

systems and their interconnections, how sensory nerves come in from the sense organs and how motor nerves go to the muscles. We know much about their micro-chemistry. We know a great deal about the correlations between behavioural and cognitive deficits and brain damage. As opposed to this impressive list of explanatory achievements, there is no significant or trustworthy evidence for non-physical mental phenomena such as parapsychology in spite of endless pronouncements in the popular press. There is not a single parapsychological effect that can be repeatedly or reliably produced in any laboratory suitably equipped to perform and control the experiment. Not one. Honest researchers have been repeatedly hoodwinked by 'psychic' charlatans with skills derived from the magician's trade. Against this record of achievements of materialist scientists, the dualist may argue that these successes concern only the meditative functions of the brain and not the central or higher-level capacities such as reason, emotion and consciousness. Concerning the latter functions, both dualism and materialism draw a blank.

So far as the capacity for reasoning is concerned, machines already exist that execute sophisticated deductive and mathematical calculations in a matter of minutes that would take the human a lifetime. If these high-level capacities take place by a distinct mental stuff, then these faculties must be invulnerable to direct control of pathology by manipulation or damage to the brain. But, in fact, the exact opposite is true. Alcohol, narcotics or senile degeneration of nerve tissue will impair, cripple or even destroy one's capacity for rational thought.

Lastly, the arguments from evolutionary history, the evidence from fossil record, comparative anatomy and biochemistry of proteins and nucleic acids leave no room to doubt that the human brain is the end product of billions of years of evolution from very simple organism to the present highly complex organism.

To sum up, it appears as though we have to give up the view that we are distinct from our animal ancestors, much hard though it is to reconcile to the fact that we are also creatures of matter.

The goals of cognitive science

G. Krishna

Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India

In this paper cognitive science is defined as the study of intelligence as a computational process. The several disciplines which contribute to the study of cognitive science are briefly described so as to glean a birds-eye view of the subject of cognitive science.

COGNITIVE science is the study of intelligence and intelligent systems, with particular reference to intelligent behaviour in terms of computational processes.

It is difficult to propose a really satisfactory intentional definition of intelligence. However, in everyday life we are willing to judge when intelligence is being exhibited.

We say that people are behaving intelligently when they choose courses of action that are relevant to achieving their goals, when they reply coherently and appropriately to questions that are put to them, when they solve problems of lesser or greater difficulty or when they create or design something useful or beautiful or novel.

We apply a single term 'intelligence' to this diverse set of activities because we expect that a common set of underlying processes is implicated in performing all of them.

Today it is quite common to attribute intelligence to both human and non-human systems and, in particular, to programmed computers. Not everyone accepts this usage, but we call programs intelligent if they exhibit behaviours that would be regarded as intelligent if they were exhibited by human beings.

Intelligence is to be judged by the ability to perform intellectual tasks, independent of the nature of the physical system that exhibits this ability.

Cognitive science defined as the study of intelligence and its computational processes has several different approaches. For example:

- We can undertake to construct an abstract theory of intelligent processes, without regard to specific physical or biological implementations (formal logic).
- We can study human or animal intelligence seeking to abstract a theory of intelligence processes from the behaviour of intelligent organisms (experimental psychology).
- We can study computer intelligence, trying to learn the computational principles that underlie the organization and behaviour of intelligent programs (artificial intelligence).

The principal contributing disciplines of cognitive science

From a sociological standpoint, disciplines are defined less by their intellectual structure and content than by the scientists who identify with them. Over time, the intellectual content of a discipline gradually defines its boundaries and membership, whereas its membership gradually redefines its content.

If we are to understand cognitive science, we must know what disciplines have contributed to its formation. Among these are:

1. Experimental and cognitive psychology
2. Artificial intelligence
3. Linguistics
4. Logic and epistemology
5. Neuroscience

Psychology

From its beginnings psychology has been concerned with intelligence. The Binet-Simon intelligence test (IQ) dates back to 1900.

The dominance of behaviourism during the first half of the century prevented experimental psychologists from being much interested in what was going on inside the organism. Hence, there was little speculation and research about the processes involved.

Brain research contributed to our knowledge of the location of functions within the brain but had little to say about the processes involved. Even the precise physiological basis of memory was not unambiguously determined.

During the high tide of behaviourism, experimental psychology focused on relatively simple cognitive performance with emphasis on sensory and motor processes such as rote verbal learning, tracking tasks requiring hand-eye coordination, memory tasks involving short-term retention, and the attainment of simple concepts.

The Gestalt psychologists attempted to develop theories of human cognitive processes for complex cognitive performances like concept formation and problem solving.

The information-processing revolution of the fifties and sixties made it possible for the different specializations of psychology – psychometrics, neurophysiology, experimental psychology and Gestalt psychology – to relate to each other. For example, a discipline like psycholinguistics used the information-processing viewpoint to establish links between psychology and linguistics.

Experimental psychology also started adopting the information-processing viewpoint as well as computer simulation. Those who studied higher mental functions such as concept formation, problem solving and

language made more use of computers than those who studied the simpler memory and perceptual tasks.

Artificial intelligence

The very term artificial intelligence (AI), coined around 1956, incorporated the belief that the concept of intelligence now had to be extended beyond human and animal performance to include artificial systems – computers.

The earliest artificial intelligence programs such as Logic Theorist of Newell and Simon are perhaps best viewed as models of abstract intelligence; nonetheless, their design borrowed ideas from psychological research on memory and problem solving such as associative structures and means and analysis for inference.

