British Journal of Psychology (1985), **76**, 291–301 *Printed in Great Britain*

Cognitive science and behaviourism

B. F. Skinner*

In this paper it is argued that cognitive scientists, claiming the support of brain science and computer simulation, have revived a traditional view that behaviour is initiated by an internal, autonomous mind. In doing so, they have (1) misused the metaphor of storage and retrieval, (2) given neurology a misleading assignment, (3) frequently replaced controlled experimental conditions with mere descriptions of conditions and the assessment of behaviour with statements of expectations and intentions, (4) given feelings and states of mind the status of causes of behaviour rather than the products of the causes, and (5) failed to define many key terms in dimensions acceptable to science.

Cognitive scientists are in no doubt about the importance of their field. A recent 'Report of the Research Briefing Panel on Cognitive Science and Artificial Intelligence' (Estes *et al.*, 1983) claims that the questions to which it addresses itself 'reflect a single underlying great scientific mystery, on a par with understanding the evolution of the universe, the origin of life, or the nature of elementary particles . . . ' (p. 21). The rapid growth of the disciplines for which the Panel speaks is said to be 'advancing our understanding of the nature of mind and the nature of intelligence on a scale that is proving revolutionary' (p. 21).

The movement is called a revolution because it is said to have overthrown behaviourism. 'Behaviorism', says Simon (1980), 'was suited to the predominantly positivist and operationalist views of the methodology and philosophy of science, and seemed to provide some guarantee against metaphysical, "mentalistic" explanations of human behavior. The price paid for these qualities was to confine experimental psychology to relatively simple memory and learning experiments, and to a preoccupation with laboratory rats rather than humans engaged in complex thinking and problem-solving tasks' (p. 76). The discipline that has replaced behaviourism, he says, shows a 'new sophistication', a 'new confidence', and a great gain in 'precision and rigor'. But what is cognitive science, what is behaviourism, and what is the difference between them?

The old mind–body problem is, of course, one issue. The Briefing Panel Report asks: 'How can the brain give rise to all the phenomena that we call mental?' and 'How can mind exist in a world totally controlled by physical laws?' (Estes *et al.*, 1983, p. 21). A more important question concerns the origin of behaviour. Cognitive science takes the traditional position: behaviour starts within the organism. We think and then act; we have ideas and then put them into words; we experience feelings and then express them; we intend, decide, and choose to act before acting. Behavourists, on the other hand, look at antecedent events in the environment and the environmental histories of both the species and the individual. The old stimulus–response formula was an attempt to give the environment an initiating role, but it has long since been abandoned. The environment *selects* behaviour. Ethologists study the contingencies of survival in natural selection that explain species-specific behaviour. Contingencies of operant reinforcement select the behaviour of the individual in a rather similar way but, of course, on a very different time scale.

Cognitive scientists clearly make the person an initiator when they adopt the paradigm of information processing. The adoption is apparently unanimous. As the Panel Report

* B. F. Skinner is Professor of Psychology and Social Relations Emeritus at Harvard University, USA. He presented a version of this paper as the Distinguished Visitor's address to the Annual Conference of the British Psychological Society, University of Swansea, March 1985.

says, 'Cognitive science and artificial intelligence stand together in taking information processing as the central activity involved in intelligent behavior. . . ' (Estes *et al.*, 1983, p. 21). Processing information is, of course, something people have done for thousands of years. They have made records of things that happen – on clay tiles, papyrus, vellum, paper, magnetic tape, and now silicon chips – and they have stored them, retrieved them, and responded to them again more or less as they responded to the originals. Cognitive scientists have taken this practice as a model or metaphor. Are they justified in doing so?

Perception

The difference is perhaps best seen in the field of perception. For cognitive science the direction of action is from organism to environment. The perceiver acts upon the world, perceives it in the sense of taking it in. (Perceive comes from a word meaning capture.) It was an important point for Greek philosophers: to know a thing was to be in contact with it. The world had to be taken in and possessed in order to be known.

