

REVIEW ARTICLE

THE ORIGINS AND NATURE OF THE COGNITIVE PARADIGM : AN OVERVIEW

SUBRATA DASGUPTA

*Institute of Cognitive Science,
Center for Advanced Computer Studies, and
Department of History,
University of Louisiana at Lafayette,
Lafayette, Louisiana 70504 - 3772 (U.S.A.)*

(Received on July 22, 2000)

Abstract : Although cognitive science, the 'science of mind' is generally regarded as of relatively recent vintage, its origins can actually be traced back to ideas from the 1930s. The purpose of this paper is to offer a view of the essential nature of what has come to be called the *cognitive paradigm* - the framework around which cognitive science has come to be constructed - and about its origins and development. Of particular note is the interdisciplinarity of the field, with its strong links to psychology, linguistics, neurobiology, computer science and philosophy. As will be explained below, this interdisciplinary character is rooted in the historical development of the science itself, and can only be understood in such a context. Thus, the present article is strongly historical in spirit and content.

Key words : cognitive science cognition psychology
artificial intelligence symbolic representation
neuroscience linguistics information processing

INTRODUCTION

In 1977, a new journal called *Cognitive Science* was inaugurated; two years later, the Cognitive Science Society was instituted at a conference in La Jolla, California (1). If the founding of a first journal and the creation of a professional society are any indicators that a field of study has come of age, then these dates project a sense of how recent has been the emergence of cognitive science as a distinct scientific arena. However, this is not to say that cognitive

science emerged dramatically, *de novo*, in the mid-1970s; on the contrary. The historical origins of cognitive science is a complex, tangled web involving many disciplines, some of which, like philosophy of mind, go back to antiquity while others, such as computer science, are almost as young as cognitive science itself. And while a few histories of the subject have been published (1, 2), a truly comprehensive, analytical study of the ways in which cognitive science emerged from its various contributing remains to be written.

The purpose of this article is vastly more modest. Taking a cue from Kuhn that the hallmark of a genuine scientific field is the presence of a *paradigm* that is, a network of theories, laws, principles, models and exemplars (3), the aim of this paper is to offer a view of the essential nature of the *cognitive paradigm*. This article is, thus, strongly historical in spirit and content.

What is the cognitive paradigm?

We all have a rough idea of what constitutes a revolution. As in the realm of politics and societies so also in the realm of ideas, certain transformations occur that are radical enough from what preceded them to be described as 'revolutionary'. More precisely, a revolution in the world of ideas involves at least two characteristics. One is that it involves a radically new way of looking at things: new kinds of questions are suggested and posed about the field of interest, new tools are employed, a new vocabulary or even language may arise, a new world view emerges. This complex of entities is precisely what Kuhn meant by the word 'paradigm'. Thus, one characteristic of a revolution in ideas is the emergence of a new paradigm.

The second characteristic is the presence of a significant community of people who are committed to the new paradigm. It is not so much the size of the community that is important but the influence its members command.

It is within this context that we can talk meaningfully about the cognitive revolution which, in essence, refers to the emergence, development and deployment of a particular

paradigm—the *cognitive paradigm*—as a way of understanding the nature and working of mind.

What is the essence of this paradigm? Well, on the one hand, we have *behavior*: observable characteristics of humans and animals, manifested as speech, interactions with the external world, generation of ideas, and so on. On the other hand, we have the *brain*: a body of physical matter, the various components and activities of which, by way of its interaction with the outside world, *give rise to behavior*.

The problem is that there is a huge *conceptual gap* between overt behavior and brain matter, which makes it rather difficult to try to understand behavior directly in terms of neurophysiological events. What scientists very often do in order to bridge a conceptual gap between two types of events that are known to be causally related and yet appear vastly removed from each other is to propose one or more *intermediate levels of descriptions* between the two extremes. Effectively, what is created is a *hierarchy of description levels*. The idea is that we are better off trying to bridge the *narrower conceptual gaps* that are thus created between the adjacent levels in such a hierarchy.

