5

Daniel C. Dennett

Co-Director and Professor

Center for Cognitive Studies, Tufts University, USA

---

1. Why were you initially drawn to computational and/or informational issues?

When I was a graduate student in philosophy in Oxford in 1964, I encountered the little anthology edited by Alan Ross Anderson, called MINDS AND MACHINES (Prentice Hall, 1964). In it, were several now-classic essays: Turing's "Computing Machinery and Intelligence," Putnam's "Minds and Machines" (the first of his brilliant series of papers on functionalism) and Lucas' "Minds, Machines and Gödel", one of the earliest attempts to make the invited link between Gödel's theorem and the denial of physicalism (a forerunner of Penrose and others). Anderson was visiting England that year, and I invited him to speak to the Voltaire Society, of which I was then President. After his talk, he and I had a long discussion of the issues raised by these essays, and I was then sure I had to pursue them. The Feigenbaum and Feldman anthology COMPUTERS AND THOUGHT (1963) had Newell and Simon's classic essay "GPS, a Program that Simulates Human Thought," and other seminal papers by founders of AI. Feldman – Julian, not Jerome – was soon to be my colleague at UC Irvine in 1965, and he helped me find my way into the early literature of AI, and introduced me to Allen Newell and others – and asked for my help in rebutting a polemical piece that had just appeared as a RAND memo: Hubert Dreyfus's notorious "Alchemy and Artificial Intelligence," the ancestor of *What Computers Can't Do*. My first publication was a rebuttal of Dreyfus, published in the computer section of the journal *Behavioral Science*. I soon found myself in demand from the AI community as a philosopher who was sympathetic to AI, and able to explain its goals and prospects better, sometimes, than they could.

2. What example(s) from your work (or the work of others) best illustrates the fruitful use of a computational and/or informational approach for foundational researches and/or applications?

I think my essay "Artificial Intelligence as Philosophy and as Psychology" (1978) did a good job articulating and defending the research program of AI, and my rebuttal of Lucas, "The Abilities of Men and Machines" (1970), was an early anticipation of the rebuttals of Penrose and others who later tried to make illicit hay out of Gödel's theorem.

3. What is the proper role of computer science and/or information theory in relation to other disciplines, including other philosophical areas?

A point I have often made is that computer science keeps cognitive science honest. If it weren't for the practical possibility of constructing and demonstrating simplified working models of cognitive processes, we'd still be at the hand-waving stage. Ironically, it was the very difficulties encountered building scalable, realistically sized models that drove home the fact that cognition is much, much more complex than many theorists had realized. This has always been a contentious point, since it is always possible that the difficulties encountered are artifactual – owing to false enabling assumptions in the computer-modellers' kit – but although this is sometimes plausible, most of the good work in computer science (and related fields such as robotics) enlarges our appreciation for just how remarkable our brains are.

4. What do you consider the most neglected topics and/or contributions in late 20th century studies of computation and/or information?

The obvious problem is that we still have no solid theory of *semantic* information. Shannon-Weaver theory is excellent, but not a theory of content at all, and hence provides no answers to such straightforward questions as: what does *Romeo And Juliet* have in common with *West Side Story*? Or, how can a diagram or picture convey (mostly) the same information as a written description? We are "informavores," to use George Miller's vivid term, always eager to obtain more information about what matters most to us,

and so, of course, we want lots of bits – but not useless bits. We want what the CIA calls "intelligence" on various topics. There is no good theory of semantic information so far as I know.

5. What are the most important open problems concerning computation and/or information and what are the prospects for progress?

The Frame Problem still looms large. How can a real time agent ignore the information that should be ignored while noticing both the changes and the constancies that demand its attention? That's not a formal expression of the problem, but it covers a range of issues that arise when one tries to make graceful use of world knowledge to generate relevant expectations about the future.