There is some ambiguity in calling perception the processing of information. Are sensory data the information to be processed or is information extracted from them? The information theory of Shannon and Weaver concerned itself with transmission of information. Processing must have a product, and for cognitive science that product is, as it was for the Greek philosopher, a *representation*. We do not see the world, we see copies of it.

In a behavioural analysis the direction is reversed. At issue is not what an organism sees but how stimuli alter the probability of its behaviour. They acquire the power to do so from the part they play in contingencies of phylogenic and ontogenic selection. What is 'seen' is the presentation, not a representation.

There is another step in the cognitive account. The perceiver must retrieve a stored history and in some way fuse it with a current representation. As Sutherland (1978) has put it: 'Everyday perception involves assessing and bringing to bear vast stores of knowledge'. In order to 'make sense of the fragmentary pattern of light reaching the eye' one must consider many possibilities, make inferences, and formulate and test elaborate hypotheses. A behavioural translation appeals only to the history of reinforcement responsible for the current effect of a presentation: everyday perception is the product of a vast number of experiences in which fragmentary patterns of light, resembling those of the moment in many different ways, have been present when behaviour has had reinforcing consequences.

Cognitive psychologists also contend that the retrieval of stored knowledge affects not only what is seen but how readily it is seen. Thus, familiar words are seen more readily than rare, 'expected' words more readily than unexpected, and decorous words more readily than obscene. A behavioural account explains the same facts as the effects of past consequences, positive or negative. As Broadbent says, ' . . . nobody disputes the fact that the perception of a word would depend very much upon its probability of occurrence' (1965, p. 11). But there is still a problem for the cognitive psychologist. As Broadbent asks, 'Why is it easier to perceive a word which is probable than a word which is improbable?' (p. 11). The notion of probability alone suffices in a behavioural account.

The inner man

A crucial difference between the two positions arises at the final stage. Once a representation has been constructed, what is done with it? A molecular biologist, Stent (1975), traces the processing of visual data, beginning with the roughly 100 million receptor cells in the retina and the roughly one million ganglion cells which, according to Stent, process the information coming from them by signalling light–dark contrasts and edge effects. The fibres of the ganglion cells connect the eye with the brain, where the 'signals converge on a set of cortical nerve cells' (1975, p. 1056) which, among other things, process information from fields too big for the ganglion cells. As Stent notes, '[The] visual system of the frog abstracts its input data in such a way as to produce only two meaningful structures, "my prey" and "my predator", which in turn evoke either of two alternative motor outputs, attack or flight' (p. 1057). (Why are 'meaningful structures' needed? Why not simply 'to produce only two alternative motor outputs?') Stent concedes: 'It is not clear at present how far this process of cerebral abstraction by convergence of visual channels can be imagined to go . . . Should one suppose that the cellular abstraction process goes so far that there exists for every meaningful structure of whose specific recognition a person is capable (for example, "my grandmother") at least one particular nerve cell in the brain that responds if and only if the light and dark pattern from which that structure is abstracted appears in its visual space?' (1975, pp. 1056–1057). The primrose path of visual perception leads the cognitive scientist to a representation and the neurologist to a grandmother cell. But if what happens when one sees one's grandmother is the result of what has happened when one has seen her and persons like her many times in the past, changes must have taken place in most of the brain. A single point in the brain is needed only to represent a putative 'meaningful structure'.

Stent concludes in a traditional way: 'No matter how deeply we probe into the visual pathway, in the end we need to posit an "inner man" who transforms the visual image into a percept' (1975, p. 1057). The Panel Report cited earlier deals with the last stage rather more cavalierly. 'The end-product of the analysis of a visual scene is subjectively familiar. We recognize the objects present in the environment and their spatial relations to each other and are able to compute distances in a way that permits us to navigate . . .' (Estes *et al.*, 1983, p. 27). Some cognitive touches survive: apparently organisms cannot move about in the world without first 'computing distances', and things they have seen before are only *subjectively* familiar, but 'we' presumably means people, not little people in the brain, and 'objects present in the environment' are presumably presentations not representations. That is good behaviourism.