The cognitive paradigm (or the idea of *cognitivism*) is based on this notion of hierarchy. The essence of the cognitive paradigm is the *development of an intermediate layer (or possibly several layers) of explanation, understanding and description of ('mental') phenomena lying between the extremities of overt behavior and the physico-chemical activities of brain matter*.

The problem of understanding the mind can thus be broken down into several smaller, more manageable subproblems: to explain overt behavior-our utterances, the generation of ideas and thoughts, the way we make sense of the external world, the fact that we learn, the manner in which children grasp their first concepts of numbers-in terms of intermediate level processes; and then explain the primitive operations from which the intermediate level processes are built in terms of neurobiological processes. Thus, if overt behavior pertains to the world at large-the social world-and if brain matter belongs to the natural world of physics and chemistry, the 'cognitive level' mediates between the social and the physico-chemical.

Origins of the cognitive paradigm

The origins of the cognitive paradigm can be traced back at the very least to the 1930s. In America, the dominant paradigm in psychology between World Wars I and II was, of course, *behaviorism*, which more or less ignored mental phenomena altogether; rather, behaviorism owed allegiance only to overt behavior. Yet, even in America, there were a handful of scientists who, while admitting to being behaviorists, were not averse to talking about 'cognition' and 'mental processes'. One such psychologist was E.C. Tolman in whose work we find the idea of mind as being some entity physically located in then brain but which, in terms of its properties, is distinct from the brain. In particular, Tolman regarded the mind as a *purposive* entity; this concept of purpose is, of course, entirely absent in the realm of physics, chemistry and neurobiology (4).

Tolman's was a relatively rare cognitivistic voice in an America dominated by behaviorism. Outside the United States, behaviorism was never quite so universally accepted and, indeed, the basic structure of the cognitive paradigm was far more in evidence. From Cambridge, England, F.C. Bartlett published his book *Remembering* (1932) in which he postulated the presence, in the head, of an "organization of past reactions or of past experiences" - in words, a schema representing one's prior experiences-which then participated in interpreting and making sense of new experiences in such mental acts as remembering (5). A decade later, Kenneth Craik suggested that thinking entails the constructions, within the nervous system, of a *model* of the external situation one is thinking about; and that "thought models or parallels reality" (6, p. 57). Craik went on to say that:

the essential feature (of thought) is... symbolism, and.... this symbolism is largely of the same kind as that which is familiar to us in mechanical devices which aid thought and calculation. (6, p. 57).

Craik was writing at a time when calculating devices were (a) mechanical or electro-mechanical, and (b) analogue devices. But he had somehow grasped the notion that thought entails the use of symbols which represent things and situations in the real world. He also appealed to the analogy of the calculating machine which itself, in some sense, models the phenomenon about which it calculates.

Craik's insight needs to be appreciated. It is not merely, he wrote, that "thought employs symbols", but also "the whole of thought (is) ... a process of symbolism" (6, p. 58). As for the nature of this 'process of symbolism', Craik went on to explain that:

If the organism carries a "small scale model" of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilise the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer and more competent manner to the emergencies which face it. (6, p. 61).

In other words, Craik conceived thinking as involving symbolic representations (or models) in the head of aspects of the external world, and symbolic representations of the human agent's actions. Thought would then entail the manipulation of the symbolic models by the represented actions – that is, by mentally *simulating* actions and their effects on external reality.

The advent and impact of the digital computer

Craik's premature death in 1945 precluded further development of his ideas and, certainly, while well known in England, his work seemed to have had virtually no effect in America – which is where a number of quite remarkable interdisciplinary sciences developed in the years immediately following World War II – including information theory, switching circuit theory based on Boolean logic, game

