile, besides designing hardware, he developed some of the first programs, programming languages, programming techniques, and numerical methods for solving mathematical problems using computers. (Much of his work on computing is reprinted in Aspray and Burks, 1987.) Shortly before his death in 1956, he wrote an informal synthesis of his views about brains. Though von Neumann left his manuscript sketchy and unfinished, Yale University Press published it as *The Computer and the Brain* in 1958. The 2000 reprint of this small but informative book is an opportunity to learn, or be reminded of, von Neumann's thoughts on the computational organization of the mind-brain.

Von Neumann began by explaining computers, which for him were essentially number crunchers: to compute was "to operate on ... numbers according to a predetermined plan" (p. 3; all page references are to von Neumann's book). In any computer, some components (input "organs") received numerical data from the environment while other components (output "organs") delivered the calculations' results. The most important components, which he called "active organs," were those that performed operations inside the machine. Each active organ performed a particular operation, which could be arithmetical (e.g., addition, subtraction, multiplication, or division), logical (e.g., conjunction, disjunction, or negation), or physical (e.g., restoring voltage levels in the machine or amplifying signals). Numbers could be "represented" in the machine, and arithmetical operations could be carried out, in two ways. The first was analog: numbers were represented by continuous physical quantities (for example, an angle of rotation or a voltage level), and operations on numbers were performed by physically manipulating those quantities. Machines that represented and manipulated numbers in this way were called analog computers. The second was digital: numbers were represented by discrete "markers" (typically, the presence or absence of an electric pulse), and operations on numbers were performed by combining markers according to logical rules (that is, when an active organ received certain pulses as inputs, it yielded certain other pulses as outputs). Machines that represented and manipulated numbers in this way were called digital computers. By hooking active organs together in appropriate ways, one could construct an arithmetic unit, which typically carried out the four arithmetical operations.

There remained the problem of logical control, i.e., determining which operations the arithmetic unit must perform on what numbers in what order. To get the machine to solve a particular problem, namely to *program* the machine, the user must do two things. First, she must reduce the solution of the problem to a series of arithmetical operations of the four basic kinds (or whatever kinds were basic in her computer); second, she must set up the machine so that it generated the desired sequence of operations. The main innovation of modern, stored-program computers over their predecessors was that all the information required for logical control (namely, the program) could be stored inside the computer in the form of numerical data, which represented the operations to be performed. This generated the need for an extensive memory internal to the computer, to store not only programs, but also data for computing and partial results. In summary, von Neumann viewed the computer, either digital or analog, as a set of devices that performed arithmetical operations on numerical data in an order determined by logical control.

Von Neumann's explanation of computers, especially digital computers, contains a puzzling aspect. For many practical purposes, the best way to interpret the markers manipulated by computers may well be that they represent numbers. However, it is users who read markers in this way: no physical property of the machine necessitates this interpretation. The markers can be interpreted as representing things other than numbers, or nothing at all; whether users read markers as representing numbers or anything else makes no difference to the functioning of the machine. Indeed, digital computers are often described not as number crunchers but *symbol* crunchers, meaning that the markers are symbols that can represent a variety of things. According to this variant picture, computers do manipulate symbols according to formal rules, but the rules need not be interpreted as arithmetical operations, and the symbols need not be seen as numerals. Surely von Neumann, eminent logician that he was, knew the distinction between symbolic systems and their interpretations perfectly well, but for some reason he did not discuss it in this context. Perhaps this is because he was aware that numbers themselves could be seen as representing other entities, such as physical quantities or instructions for the computer. Be that as it may, von Neumann's comparison between artificial computers and brains was based on the notion that computers manipulated numbers, not symbols.

When it comes to comparing computers and brains, von Neumann was very cautious. He said he didn't have a theory of automata yet, and he was speculating about fruitful lines of attack for a mathematical understanding of the brain. But regardless of his caution about the details of brain function, von Neumann appears sure of one thing: brains compute.