It is obviously not the whole organism that is said to view a representation of reality. It is some lesser part inside. Cognitive science is forced into saying so when it takes the processing of information as a model. In a behavioural account the whole organism responds, and it responds to the world around it – for reasons which neurology, not cognitive science, will eventually discover.

Rules

Constructing representations of sensory data is only one kind of information processing. Another kind plays a far more important role in cognitive science. Consider a familiar experiment. In an operant chamber a hungry rat occasionally presses a lever extending from one wall. When, for the first time, a bit of food is delivered after a single press, the rat eats and then presses the lever more rapidly. That is a simple fact, but cognitive psychologists want to say more about it. They want to say that 'the rat has learned, and now *knows* that pressing the lever brings food'. 'Pressing the lever brings food' is a description of the contingencies of reinforcement. In some form or other, it is said to pass into the head of the rat in the form of knowledge, a synonym of cognition. Processing has led not to a representation but to a rule.

There is no evidence that anything of the sort happens in a rat, but what about a person? Given the necessary exposure to a verbal community, a person could certainly say, 'When I pressed the lever, food appeared' or, more generally, 'When the lever is pressed, food appears'. People come to 'know' the contingencies in the world around them in that sense. The behaviour is verbal and must be analysed as such.

Verbal behaviour evidently came into existence when, through a critical step in the evolution of the human species, the vocal musculature became susceptible to operant conditioning. Processes shared with other species were already well established, and they can be used to explain the other properties of verbal behaviour (Skinner, 1957). One result is particularly important here. A very large part of the social environment we call a culture consists of descriptions of contingencies of reinforcement in the form of advice, maxims, instructions, rules of conduct, the laws of governments and religions, and the laws of science. With their help members of a group transmit what they have learned to new members, who then behave for either of two reasons: their behaviour is either directly shaped and maintained by contingencies of reinforcement or controlled by descriptions of such contingencies. When we learn to drive a car, for example, we begin with responses to verbal stimuli. Our behaviour is rule-governed. We flip switches, push pedals, and turn the wheel as we are told to do. But consequences follow, and they begin to shape and maintain our behaviour. When we have become skilful drivers, the rules no longer function.

When contingencies are not adequate, however, we return to the rules. Most of the time we speak grammatically, for example, because of the prevailing contingencies in our verbal community, but when the contingencies are not sufficient, we turn to the rules of grammar. (An unnecessary return to rules may be troublesome. Mrs E. Craster (d. 1874) suggested that when the toad asked the centipede, 'Pray, which leg goes after which?' the centipede 'worked her mind to such a pitch/She lay distracted in the ditch/Considering how to run'.)

Cognitive scientists argue that rules are *in* the contingencies and that a person can learn about them in either way. They have therefore felt free to take the easier path and study rule-governed behaviour. Settings are often merely described rather than presented; subjects are asked to imagine themselves in given situations and to say what they would do rather than do it. Contingent consequences are suggested rather than imposed: 'Assume that if you do *A*, *B* will follow'. Yet descriptions of settings are never quite accurate or complete, what people say they will do is not always what they actually do, and a description of contingencies (for example, a statement of the odds in a wager) seldom has the same effect as exposure to the contingencies (in repeated betting with those odds).

In an experiment on reaction time, for example, the subjects were said to have 'understood that their task was to depress the response button as quickly as possible following the onset of the signal lights' (Estes, 1972, p. 726). But how accurately does 'understanding' describe the effect of a history of reinforcement, and how well are contingencies put into effect by saying 'respond as quickly as possible?' Many years ago I arranged contingencies under which a pigeon pecked a key 'as quickly as possible' and found times in the human range (Skinner, 1950). I do not believe one can learn how to do so from the cognitive literature. The response of which a reaction time is measured is quite different from pulling the hand away from a painful stimulus, for example, and to understand the difference (and give neurologists their proper assignment) the contingencies must be specified. Paying subjects according to their reaction times is a move in the right direction but only a short one. Perhaps cognitive psychologists do not care why their subjects respond as quickly as possible; a measure of the speed of a cognitive process may be enough. But in the world at large people behave as quickly as possible for reasons which, for many purposes, need to be understood.

