CSC: Classic Paper Review/Analysis

Title and Author

Title Computer Science as Empirical Inquiry: Symbols and Search

Author Allen Newell and Herbert A. Simon

Summary/Hook

In this paper/lecture given at the 1975 ACM Turing Awards, Newell and Simon describe and explain the field of computer science, some history concerning its development as a field of study, some influential figures and their contributions to its founding (Emil Post, Alan Turing, Alonzo Church, etc.), some fundamental themes and characteristics pertaining to its emergence as an influential field, and its many applications, in areas such as artificial intelligence, philosophy, psychology and human cognition, and list processing. In particular, they discuss the reasoning behind and the key elements of their Physical Symbol System Hypothesis (PSSH), which fundamentally states that a physical symbol system has the necessary and sufficient means for general intelligent action. Throughout their lecture, Newell and Simon provide precise definitions in describing a physical symbol system, present and explain a variety of examples of physical symbol systems, and offer explanations for their importance in characterizing and appropriately representing human cognition, behavior, intelligent action, and the processes that we undertake during the execution of intelligent action. Most notably, Newell and Simon outline areas of research that are lacking in the study of artificial intelligence and in our efforts to map problem-solving abilities, semantic information, heuristic search procedures, and appropriately represent information via machine programming.

Knowledge Relating to the Cognitive Science Program Learning Outcomes

1. Symbol Systems

Let us return to the topic of symbols, and define a physical symbol system. The adjective "physical" denotes two important features: (1) Such systems clearly obey the laws of physics--they are realizable by engineered systems made of engineered components; (2) although our use of the term "symbol" prefigures our intended interpretation, it is not restricted to human symbol systems. A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure). Thus, a symbol structure is composed of a number of instances (or tokens) of symbols related in some physical way (such as one token being next to another). At any instant of time the system will contain a collection of these symbol structures. Besides these structures, the system also contains a collection of processes that operate on expressions to produce other expressions: processes of creation, modification, reproduction and destruction. A physical symbol system is a machine that produces through time an evolving collection of symbol structures. Such a system exists in a world of objects wider than just these symbolic expressions themselves.

2. Formal Systems and Theories of Computation

We should marvel that two of our deepest insights into information processing were achieved in the thirties, before modern computers came into being. It is a tribute to the genius of Alan Turing. It is also a tribute to the development of mathematical logic at the time, and testimony to the depth of computer science's obligation to it. Concurrently with Turing's work appeared the work of the logicians Emil Post and (independently) Alonzo Church. Starting from independent notions of logistic systems (Post productions and recursive functions, respectively) they arrived at analogous results on undecidability and universality-- results that were soon shown to imply that all three systems were equivalent. Indeed, the convergence of all these attempts to define the most general class of information processing systems provides some of the force of our conviction that we have captured the essentials of information processing in these models. In none of these systems is there, on the surface, a concept of the symbol as something that designates. The data are regarded as just strings of zeroes and ones--indeed that data be inert is essential to the reduction of computation to physical process. The finite state control system was always viewed as a small controller, and logical games were played to see how small a state system could be used without destroying the universality of the machine. No games, as far as we can tell, were ever played to add new states dynamically to the finite control--to think of the control memory as holding the bulk of the system's knowledge. What was accomplished at this stage was half the principle of interpretation--showing that a machine could be run from a description. Thus, this is the stage of automatic formal symbol manipulation.

3. Consciousness and Controversies

There is a steadily widening area within which intelligent action is attainable. From the original tasks, research has extended to building systems that handle and understand natural language in a variety of ways, systems for interpreting visual scenes, systems for hand-eye coordination, systems that design, systems that write computer programs, systems for speech understanding--the list is, if not endless, at least very long. If there are limits beyond which the hypothesis will not carry us, they have not yet become apparent. Up to the present, the rate of progress has been governed mainly by the rather modest quantity of scientific resources that have been applied and the inevitable requirement of a substantial system-building effort for each new major undertaking.

4. Algorithms and Automata

The search for generality spawned a series of programs designed to separate out general problem-solving mechanisms from the requirements of particular task domains. The General Problem Solver (GPS) was perhaps the first of these; while among its descendants are such contemporary systems as PLANNER and CONNIVER. The search for common components has led to generalized schemes of representation for goals and plans, methods for constructing discrimination nets, procedures for the control of tree search, pattern-matching mechanisms, and language-parsing systems. Experiments are at present under way to find convenient devices for representing sequences of time and tense, movement, causality and the like. More and more, it becomes possible to assemble large intelligent systems in a modular way from such basic components.

5. Psychological Investigations

The search for explanations of man's intelligent behavior in terms of symbol systems has had a large measure of success over the past twenty years; to the point where information processing theory is the leading contemporary point of view in cognitive psychology. Especially in the areas of problem solving, concept attainment, and long-term memory, symbol manipulation models now dominate the scene. Research in information processing psychology involves two main kinds of empirical activity. The first is the conduct of observations and experiments on human behavior in tasks requiring intelligence. The second, very similar to the parallel activity in artificial intelligence, is the programming of symbol systems to model the observed human behavior. The psychological observations and experiments lead to the formulation of hypotheses about the symbolic processes the subjects are using, and these are an important source of the ideas that go into the construction of the programs. Thus, many of the ideas for the basic mechanisms of GPS were derived from careful analysis of the protocols that human subjects produced while thinking aloud during the performance of a problem-solving task.