#831

The Information-Processing Theory of Mind

Herbert A. Simon
Carnegie Mellon University

I am deeply honored to receive this mark of my profession's esteem, but I must also confess that a "lifetime" award has a slightly ominous sound to it. It might be interpreted to imply that the lifetime has been completed and evaluated. The APA is taking something of a risk in this matter: You really do not know how I am going to complete my lifetime or what outrageous thing I will do next. I make no promises. I have a large uncompleted agenda that I will whittle at, but I will try not to scandalize you.

The very generous citation for the award credits me with contributions to a number of sciences. Mr. Freud taught us that not everything is as it seems, and perhaps if you peered into my subconscious you would find that my case is no exception: that far from being a generalist, I have been something of a monomaniac throughout my career. As I have told the story in my autobiography, *Models of My Life*, I can be brief here. Sixty years ago, as a University of Chicago undergraduate, age 19, I undertook for a term project some research on the operation of the municipal government in my native city of Milwaukee and discovered that the way in which decisions were being made in that setting bore little relation to the account of rational decision making I had found in my economics textbooks.

That ignited in me a passionate interest in human decision making (and its attendant processes of problem solving and thinking) that has not yet been quenched. If you are interested in decision making, you go to political science, economics, organization theory, science (scientific discovery), the arts (musical composition or painting) to observe the relevant phenomena.

Decision making is at the center of all of these human activities (and many more). You go to philosophy of science to sharpen your methods and to computer science to find a formal language for expressing your theories. And in particular, you go to psychology to study the underlying processes that enable people to make decisions, solve problems, and generally to think. In fact, if your goal is to understand human decision making, there are very few activities you cannot engage in on company time. Just keep the tape recorder handy.

In the century since its birth, scientific psychology has learned a great deal about human thinking. Thought processes, at least at the symbolic level (second by second), are not the mysteries they are sometimes alleged to be. How much we know about the neurological foundations is another question, one that I will not address. One test of our knowledge of thinking is that we have created expert systems that can and do perform a substantial number of human tasks at a professional level: diagnosing illnesses, designing electric motors and transformers, judging credit risks, and many others.

Not all of these expert systems behave in a humanoid manner—some of them (e.g., the grandmaster-level chess-playing program *Deep Thought*) take advantage of computer speed and memory capacity that we humans simply do not possess. But others (e.g., the chess playing program *MATER*), staying within severe limits on speed of processing, on short-term memory capacity, and on the time required to transfer knowledge to long-term memory, have been shown to behave very similarly, on almost a second-by-second basis, to humans performing the same tasks.

In recent years, a number of systems have been built that simulate human skills over whole ranges of tasks, thereby demonstrating the generality of the mechanisms that are used. In this category, I would mention the General Problem Solver that Newell and I built; John Anderson's ACT programs; Feigenbaum's Elementary Perceiver and Memorizer (EPAM); Newell, Rosenbloom, and Laird's Soar; and others. As these systems all operate within the same general information-processing framework and share many mechanisms, they should not be viewed as competitive theories but as modestly different representations of the same underlying themes.

What are those themes, as they express themselves in human expertise? It turns out (as good science should) that the very complex behaviors of the human expert are produced by a small number of rather simple processes. First, the expert needs a large memory (some millions of chunks) indexed by a discrimination net (EPAM net) that recognizes a corresponding number of different kinds of stimuli. All studies show that it takes a motivated person at least 10 years of intensive study and practice to acquire this memory store and reach "world-class" level in any domain of expertise. Second, the expert needs the ability to solve problems by selective (heuristic) search through spaces of possibilities, using such general heuristic "tricks" as means–ends

Editor's note. This article was originally presented as a Lifetime Contribution to Psychology award address at the 101st Annual Convention of the American Psychological Association in Toronto, Ontario, Canada, on August 20, 1993.

Author's note. Correspondence concerning this article should be addressed to Herbert A. Simon, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213.

Vol. 50, No. 7, 507–508

Herbert A. Simon
Photo by Carnegie-Mellon University.

analysis and such specific ones as the calculus or diagrams or legal reasoning. To say that is all there is to expertise would be an exaggeration, but not a great exaggeration.

In particular, the two mechanisms I have mentioned, recognition and heuristic search, are adequate to explain not only everyday problem solving but also such phenomena as intuition, insight, and the cognitive aspects of creativity. Skills of recognition (based on a large body of knowledge about the subject matter), combined with skills of heuristic search (based on a substantial body of search techniques), are what it takes to be an expert. An important body of evidence in support of this claim resides in the now-numerous computer programs that have been built to simulate closely the processes (not just the products, but the processes as well) of experts in a growing set of domains.

In the light of these developments over the past 35 years, it is time to stop philosophizing about whether, in principle, a computer program can simulate human thinking. Demonstrated fact refutes armchair speculation. Let us treat these new theories just as we treat all other theories: continue to test them against the data of human performance, modify them when they are wrong, and extend them to tasks they have not been shown to handle or where they have not yet been tested against human data.

In speaking this way, I am not trying to describe cognitive psychology as a finished science. Far from it. I am describing it as a progressive science that, building on behaviorist and Gestalt foundations, has continually broadened and deepened the phenomena that it can predict and explain. Let me conclude by mentioning a few items that I think should be near the head of our agenda for extending this science.

First, before we think about a problem, we must have a way of representing it in our minds. That representation may be verbal (or propositional), it may be pictorial (or diagrammatic), or it may take other forms. Steve Kosslyn and others in recent years have explicated what pictorial and diagrammatic representations are all about (the "mind's eye"), building good foundations for this research. The representation of visual perceptions may be the most promising area for building the first real bridges between the symbolic and neural levels of explanation.

Second, artificial intelligence research has revealed a wide range of learning mechanisms that can be used by a system like the human mind: classical reinforcement, adaptive production systems that learn from examples, chunking mechanisms, connectionist networks, EPAM-like discrimination nets, and others. What role, if any, does each of these mechanisms play in human development and learning, and what are the implications of our understanding for instructional methods? How is the "situatedness" of action linked simultaneously to sensory stimuli and stored memories?

Third, we need to reconnect cognition with affect and motivation, probably via the mechanisms that determine the focus of attention. To do this will require not only refocusing the attention of cognitive psychologists, but also broadening the subject matter addressed by specialists in social cognition. It is imperative for progress in instructional methods that we deal simultaneously with cognition and motivation in our research and our prescriptions. We already have too much medicine that is (cognitively) good for the patient—who will not take it—and medicine that patients find delicious—but that contributes little to their cognitive abilities.

This list is already long enough to last me through the rest of this lifetime, and probably the next one, so I need not add to it. Besides, you undoubtedly have some items of your own to which you attach high priority. So let me again express my deep thanks for this award (and for a platform from which to propose this research agenda). Now let us get back to the lab and get busy pursuing it.