theory, general systems theory, and operations research. In the tangled history of the origins of cognitive science, these would all leave their imprint in one way or another. But the most influential development was that of the electronic digital computer between 1946 and 1949. We find evidence of the immediate appeal of, what was apparently, a purely technological development for thinkers on mind, brain and thought in the interdisciplinary conference, sponsored by the Hixon Foundation on "Cerebral Mechanisms in Behavior", held in Pasadena, California in 1948 (7). Here, John von Neumann, mathematician and co-inventor of game theory, as well as co-inventor of the concept of the stored program digital computer, compared certain features of the computer to the brain (8). As an example, he pointed out that *functionally*, neurons transmit impulses which possess an "all-or-nothing" character, similar to the digital elements in a computer. *Functionally*, the neurons are analogous to digital electronic devices, even though the matter of which they are made bear no resemblance to each other. In fact, in analogizing between brain and computer, von Neumann was drawing upon earlier work by Pitts and McCulloch who had shown that the functional behavior of a network of neurons could be modeled in terms of the principles of mathematical logic, the resulting model being termed a 'formal neural network' (9).

But the Hixon Symposium participants had more on their minds than computers and digital circuit elements; they had *mind* on their minds. Behaviorism, the dogma that had eschewed any discussion of mind, mentalism or cognition as being within the

purview of scientific psychology, came under serious attack by such scientists as the psychologist Karl Lashley and the mathematician-neurophysiologist Warren McCulloch. Thus, Lashley would argue that the nervous system, far from being a "quiescent or static system" is organized and 'actively excited', and that 'behavior is the result of interaction of this background of excitation with input from any designated stimulus' (10, p. 112). What one needed to know, then, was the "general character of this background of information"; for only then can the effect of some input stimulus be understandable. And McCulloch, for whom 'mind' meant "ideas and purposes", argued that the 'mind is in the head' precisely because:

there, and only there, are hosts of possible connections (between neurons) to be formed as time and circumstance demand. Each new connection serves to set the stage for others yet to come and better fitted to adapt us to the world. (11, p. 56).

Two years after the Hixon Symposium, the logician and computational theorist Alan Turing suggested an experimental situation in which one might not be able to distinguish between the manifestation of human intelligence and the behavior of a computer (12). Turing's idea was the following. Imagine an experimental setup in which a computer is in one room and a person in another; in a third room there is a second person who serves as an interrogator. The computer is programmed to receive and answer questions from the interrogator; and the latter, using appropriate devices, puts questions to both

the computer and the other person, though he does not know where either is located. Through appropriate devices, the interrogator also receives answers to his questions from human and machine. Under this set up, if the interrogator is unable to distinguish between the computer's answers from those of the person, then the computer can be said to think and be intelligent in the sense that human beings are normally said to be intelligent.

Turing's essay remains one of the genuine classics in the early history of cognitive science; his testing criterion (known as the 'Turing test') set the ultimate standard of what would count as 'artificial intelligence'.

Annus mirabilis: 1956

Turing's essay, though provocative and, ultimately, influential, was fundamentally speculative. If we are to identify an *annus mirabilis* for the cognitive revolution, that year must be 1956, for in that year, a series of relatively independent developments became public, the collective impact of which would be felt very quickly.

Of these developments, one was entirely linked to the digital computer; and its specific nature was to utilize the computer as a means for talking about mental processes, and as a physical instrument with which experiments concerning mental phenomena could be performed.

The work being alluded to was due to Herbert Simon and Allen Newell, one a social scientist, the other a mathematician. At a symposium held in September 1956 at

the Massachusetts Institute of Technology, Newell and Simon described a 'machine'—actually a computer program—called the Logic Theorist that was capable of discovering proofs for theorems in mathematical logic (13). Here, then, was a machine exhibiting behavior akin to the highest form of human intelligence, able to prove the first few theorems stated in Russell and Whitehead's *Principia Mathematica*. Leaving aside the question of whether the 'architecture of cognition' functioned in ways similar to the program, the Logic Theorist suggested the possibility of *artificial intelligence*; and the possibility that, perhaps, not all minds have to be biological in origin. Logic Theorist also demonstrated rather formally and precisely what such concepts as *symbol processing*, *symbolic representation*, and *information processing* might mean in an empirical or operational sense. It is important to note that amongst the sources of ideas that would directly or indirectly feed into the design of the Logic Theorist were Simon's earlier work on the psychology of human decision making in organizations (14), and Von Neumann and Morgenstern's seminal work on game theory (15)—both, seemingly very unlikely source indeed (Dasgupta, unpublished)!

