BOOKS: PSYCHOLOGY

The Meaning of Make-Believe

Alison Gopnik

The Work of the Imagination
by Paul L. Harris

Blackwell, Oxford, 2000. 234 pp. $59.95, £50. ISBN 0-631-21885-8. Paper, $26.95, £14.99. ISBN 0-631-21886-6.

Imagine the following scenario: You walk into a room and are greeted by a wild-haired person wearing a piece of sparkly cloth around her shoulders and a cardboard crown on her head. In an unnaturally deep and loud voice she informs you that she is the Queen of Fairyland and proceeds to summon invisible spirits to serve her. Where are you? Three possibilities spring to mind. You are in the schizophrenia ward of a psychiatric hospital. You are attending a performance of *A Midsummer Night's Dream*. Or you are in preschool.

Young children typically spend hours pretending. Why? Until recently, the standard answer was that they behaved in this strange way because they thought in a strange way. Both Freud and Piaget, for example, believed that young children were fundamentally irrational. They were supposed to loosely associate ideas rather than organizing them logically. Their thoughts were supposed to fulfill wishes, rather than reflect reality. Children were supposed to be unable to think causally or to discriminate fantasy from fact.

Paul Harris's intriguing and lucid book summarizes twenty years of empirical research, much of it from his own lab, that definitively disproves these views. In the author's simple but clever studies, even two- and three-year-olds turn out to be adept at distinguishing pretense from reality. They may spend hours pretending, but they know they are pretending. They don't try to eat the pretend ice cream or talk on the pretend telephone, and they can describe with great accuracy how thoughts and facts differ. Moreover, children can make sophisticated causal inferences and they can even reason counter-factually. In Harris's studies, for example, three-year-olds can explain that if an imaginary Teddy hadn't spilled the ink, his hands would not be dirty. Other developmental psychologists have come to similar conclusions. Marjorie Taylor (*1*) has found that children with imaginary friends are perfectly aware that their companions really are imaginary (though they are no less beloved because of it); Henry Wellman showed that even the youngest children can provide appropriate causal explanations for simple events (*2*).

These empirical advances have completely reversed the traditional idea of the irrational, fantasy-ridden child. But they leave us with a puzzle that is the central theoretical question of *The Work of the Imagination*. If young children are so good at understanding reality and differentiating it from pretense, why do they spend so much time pretending? Preschoolers understand the real world very well, but apparently (all things considered) they would prefer to live in the unreal one.

Harris's answer is that the children's pretenses really are more like theater than like psychosis. He points to studies that show how deftly adults take on the perspective of a fictional narrative. Ordinary, everyday, unhistrionic adults are remarkably good at putting themselves into the shoes of a fictional character. They even generate, in detail, the appropriate fictional spatial orientation, and they genuinely experience the appropriate fictitious emotions. Becoming caught up in a story, it seems, is not just a rarified experience of Shakespeare-lovers, but a deep part of our human cognitive competence.

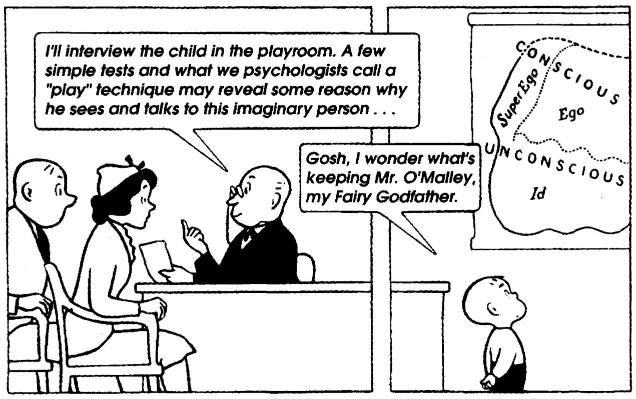

CREDIT: CROCKETT JOHNSON, *BARNABY* (HOLT, NEW YORK, 1943)

Harris argues convincingly that children are doing just this in their pretend play. Like adults, they can follow through on fictional premises even when they are perfectly aware that these are fictional. In fact, children do so in ways that reveal the extent of their causal knowledge. Two-year-olds predict that imaginary tea spilled on an imaginary Teddy will make the Teddy wet and mad, not dry and happy. And although children may be genuinely frightened by the imaginary monster under the bed, we adults, after all, genuinely weep over Romeo and Juliet.