Storage

Are the contingencies of reinforcement to which a person is exposed stored as representations and rules, or is behaviour simply changed by them? When physical records are stored, the records continue to exist until they are retrieved, but is that true when people 'process information?' A storage battery would be a better model of a behaving organism. We put electricity into a battery and take it out when needed, *but there is no electricity in the battery*. When we 'put electricity in', we *change* the battery, and it is a changed battery that 'puts out electricity' when tapped. Organisms do not acquire behaviour as a kind of possession; they simply come to behave in various ways. The behaviour is not in them at any time. We say that it is emitted, but only as light is emitted from a hot filament; there is no light in the filament.

How organisms are changed by contingencies of reinforcement is the field of a behavioural analysis. What is happening inside is a question to be answered by neurology, with its appropriate instruments and methods. Cognitive scientists cannot leave it to neurology because processing information is part of the story they want to tell. As the previously cited Panel Report puts it, '[a] prominent component of artificial intelligence systems is the memory organizations that hold the knowledge and permit finding the right bit of knowledge at the right time' (Estes *et al.*, 1983, p. 25).

Surrogates

Cognitive scientists presumably appeal to the storage and retrieval of representations and rules because they can then explain behaviour by pointing to conditions present at the time the behaviour occurs. The rat was changed when pressing the lever was reinforced, but it presses now 'because it knows that pressing brings food'. The knowledge is a current surrogate of the history of reinforcement. A classical example of a current surrogate is purpose. We say that we have eyes in order to see, but biologists have long since learned to say only that eyes have evolved because variations which enabled organisms to see better were selected. A similar mistake is made when cognitive psychologists call operant behaviour purposive or goal-directed. Features suggesting direction toward a goal are the products of consequences experienced in the past.

Cognitive psychologists need current surrogates because they pay little attention to selection as a causal principle. A new discipline, cognitive ethology, makes that mistake. Animals do many complex things: they use tools, lay traps, entice their victims. How can such behaviour be explained? Ethologists answer by pointing to contingencies of survival in natural selection. Cognitive ethologists must do so, too, to explain features like gross anatomy (for which mental processes could scarcely be responsible), but they contend that not all complex behaviour can have evolved. Instead, what evolved are said to be mechanisms which enabled animals to have expectations and solve their problems by processing information. Behaviour of comparable complexity can, however, be shaped by arranging sequential contingencies of reinforcement, and it is then easier to understand that contingencies of natural selection could have sufficed in phylogenic instances.

Rationality

Expectation as a current surrogate of a history of reinforcement figures in recent studies of rationality. Simon calls the elucidation of the concept of rationality '[one] of the crowning achievements of the social sciences in the past two or three generations' (1980, p. 75). 'The core of the formal treatment of rationality', he writes, 'is the so-called subjective expected utility (SEU) theory' (p. 75). According to Simon, 'human behavior is manifestly directed towards goals and the satisfaction of wants and needs' (p. 75). (In behavioural terms: people behave in ways which have had reinforcing consequences, and the susceptibilities which make those consequences reinforcing arose through natural selection or operant conditioning. 'Wants' and 'needs' are current surrogates of a history of deprivation.) Simon continues: 'the rational actor is assumed to have a consistent preference ordering of all possible states of the world' (1980, p. 75). (In behavioural terms: 'reinforcers can be arranged according to their power to reinforce'.) The axioms of SEU theory imply that 'the actor maximizes his or her expected utility in the light of subjective estimates of the probabilities of events' (Simon, 1980, p. 75). (In behavioural terms: 'people act according to rules which describe or imply contingencies of reinforcement'.) Missing from the cognitive account is any reference to the effect of the contingencies themselves, quite apart from the rules derived from them.