This paper was not the only contribution this particular conference made to the development of the cognitive paradigm. A young linguist named Noam Chomsky presented a paper which would become the root of yet another major movement in the cognitive field, one which is sometimes called the 'Chomskyan revolution' in linguistics (16). Chomsky would argue in his various publications that the creativity we observe in linguistic behavior, our ability to *generate*

sentences we may not have ever uttered before, and the ability to understand sentences we may not have ever heard before, can only be explained by the existence, in the mind, of a system of rules and representations that encode *universal* principles of language that are common to all specific languages. Chomsky connected linguistics, the systematic study of language structure and use, with cognitivism in the following terms:

The person who has acquired knowledge of a language has internalized a system of rules that relate sound and meaning in a particular way. The linguist constructing a grammar of a language is in effect proposing a hypothesis concerning this internalized system. The linguist's hypothesis, if presented with sufficient explicitness and precision, will have certain empirical consequences with regard to the form of utterances and their interpretations by the native speaker. (17, p. 26).

For Chomsky, the 'internalized system of rules' is knowledge of language.

There were also other seminal publications in that same year that would contribute significantly to the emergence of the cognitive paradigm. The psychologist George Miller suggested that there were some basic limitations on human information processing capacity (18). If Newell and Simon were influenced directly by the workings of the computer, Miller was influenced by the new mathematical theory of information which had been invented in the 1940s. Also in 1956, Jerome

Bruner, Jacqueline Goodnow and George Austin published a work in which the task of categorization or concept formation was explained as a constructive process involving the deployment of strategies (19).

The important lesson that emerges from these landmark studies, studies that led to the shaping of the cognitive paradigm, is that the latter, unlike most scientific paradigms, is fundamentally interdisciplinary. Indeed, one widely accepted view is that cognitive science lies at the intersection of five 'mainstream' disciplines, namely, psychology, philosophy, linguistics, artificial intelligence and neuroscience. Some would disagree that these constitute *all* the participating disciplines at the intersection. For example, Gardner includes anthropology as another contributing field (2), while this writer has come to believe, on the basis of historical case studies of scientific creativity, that some aspects of cognition cannot be explained without appeal to the social sciences.

However, the one shared belief that is more or less unchallenged by those philosophers, psychologists, linguists, artificial intelligence researchers and (perhaps to a lesser extent) neuroscientists who call themselves cognitive scientists, is that cognition entails manipulation and processing of *representations* (also called *information*).

Indeed, a large part of the research program in cognitive science, whether conducted by psychologists, linguists, computer scientists or philosophers, is concerned with the question of what can be

the nature of mental representation for different kinds of cognitive tasks. For instance, researchers in the field of cognitive development studying how and when infants first develop the concept of numbers debate on the form of the infant's representation of numbers. One proposal suggests that each distinct *number* is represented in the infant's mind by a distinct symbol; an alternative hypothesis is that each distinct *object* perceived by the infant is represented by a distinct symbol (20). Psycholinguists have used certain formal structures called trees (that look roughly like organization charts) to depict how humans may represent, construct, parse and understand sentences (17, p. 29). Cognitive psychologists and artificial intelligence researchers interested in how humans (or computers) learn concepts—for example, the concept 'computer'—have suggested representations called *semantic networks* that show relationships between the concept of interest and other concepts, facts, objects, and properties, so that the concept of interest (such as 'computer') is meaningful by virtue of these various relationships. Psychologists inquiring into how humans solve problems may suggest that the subject possesses knowledge of the task domain (e.g., in the domain of chess or other board games) in the form of rules that specify 'IF such and such condition is the case THEN take such and such action' (21).