To say that brains deal with information is a truism. But to maintain that brains are computers in von Neumann's sense—that is, mechanisms for performing arithmetical operations on numerical data in accordance with previously specified programs—sounds contentious. Von Neumann's comparison of computers and brains was couched in terms of how they transmit numerical data, how they perform arithmetical operations, and how their logical controls operate. This is a comparison of different *computational* organizations—it only makes sense under the assumption that brains compute. What's interesting is that von Neumann, in this book or in his previous work on automata theory, offered no justification for this assumption. Some readers won't be troubled—in their foreword to the new edition, Paul and Patricia Churchland take it as "*obvious* that ... the brain somehow manages to perform a great variety of sophisticated *computations*" (p. xviii, emphasis added). Other readers, those who don't find it obvious that brains compute, will be disappointed. They may feel that discussing the motivation for this assumption was particularly warranted by the nature of the comparison. In the case of artificial computers, their users are free to interpret the markers in the machines as numbers and their manipulations as arithmetical operations; after all, they *designed* them to fit that interpretation. So, saying that artificial computers manipulate numbers involves no metaphysical mystery. But in the case of brains, *prima facie* it's unclear what it means that—in the absence of users who interpret their processes—brains are said to perform arithmetical operations on numerical data. Over the last two decades, the resulting theoretical question, namely the question of how cognitive or neural states can have semantic content, and what content they have, spurred a huge literature in philosophy of mind. Writing in the 1950s, von Neumann didn't clarify what he had in mind, and many scientists adopted his language with the same lightheartedness.

While taking for granted that brains computed, von Neumann devoted his attention to *how* they did this. The first question was whether brains were digital or analog, which was hotly debated at the cybernetics conferences organized by von Neumann and others between 1946 and 1953 (Heims, 1991). In *The Computer and the Brain*, von Neumann laid out the digital position: given the all-or-none character of neural impulses, and given that neurons were the basic "active organs" from which neural circuitry was built, von Neumann concluded that the brain was fundamentally digital. (However, he pointed out that idealizations were involved in the digital position; he also stated that brains participated in "mixed" processes that involved digital as well as analog elements, such as muscle contractions and homeostasis, and he didn't rule out that some analog electrical phenomena might influence neural computations.) Moreover, von Neumann compared the size, speed, and energy consumption of neurons to those of vacuum tubes, which were the components of electronic computers at the time. According to von Neumann, the result was that neurons were much slower but much more efficient than vacuum tubes. In addition, he noted, a brain contained many more neurons than a computer contained vacuum tubes.

Then, von Neumann used his quantitative comparisons to reach general conclusions about neural computing. In going over his arguments, readers not committed to the computational theory of the mind-brain may find reasons for skepticism. Here are two examples.

First, in his experience with constructing artificial computers, von Neumann carefully considered the pros and cons of serial and parallel architectures (Aspray, 1990). He realized that parallel machines, where many operations were executed at the same time, were very difficult to program, and that many interesting problems—as far as anyone saw—could be solved only by a serial sequence of operations, each of which needed intermediate results obtained by previous operations. Hence, in building computers, he settled for a (mostly) serial architecture. When comparing computers and brains, though, von Neumann argued that neural components were too slow to support efficient serial computations, and that brains did many things at once anyway. So brains should be parallel computers. However, he pointed out the severe difficulty, or in some cases the impossibility, of carrying out computations in parallel. (Contemporary supporters of the view that brains are parallel computers rarely discuss this problem.)

Assuming von Neumann's premises, skeptical readers could take this as evidence that brains are not computers after all. If brains are too slow to be serial, and interesting computations are difficult or impossible to be done in parallel, doubters might consider the possibility that whatever brains do, and surely they do many things at once, isn't computing. Not von Neumann. Holding on to his computational view, he concluded, rather vaguely: "[T]he logical approach and structure in natural automata may be expected to differ widely from those in artificial automata" (p. 52).

The other example comes from the level of precision needed for effective digital computation. Again, given his extensive experience with solving mathematical problems with computers, von Neumann knew that high precision was needed. For many important applications, numbers needed to be represented with 10–12 decimals, or else the calculations' results would be spoiled by errors. If the brain were an ordinary digital computer, working with a digital notational system, this precision requirement would be easy to satisfy. But von Neumann argued that, although neurons were digital organs, brains employed a different system of notation from that of artificial computers. "[T]he nervous system," wrote von Neumann, "transmits numerical data . . . by periodic or nearly periodic trains of pulses . . . [U]nder these conditions, precisions like the ones mentioned above (10–12 decimals) are . . . out of question . . . [O]nly precision levels of 2–3 decimals are possible" (pp. 76–77). Thus, according to this estimate, there was a big gap between the precision needed to perform effective digital computations and that offered by the brain's notational system. To readers who don't share von Neumann's computational assumption, this gap might suggest that brains aren't digital *computers* after all. But von Neumann's conclusion is the opposite: "The nervous system is a computing machine which manages to do its exceedingly complicated work on a rather low level of precision . . . This fact must be emphasized again and again because no known computing machine can operate reliably and significantly on such low precision level" (p. 77). He offered no suggestions about how the brain might accomplish this feat.