But this still leaves the puzzle of why human beings evolved this fictive capacity at all. It is easy to see why understanding the real world would be adaptive. But why would telling and understanding patently untrue stories help us to survive? Harris suggests that the answer involves language. Children begin to pretend at about the same time they begin to speak, and there is some evidence that language, drama, and art all evolved at around the same time. Harris points out that in order to take advantage of linguistic information from others we need to take their narrative perspective, even if the story they tell is far removed from our immediate experience.

This definitely seems to be on the right track, but the problem that remains is why so much pretense and fiction should be so drastically removed from any possible experience, our own or others. Why should evolution have designed us to believe in fairies, or at least to act as if we do? The promise of Harris's book is that further empirical developmental research will eventually tell us how our sojourns in the imaginary worlds of Teddy and Titania help us to get along in the real one.

The author is in the Department of Psychology, University of California at Berkeley, Berkeley, CA 94704, USA. E-mail: gopnik@socrates.berkeley.edu

References

1. M. Taylor, *Imaginary Companions and the Children Who Create Them* (Oxford Univ. Press, Oxford, 1999).
2. H. M. Wellman, A. K. Hickling, C. A. Schult, *New Dir. Child Dev.* No. 75 (spring 1997), p. 7.

BOOKS: COMPUTING

Grappling with Qubits

Günter Mahler

Is physics useful? Considering the endless list of technical devices that support our everyday life, there can be no serious doubt. Is quantum physics useful? Surprisingly, the answer to that question is far less clear. Although any physical system, when studied on a sufficiently fundamental level, is believed to be quantum in nature, more often than not its function can be described in classical terms. This approach even holds for such macroscopic

The author is at the Institute für Theoretische Physik I, Universität Stuttgart, Pfaffenwaldring 57, D-70550, Stuttgart, Germany. E-mail: mahler@theo1.physik.uni-stuttgart.de

quantum phenomena as superconductivity. But it does not hold for a quantum computer. Understanding the operation of such a machine explicitly calls for a quantum mechanical description.

There has been much uncritical praise of quantum computation, in particular (though not exclusively) in the semi–popular media. Most readers of *Science* have probably been exposed to some explanations of the subject. However, even experts in the field often find themselves confused when asked to identify the most decisive facts behind the computational concepts, let alone to assess future directions of development.

Quantum Computation and Quantum Information
by Michael A. Nielsen and Isaac L. Chuang

Cambridge University Press, New York, 2000. 702 pp. $130, £80. ISBN 0-521-63235-8. Paper, $47.95, £29.95. ISBN 0-521-63503-9.

Here the book by Michael Nielsen and Isaac Chuang comes to rescue. But be warned, the insight it provides has a price. *Quantum Computation and Quantum Information* is a challenging text that offers a thorough discussion of the relevant physics and a reference book that guides readers to the original literature. One can certainly profit from reading selected sections or browsing about and letting one's attention be captured by details here and there. Perhaps the best way to use the book, though, is to ask questions and then search within it for answers. Such a self-guided tour can keep one from getting lost in details and can provide a rewarding journey.

Among the most pressing of such questions, I think, are these: (i) What is the difference between quantum information and classical information? (ii) What makes quantum computation more efficient than classical computation? (iii) Why is the actual implementation of a quantum computer so difficult?

To address the definition of quantum information, we first consult chapter 1, "Introduction and overview." Here we are advised to "think physically about computation" and "computationally about physics." Information emerges as the specification of a physical state (on a given level of description), and information may be said to be carried by that state. Classically, we are accustomed to think of a two-level system (states 0 and 1) as representing a unit of information, the bit. This notion is readily carried over to quantum mechanical two-level systems (a spin, say). The encoded information is then called a qubit. It differs from a bit because it specifies a state that is allowed to be in a superposition of 0 and 1. Easy enough, but is this all?