How does one *test* the validity or plausibility of these abstract theories about formal representations? Cognitive scientists devise various laboratory experiments on human (or sometimes primate) subjects; or they may construct computer models based on the representation and simulate

them on the computer. In either case, the purpose is to corroborate or refute certain properties (e.g., behavior or performance) of the subject or of the simulated model as are predicted by the hypothesized representation. Sometimes, computer simulation and laboratory experiments may be both performed to test the validity of the representation. As a specific example, Langley and his collaborators postulated a number of general rules of the "IF ... THEN ..." type to represent part of the knowledge for a computer program (called 'BACON') which was able to use these rules to 'discover' Kepler's third law of planetary motion (22). Subsequently, Qin and Simon presented the same type of data that had been fed to BACON to human subjects who were given the task to 'discover' possible laws connecting the data. The subjects were also asked to verbalize their thoughts and strategies as they went about solving the problem; from an analysis of such *verbal protocols*, Qin and Simon argued that the human subjects were deploying the same kinds of rules as had been implemented in the computer models (23).

It must be noted that while the nature of representations occupy the attention of much of cognitive science, there are other aspects of cognition that cannot necessarily be investigated at the *level* of representation. For example, in studying the nature of *creative thought*, the questions of interest might be: What kinds of goals initiated a particular act of creation? How did these goals *originate*? What kinds of *knowledge* did the creative being draw upon in the course of creative work? What kinds of *inferences* or *mental actions* can

we infer for a given act of creation? Here, the researcher assumes that such entities as goals and knowledge are represented in the person's 'working memory' and 'long term memory' respectively. But what these representations are, are not so much of interest as what specific entities are being represented (i.e., what goals? what knowledge? what inferences?) These kinds of questions are more appropriately addressed at a level of cognitive abstraction termed by Newell the 'knowledge level' (24); one might say that the knowledge level of abstraction is at the very edge of the cognitive paradigm, where it intersects with the paradigms of the social sciences. Studies of the cognitive process of creativity, asking questions of the above sort, are typically of the sort that reside at this 'edge', and have been carried out at the knowledge level (25).

The computation-cognition connection - revisited

It is not just representation that is taken to be the core of cognition but the processing and transformation of representations also; and here *computation* and the *computer metaphor* became the principle conceptual tools by which most cognitive scientists envisioned how representations could be manipulated and processed.

The complex relationship between technology and science is one of the enduring and fascinating themes in the history of human culture. The appeal to artifacts and their workings as a *model* or *metaphor* for the workings of nature is one aspect of this relationship; yet another, more obvious, one is the use of technology as scientific *instruments*.

The one helps to clarify our thoughts about the natural world, the other extends our capabilities to perceive the natural world; but both serve a common purpose: as amplifiers of human ability to comprehend nature.

It is in the former sense in which the computer has most interestingly served the cause of cognition. First, it provided a significant insight into how the symbolic processing and mental simulation that Craik had suggested as the basis for thought could be *more precisely envisioned*-viz., in the form of computer programs executed on a physical computer. Well before the mid-1950s, people like Turing had realized that the computer was not a number processing device but a general purpose symbol processing machine. Here, then, was a model of how symbols might be represented in the head and how they might be manipulated, processed and transformed into other symbols. A new *vocabulary* was made available, one that could be used to talk about the functional character of cognition, a vocabulary that included such terms as 'program', 'algorithm', 'storage register', 'communication channel', 'processor', 'input-output unit', and so on.

Second, the computer and its operations suggested how something purely physical like the brain, obeying physico-chemical laws and harboring nothing remotely like purpose could develop purposive properties of the kind Tolman had suggested (4). The following analogy was identified:

Computer: HARDWARE->SOFTWARE->FUNCTION

Cognition: BRAIN->MIND->BEHAVIOR

In the case of a computer, there are physical circuit elements which intrinsically have no purposive attributes; they obey the laws of physics. But by virtue of *organizing* these circuit elements, and encoding and storing signals in a memory like device that could then be interpreted by the circuits as instructions or data, the resulting machine could be made to manifest *purposeful behavior*. As a way of re-examining the classical mind-brain problem, it is easy to imagine how appealing and compelling this analogy must have been.