In a fairly typical experiment, two actions having probabilistic outcomes (for example, purchasing two kinds of lottery tickets) are described, and the subject is asked to choose between them. (As in much cognitive research, choice is used because a more direct measure of probability of action is lacking.) The choice is then compared with what would be called rational in the light of the actual contingencies. But what ticket would subjects actually buy – either after the probabilities were first described to them (as rule-governed action) or after they had purchased many tickets and consequences had followed? To the extent that economic theory is concerned with what people *say* they will do, SEU theory may be adequate, but a behavioural scientist (and, one would suppose, an entrepreneur) must be concerned with what they actually do. People most often do those things which have been most abundantly reinforced, without making subjective estimates of the probabilities of reinforcement; and the reinforcers are real, not expected.

Since contingencies are usually more effective than rules derived from them, we should not be surprised that, as Simon reports, 'A fair summary of the findings of these experiments is that actual human choices depart radically from those implied by the axioms except in the simplest and most transparent of situations' (1980, p. 75). This does not mean that people are 'irrational' in the sense that contingencies of reinforcement are not effective. It means that describing contingencies and acting upon the descriptions can obscure the effect of the contingencies themselves.

Feelings

Feelings are among the current surrogates of a history of reinforcement. According to cognitive psychologists, the rat not only learns and hence knows that pressing the lever brings food; it not only expects food to appear when it presses the lever; it *feels like* pressing the lever. A human subject would say as much. In a behavioural account what one feels is various states of one's body, and what one feels when one is behaving or likely to behave is therefore a collateral product of the causes of the behaviour. It should not be mistaken for a cause.

That position is sometimes challenged by asserting that one feeling or state of mind can cause another feeling or state of mind. Fodor (1981) contends that 'mental causes typically give rise to behavioral effects by virtue of their interaction with other mental causes. For example, having a headache causes a disposition to take aspirin only if one also has the *desire* to get rid of the headache, the *belief* that aspirin exists, the *belief* that taking aspirin reduces headaches and so on. Since mental states interact in generating behavior, it will be necessary to find a construal of psychological explanations that posits mental processes: causal sequences of mental events. It is this construal that logical behaviorism fails to provide' (Fodor, 1981, p. 116, italics added).

But what is felt in each case can be construed as either a state of the behaving body or an external controlling variable, and it is *they* that interact. Consider a simpler example. When one pulls one's hand away from a hot object, must one have a disposition to do so, a desire not to be burned, a belief that the hot object exists, and a belief that pulling the hand away will terminate the stimulus? Contingencies of survival in natural selection offer a much simpler explanation of that 'flexion reflex'. Contingencies of reinforcement offer a similar explanation of the fact that one turns the steering wheel of a car to avoid a head-on collision. One need not have a belief in the existence of the oncoming car, a desire to avoid collisions, and a belief that turning the wheel will avoid one. Turning the wheel occurs because it has had fairly similar consequences in the past.

Neglected contingencies of reinforcement can be subtle. Kahneman & Tversky (1984) have reported that people say they would be less likely to buy a second ticket to the theatre if a first had been lost than to buy a ticket after losing the money they had set aside for that purpose. The difference is said to be due to a difference in categorization. A difference in relevant contingencies of reinforcement should not be overlooked. A boy who usually washes his hands before sitting down to dinner quite justly protests when told to wash them if he has already done so: 'I *did* wash them!' We pay a bus fare and protest at a subsequent request to do so again. Buying a ticket to the theatre is in that 'category' of contingencies. Contingencies involving loss are different. We spoil one sheet of paper and take another. And, quite as readily, we use other money to buy a ticket to the theatre.

The black box

Shevrin & Dickman (1980) argue that 'Behaviorism. . . must accommodate itself to accepting the importance of what goes on inside the "black box", especially since we now have methods for investigating its contents' (p. 432). What are those methods?