Jumping to chapter 12 ("Quantum information theory"), we find ourselves confronted by strange limits with respect to accessible information about quantum states. "Not only are quantum states impossible to copy, generally, but also they cannot be perfectly distinguished." We are then led to contemplate reality. Although the description of the state, being classical, can be copied and distinguished at will, the actual referent of that description cannot. This is why a classical simulation of a quantum system, even if perfect in terms of the underlying equations, cannot substitute for the behavior of the "real thing." In our scientific training we are always warned against mistaking a description of reality for the reality itself. Nevertheless, it comes as a surprise to see how technically relevant this warning can be.

This brings us to the question about the origin of quantum computational efficiency. Chapter 1 introduces us to so-called quantum algorithms, i.e., algorithms which can be shown to be "better" than any classical algorithm for the same task. So far, the list of such algorithms is depressingly short. Nielsen and Chuang offer some clues as to why this may be so. Our thinking has been trained within the classical world, so our intuition is of little help to suggest useful applications of quantum rules. It remains open whether this will change with time and whether there might even be a real breakthrough. At present, central roles are played by the quantum Fourier-transform and the quantum search algorithm, which are discussed in separate chapters.

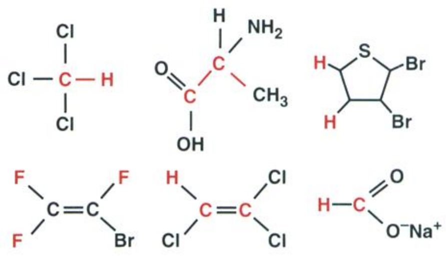

Testbits. Some of the molecules (and their qubits, highlighted) that have been used to demonstrate quantum computation tasks with NMR.

Is there a common feature underlying all known quantum algorithms? The authors introduce a naïve explanation only to reject it after second thoughts. The quantum superposition principle seems to give rise to a form of "quantum parallelism," which, unlike the well-known classical parallelism (separate processors working in concert), would be effective within an individual "processor." Unfortunately, this wonderful multiple processing of information is not accessible to any single experiment, and resorting to ensembles would require resources comparable to those of classical parallelism.

An alternative explanation for the increased efficiency is based on "entanglement," which is a direct consequence of applying the superposition principle to the states of composite systems. In chapter 12, entanglement is treated as a nonclassical "resource"; its transformation, distillation or dilution "emphasizes the interchangeability of different resources in quantum mechanics." Thus, we tentatively conclude that it is this web of interrelations (rather than any singled-out property) that eventually gives rise to the power of nonclassical computation.

Turning to our third question, we read the chapter "Quantum circuits" to familiarize ourselves with what needs to be done. Proceeding to the chapter on physical realization, we learn about basic models taken from different branches of physics including ion-trap physics, cavity electrodynamics, nuclear magnetic resonance (NMR) schemes, and semiconductor quantum dot-arrays. To date, there is no obvious best choice for a quantum computer.

What are the main obstacles an actual implementation must overcome? As for any physical device, there will be control problems. Studying their quantum variants (discussed in the chapter on quantum noise) sheds new light on these. For example, error correction, as in the classical case, is based on clever coding, which makes sure that some memory of the "true state" is still available even after errors of a given class have occurred. Quantum error correction had first been considered impossible because testing (that is, direct measurements) would destroy the precious superpositions. Surprisingly, this shortcoming can be overcome.

In thermodynamic language, we are thus proposing to control micro-states rather than macro-states. Chapter 12 includes an appealing analogy in which error correction is compared to a kind of refrigerator, the common task being to keep the entropy down. In neither case does this contradict the second law of thermodynamics, as we are dealing with open systems. However, although the refrigerator reduces entropy basically by reducing the accessible state space of the cooled system, the trajectory of the quantum computer should remain freely disposable. There is apparently no law of physics forbidding this goal, but also no guarantee it can be achieved at acceptable cost.

Much work in the field is still in progress, and unexpected results may emerge any day. Nevertheless, *Quantum Computation and Quantum Information* will not become immediately outdated. And even after a few years, some addenda and mild revisions should suffice for a second edition. Nielsen and Chuang have set a high standard.

CREDIT: AFTER FIG. 7.21FROM *QUANTUM COMPUTATION AND QUANTUM INFORMATION*