Third, if the computer and its organization yield some sense of how brain matter might be organized to give rise to mental properties and processes leading to behavior, then perhaps the logical organization (or, technically, the 'architecture') of the computer might suggest possibilities for the cognitive organization of the brain. Cognitive scientists, thus, began to systematically investigate the *architecture of cognition*. For example, just as a computer's architecture consists of memories of various capacities and access times, special and general purpose processing units, control units, input/output transducers, instructions, and so on, so also, it was argued, a cognitive architecture might consist of analogous functional components such as 'working' and 'long term' memory, rules or operations, input (perceptual) and output (motor) systems, etc. (26).

Fourth, the very same characteristics of the computer that had provided a metaphorical role for understanding cognition also afforded an experimental apparatus, a kind of telescope for the inner eye, with which one could "observe" mental phenomena. By programming the computer, one could attempt to *simulate* specific, higher level cognitive tasks-such as language understanding, game playing, scientific creativity, and so on. If the simulation produced behavior that was consistent with what was predicted, then the program itself became a *theory* of how humans might think and carry out relevant cognitive processes.

The case of David Marr's 'Vision'

The power of computation as a metaphor for understanding cognition has been, historically, compelling. For many scientists not inherently computationally minded but interested in one or another kind of cognitive phenomenon, the emergence of the computational point of view truly constituted a genuine paradigm shift. A vivid example of this influence of the computer-as-metaphor is the work of David Marr. Like Craik a generation before him, Marr died young. And just as *The Nature of Explanation* was Craik's legacy to cognitive science, so was Marr's *Vision* (1982), published posthumously (27). For our present purpose, the most significant aspect of this book lay not so much in the specific theories he developed to explain certain aspects of the visual process, but in the extent to which computational ideas permeated his explanation of visual

processing. Basically, he posited that vision entails symbolic representation of the external world in the human nervous system in some fashion and the processing of this representation in some manner that allows one to make visual sense of what is looked at. Marr, further, appealed to three different levels of abstraction to explain vision. At the highest, most abstract, level, one explains the visual system in terms of the *goals* of the system, and the logic or *strategy* by which system performs its task. In fact, this corresponds to the 'knowledge level' of explanation mentioned earlier (although Marr termed it the 'computational' level). At the next lower level, one attempts to explain how the higher knowledge level theory might come about in terms of the *representation* of symbols and the *algorithms* used to effect the transformations. Marr termed this the 'algorithmic' level. Finally, at the lowest *hardware* level, one tries to explain how the representations and algorithms are physically wired into the nervous system.

One cannot articulate a more explicitly computational view of cognition than this! It is important to note that to a computer scientist, there is nothing startling about Marr's three levels of abstraction, since much has been written on how computers might be effectively designed and described in terms of multiple levels of abstraction (28). But to a psychologist like Marr, and to other psychologists exposed to his work, these computational modes of explanation were a significant contribution of *Vision*.

Cognition using neural networks

Not all cognitive scientists take computationalism as the *sine qua non* of cognitive science. For instance, Bruner, one of the founding fathers of cognitive science in the 1950s, has in recent years reacted sharply to what he feels is an extreme computational attitude in contemporary cognitive science (29). And even within the computational camp, there continues to rage a fierce debate about the kind of computation or information processing to adopt as the appropriate model. There are those who argue that the usual model of the digital computer is far too removed from the neuronal structure of the brain; and that the extent of *parallel procesing* that goes on in such cognitive tasks as pattern recognition, vision, categorization, etc., cannot be adequately characterized by the 'standard' symbol processing computational model. Instead, a cognitive architecture more akin to the neuronal network anatomy of the brain should be pursued.