On the basis of these considerations, von Neumann presented further hypotheses about the machine language, arithmetic system, and logical rules employed by the brain in its computations. Given his analysis, von Neumann didn't see how the brain could compute effectively using any of the presently known languages, arithmetic operations, and logical rules. He concluded that the brain must be using a language, arithmetic, and logic that were radically different from those invented by humans.

Von Neumann was one of the top experts in all aspects of computing—hardware, programming, numerical analysis, and computability theory—and one of the most rigorous minds ever to discuss the computational organization of brains. His last book presents one of the most sophisticated comparisons ever made between computers and brains, with the hope that a new discipline, automata theory, would soon develop to explain them both. Automata theory never materialized in the form that von Neumann dreamt: mathematicians developed A. Turing's computability theory in a way that applies to artificial computers; engineers went on to use C. Shannon's theory of information for their communication problems; while neuroscientists studied neural mechanisms by their own means. But the idea that brains are, somehow, computers has remained a fundamental assumption of many psychologists, philosophers, and neuroscientists of today. *The Computer and the Brain* is an important historical document, from the era in which the computational theory of the mind-brain was introduced and explored in detail. It's a landmark in the history of computing, psychology, and neuroscience, and it's required reading for anyone interested in the foundations of those disciplines.

References

Aspray, W. (1990), *John von Neumann and the Origins of Modern Computing*, Cambridge, MA: MIT Press.

Aspray, W. and Burks, A. W., eds. (1987), *Papers of John von Neumann on Computing and Computer Theory*, Cambridge, MA: MIT Press.

Burks, A. W., von Neumann, J., et al. (1946), *Preliminary Discussion of the Logical Design of an Electronic Computing Instrument*, Princeton: Institute for Advanced Studies.

Heims, S. J. (1991), *Constructing a Social Science for Postwar America: The Cybernetics Group, 1946–1953*, Cambridge, MA: MIT Press.

von Neumann, J. (1945), *First Draft of a Report on the EDVAC*, Philadelphia, PA: Moore School of Electrical Engineering, University of Pennsylvania.

von Neumann, J. (1951), 'The General and Logical Theory of Automata', in L. A. Jeffress, ed., *Cerebral Mechanisms in Behavior*, New York: Wiley, pp. 1–41.

von Neumann, J. (1956), 'Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components', in C. E. Shannon and J. McCarthy, eds., *Automata Studies*, Princeton, NJ: Princeton University Press, pp. 43–98.

von Neumann, J. (1966), *Theory of Self-Reproducing Automata*, Urbana, IL: University of Illinois Press.

GUALTIERO PICCININI
Department of History and Philosophy of Science
University of Pittsburgh
Pittsburgh, PA 15260, USA
E-mail: gupst1@pitt.edu

Kepa Korta, Ernest Sosa, and Xabier Arrazola (eds.), *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*, Philosophical Studies Series 79, Dordrecht/Boston: Kluwer Academic Press, 1999, xi + 187 pp., $93.00 (cloth), ISBN 0-7923-5973-9.

News Bulletin: Scientific Community Reaches Consensus—Human Beings are Essentially Irrational! What could such a headline mean? The traditional definition of 'man' as 'the rational animal' is based on the idea that human beings possess some capacity that separates them from other animals. So perhaps we should read our headline as declaring that we, after all, have exactly the same cognitive abilities as say, monkeys. From what we know, this seems an unlikely development. How then are we to understand the view of some cognitive scientists that human beings are 'essentially' or 'systematically' irrational?

Over the last 30 years or so, a series of experiments have been conducted aimed at testing subjects on such 'rules of reasoning' as the conjunction rule in probability theory and modus ponens in propositional logic. These experiments have produced a body of results that, it has been argued, support the claim that human beings are essentially irrational. The details of the various experiments are important, but I do not have space to describe all of them. In what follows I will discuss