In turn AI research has made numerous contributions to cognitive psychology. AI programming languages like Lisp permit modelling of elaborate associative structures – schemata, scripts, frames – to simulate important properties of human semantic memory. Production systems, an important tool in AI, are sophisticated versions of the classical stimulus-response relations and stimulus recognition processes. Robotics has employed ideas in sensory and perceptual psychology and the psychology of vision and speech recognition has borrowed many ideas from AI.

In fact, there has been a close continuing relation between AI and cognitive simulation during the whole thirty odd years of the history of both the subjects and their mutual relevance and synergy was a major motivation for creating a common meeting ground in cognitive science.

Linguistics

The study of language is represented in cognitive science under the labels 'computational linguistics' and 'psycholinguistics'.

Computational linguistics is concerned with the use of computers to process language, for example, in parsing and translation algorithms.

Psycholinguistics, the study of language as a psychological phenomenon, although very much a discipline of cognitive science has followed an autonomous path. Even to this day there appears to be a gulf in communication between cognitive scientists interested in problem solving and concept forming and cognitive scientists interested in language. This is in spite of the fact that both are investigating the same phenomenon: human thinking.

Neuroscience

Neurophysiology and neuroscience generally occupy a very complex place in cognitive science. Probably, a

large majority of psychologists believe that the processes of thinking are, in principle, explainable in terms of the electrochemical processes of the brain. Many, however, believe that theories of intermediate level – such as theories that take information processes and not neurological processes as their primitives – are absolutely essential to the understanding of human thinking. Just as biochemistry is not ‘simply’ physics but must be pursued independently, thinking is not ‘simply’ neurophysiology but requires levels of theory that can link to neurophysiology through a sequence of connecting theories.

Apart from this question of the relation between information processing and direct neurological explanation of thinking, neurophysiology plays a secondary role in cognitive science – particularly in providing hypotheses about fruitful architectures for machine intelligence and perhaps human intelligence.

The architecture of cognitive systems

The fundamental design specifications of an information-processing system are called its architecture.

The components of the architecture represent the underlying physical structures but only abstractly. For example, an architecture for modelling the human brain might contain neurons as components, but the neurons may be characterized quite grossly as binary on-off elements with certain switching speeds.

Another architecture might characterize the brain even more aggregately, with units such as long-term memory, short-term memory, sensory organs and so on.

The amount of detail incorporated in an architecture depends on what questions it seeks to answer, as well as how the system under study is actually structured.

The notion that architectures may be specified at different levels of abstraction is best seen in computers. We speak of a computer as having Von Neumann architecture when it has addressable memory capable of storing both program and data, input and output devices and symbol-processing capabilities that operate serially, including operators for comparing symbol structures and branching. The specification does not say anything at all about the physical devices realizing this scheme. They may be of the most varied and disparate kinds.

At the next level of abstraction, the architecture of a system may be described by defining a specific language. For example, Lisp defines an architecture for a list-processing system. Memory is organized into associative structures, lists and property lists. The basic operators of the language allow manipulations of these lists, whose elements are symbols which themselves are capable of indexing or pointing to other symbol structures.

In the contemporary practice of cognitive science, models of the human nervous system are defined at one of two different levels, connectionist and symbolic.

The elements of the connectionist systems are conceptualized as highly simplified and schematized neurons interconnected in a network. The operators modify the network and, in particular, modify the strengths of the connections between elements.

The elements of symbolic systems may be conceptualized as symbols held in one or more memories. The memories may be organized as lists or networks but the elements of these are interpreted as symbols rather than as neurons or perceptual features. The symbolic network architecture describes the system at a higher, more abstract level than does the connectionist architecture.

In evaluating the plausibility of different architectures as responsible models of the human nervous system, it is good to keep in mind some of the parameters a system must fit if it is to claim that it describes human cognition.

It takes about one millisecond for a signal to cross the synapse between two neurons and longer, of course, for a sequence of such transmissions. A simple act of recognition takes roughly a second. Hence, all of the activities between simple neural events and overt behaviours emanating from a few elementary information processes must be squeezed into a time range of only three orders of magnitude. Stated otherwise, an act of recognition cannot require more than a thousand successive synaptic crossings.

When our concern is with modelling the human brain, these parameters put important constraints on the architecture of a serial system. But they also put severe constraints on the parallel systems that have been proposed, because the equilibration processes used by the connectionist systems require numerous rounds of successive approximation, or ‘settling down’. It is not evident one way or the other whether their time requirements are greater or less than those of a serial recognition system²⁻⁸.

1. Newell, A. and Simon, H. A., *IRE Trans. Inf. Theory*, 1956, **IT-2(3)**, 61–79.
2. Posner, Michael I., *Foundations of Cognitive Science*, MIT Press, Massachusetts, 1989.
3. Johnson-Laird, P. N., *The Computer and the Mind: An Introduction to Cognitive Science*, Fontana Press, 1988.
4. Jackendoff, R., *Consciousness and the Computational Mind*, MIT Press, Massachusetts, 1987.
5. Stillings, N. A. et al., *Cognitive Science: An Introduction*, MIT Press, Massachusetts, 1987.
6. Churchland, P. M., *Matter and Consciousness*, MIT Press, Massachusetts, 1988.
7. Dennet, D. C., *Brain Storms, Philosophical Essays on Mind and Psychology*, MIT Press, Massachusetts, 1988.
8. Johnson-Laird, P. N., *Mental Models*, Cambridge University Press, Cambridge, 1983.