What is being referred to here is *connectionism* and what its major advocates have called "brain style computation" (30). In fact, this approach has its roots in the work by Pitts and McCulloch (9) in the 1940s mentioned in an earlier section. These scientists, it will be recalled, showed that the operation of a single nerve cell and its connections with other nerve cells could be modelled in terms of a formal neural network and the principles of mathematical logic. The basic idea of 'modern' connectionism is that information processing is effected by the action of very simple processing elements that are abstract

versions of biological neurons, along with the interaction between these processing elements which communicate with each other by sending and receiving signals. Computation is, thus, "distributed" in a highly parallel fashion within a network of such abstract neurons. In connectionist architectures of cognition-unlike the situation in the conventional symbol representation and processing model-memory resides in a distributed fashion in the *connections* between the neurons rather than in some centralized storage system. As noted above, the debate continues between the proponents of the 'classical' and the connectionist models. In the opinion of this writer, the two models are not necessarily in conflict; they are, rather complementary, signifying two distinct levels of explanation and abstraction of cognitive phenomena.

The role of neuroscience

Implicit in the emergence of the cognitive paradigm is the notion that while *in principle* cognitive phenomena should be explicable in terms of neurobiology, in practice this is neither feasible nor even desirable. In particular, to explain such higher level cognitive phenomena as problem solving, planning or creativity in terms of neurobiological processes poses the same kind of complexity issues as explaining, say, the behavior of a bridge at the level of quantum mechanics. Just as there will always prevail a chemical level of explanation of matter even when these chemical concepts can be explained in physical terms, and just as there will always prevail a biological level of explanation despite advances in chemistry and

physics of life, so also there will remain a cognitive level of explaining the mind-brain phenomena, even when the elementary concepts of cognition have been reduced to neurobiology. The different 'levels of explanation' serve different purposes and will continue to have their unique uses. But this obviously does not mean that the cognitive scientist can afford to ignore the findings of the neurosciences; indeed, any theories that purport to explain cognitive phenomena in terms of information processing models must, at the very least, have what Thagard has referred to as *neurological plausibility* (31).

Conversely, the cognitive neuroscientist cannot afford to ignore the results, theories and models of cognitive science. In fact, a major part of the agenda of *cognitive neuroscience* is to establish and identify the neurological correlates of models and theories of cognitive phenomena. Of course, this agenda itself has a history going back to the identification or *mapping* of brain areas responsible for such cognitive functions as vision and language. A modern, albeit speculative, instance of reducing cognitive phenomena to a neurological mechanism is Francis Crick's conjecture about the mechanism by which the visual cortex attends to visual awareness. Here, Crick drew on the psychological evidence that awareness entails attention, and hypothesized that the shifting of attention from one object to another may be due to the coordinated firing of a group of neurons and the suppression of another group of coordinated neurons (32).

A very different aspect of the neurobiology/cognitive science relationship is exemplified by research on the effect of hormones on cognitive functions. Since the late 1970s, evidence has been accruing that the same hormones that activate and control brain mechanisms underpinning sexual behavior also influence areas of the brain involved in learning and memory (33)

Conclusion

Both historically and substantively, cognitive science, the science of mind, is a genuinely 'multicultural' science: many intellectual cultures have participated in the emergence of the cognitive paradigm, and continue to do so today; each culture has brought to the problem its own paradigms, world views, techniques, tools, traditions and vocabulary. One of the real challenges for the cognitive scientist is to explicitly recognize and accept this multiculturalism. This does not mean that cognitive science will necessarily be a fragmented compendium of independent disciplines; only that there may not be, or it is unlikely there will be, just one mode of explaining cognitive phenomena. Rather, what seems far more likely is that something as complex as the mind will necessarily demand multiple modes of inquiry, multiple metaphors and models, and multiple levels of explanation. If the recognition of scientific multiculturalism is one challenge for the cognitive scientist, the other is to make the explanations at these different levels 'hang together' in a mutually consistent fashion. Needless to say cognitive science can look forward to a very long future.

REFERENCES

1. Bechtel W, Abraharnson A, Graham G. The Life of Cognitive Science In: Bechtel W, Graham G, eds. *A Companion to Cognitive Science*, Oxford, Blackwell, 1998: pp 1-104.
2. Gardner H. *The Mind's New Science*. New York, Basic Books, 1985.
3. Kuhn TS. *The Structure of Scientific Revolutions*. Chicago, University of Chicago Press, 1970.
4. Tolman EC. *Purposive Behavior in Men and Animals*. New York, The Century Co., 1932.
5. Bartlett FC. *Remembering*. Cambridge, Cambridge University Press, 1932.
6. Craik K. *The Nature of Explanation*. Cambridge, Cambridge University Press, 1943.
7. Jeffress LA. ed. *Cerebral Mechanisms in Behavior. The Hixon Symposium*. New York, Wiley. 1951.
8. Von Neumann J. The general and logical theory of automata. In Ref. 7, 1-41.
9. McCulloch WS, Pitts WH. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Math Biophysics* 1943; 5: 115-133.
10. Lashley KS. The problem of serial order in behavior. In: Ref. 7, 112-146.
11. McCulloch WS. Why the mind is in the head. In: Ref. 7, 42-81.
12. Turing AM. Computing machinery and intelligence. *Mind* 1950; 59, 2236: 433-460.
13. Newell A, Simon HA. The logic theory machine: A complex information processing system. *IRE Trans on Inform Theory* 1956; IT-2, 13: 61-79.
14. Simon HA. *Administrative Behavior*, New York, The Free Press, 1947.
15. von Neumann J, Morgenstern O. *Theory of Games and Economic Behavior*. Princeton, Princeton University Press, 1947.
16. Chomsky N. Three models for the description of language. *IRE Trans on Inform Theory* 1956; IT-2, 3: 113-124.
17. Chomsky N. *Language and Mind*. New York, Harcourt Brace Jovanovich, 1972.
18. Miller GA. The magical number seven plus or minus two. *Psychol Rev* 1956, 65: 81-97.
19. Bruner JS, Goodnow, JJ, Austin GA. *A Study of Thinking*. New York, John Wiley, 1956.
20. Uller C, Carey S, Huntley-Fenner G, Klatt L. What representations might underlie infant numerical knowledge. *Cognitive Development* 1999; 14: 1-56.
21. Newell A, Simon HA. *Human Problem Solving*. Prentice-Hall, Englewood-Cliffs, 1972.
22. Langley P, Simon HA, Bradshaw Q, Zytkow J. *Scientific Discovery*. MIT Press Cambridge, MA, 1987.
23. Gin Y, Simon HA. Laboratory replication of scientific discovery processes. *Cognitive Science* 1990; 14: 281-512.
24. Newell A. *Unified theories of cognition*. Harvard University Press, Cambridge, MA 1990.
25. Dasgupta S. *Creativity in Invention and Design: Computational and Cognitive Explorations of Technological Originality*. Cambridge University Press, New York, 1994.
26. Newell A, Rosenbloom PS, Laird JE. Symbolic architectures for cognition. In: Posner MI, ed.. *Foundations of Cognitive Science*. MIT Press, Cambridge, MA: 1989, pp 95-132.
27. Marr D. *Vision: A computational Investigation into Human Representation and Processing of Visual Information*. W.H. Freeman, San Francisco, 1982.
28. Dasgupta S. *Computer Architecture, Vol. 1. Foundations*. John Wiley, New York, 1989.
29. Bruner JS. *Acts of Meaning*. Harvard University Press, Cambridge, MA, 1990.
30. Rumelhart DE. The architecture of mind: A connectionist approach In: Posner MI, ed. *The Foundations of Cognitive Science*. MIT Press, Cambridge, MA, 1989: pp 135-159.
31. Thagard PR, *Mind: An Introduction to Cognitive Science*. MIT Press, Cambridge, 1996.
32. Crick FHC. *The Astonishing Hypothesis*. Touchstone/Simon Schuster, New York, 1994.
33. Williams CC. Estrogen effect on cognition across the lifespan. *Hormones Behavior* 1998; 34, 2: 80-84.