Introspection is presumably one of them. An early version of behaviourism, like logical positivism, held that science must confine itself to publicly observable events, but radical behaviourism accepts the fact that parts of our bodies enter into the sensory control of what we do, not only in behaviour like figure skating but in the self-observation that is shaped and maintained by a verbal community. But can we actually observe ourselves engaging in mental processes? Do we see ourselves extracting information from sensory data or merely the settings in which we are said to do so and the 'representations' that result? Do we see ourselves storing and retrieving information or merely the information stored and retrieved? We observe the conditions under which our behaviour changes and the changed behaviour, but do we see the process of changing? We have more information about ourselves than about other people, but it is only the same *kind* of information – about stimuli, responses, and consequences, some of them internal and in that sense private. We have no sensory nerves going to the parts of the brain that engage in 'cognitive processes'.

Brain science looks into the black box in another way. It has made great progress in discovering both the architecture and the chemistry of the brain, but what can it say about the processing of information? Visual representations are puzzling enough. What about *auditory* ones? Do we construct representations while listening to the *Emperor Concerto*, and how is it stored in the head of the pianist who is playing it? A behavioural account is incomplete, in part because it leaves a great deal to neurology, but at least it avoids the unnecessary problems of storage and retrieval. Cognitive psychologists, like psychoanalysts, observe causal relations between behaviour and genetic and personal histories and invent mental apparatuses to explain them, but one may doubt that neurologists, with their very different and much more appropriate techniques of observation, will find anything that resembles them. In confining itself to the facts to which its instruments and methods are suited, behavioural science offers brain science a much clearer statement of its assignment.

Simulation is said to be another source of information about the black box. Indeed, the Panel Report cited earlier asserts that cognitive science and artificial intelligence stand together not only in taking information processing as the central activity involved in intelligent behaviour but in 'taking the framework of modern computer science as the foundation for understanding information processing systems' (Estes *et al.*, 1983, p. 21). Apart from the advisability of using a framework as a foundation, are we to understand that the study of computers can tell us what we need to know about human behaviour?

In constructing an artificial organism – that is, a system that exhibits artificial intelligence – one has a choice. One may simulate a non-verbal organism – say, a pigeon – and construct a sensory-motor system the behaviour of which is selected and strengthened by its consequences. Or one may construct a rule-following system that responds as it is directed to respond and changes its behaviour as it is directed to do. The first system would resemble a person who drives a car skilfully because of the way it behaves on the highway. One could study it by arranging contingencies of reinforcement and observing the changes in its behaviour which follow.

It is perhaps natural that those who are concerned with artificial intelligence should choose the rule-following alternative – resembling people who drive cars by making only the moves they are told to make. That kind of artificial organism could be instructed to behave in intelligent ways because specialists in artificial intelligence are intelligent, and it could then be credited with intelligence.

The first of the two systems could, of course, learn to behave verbally, given the necessary verbal environment, and it would then resemble the second. The second, however, would remain forever simply rule-following. A computer may function as either type of system, but if cognitive scientists have actually programmed computers to 'think creatively' and to make 'scientific discoveries', they have simulated the first type. And they must have known a great deal about it in order to do so.

Linguistics is also said to have illuminated the black box, but most of linguistics is itself an offshoot of an earlier cognitive position (Skinner, 1957). According to the Panel Report, speakers 'learn to cope with language', acquire a vocabulary, and '[master] a complex of grammatical rules' (Estes *et al.*, 1983, p. 29). Listeners possess a 'natural language understanding system' with which they produce 'internal representations of the information conveyed by the system' (p. 30) and make 'a semantic analysis of the message conveyed by the language' (p. 29). That is a long way from throwing much light on what speakers and listeners actually do.

In a behavioural analysis meaning is not *in* what speakers say; it is at best in the personal histories and current settings responsible for their saying it. Meaning for the listener is what the listener does, as the result of a different personal history. Rules of grammar describe the contingencies maintained by verbal communities. Languages share 'universals' because they serve universal functions. Nothing is communicated in the sense of being transmitted from speaker to listener. Speakers create settings in which listeners respond in given ways. Sentences are usually 'generated' by the contingencies, and only occasionally with the help of rules extracted from them. Only when the contingencies are inadequate do speakers turn to rules.

The belief that meanings reside in words is hard to resist. A news commentator recently described the distressing problems some women face just before their periods as being caused by the pre-menstrual syndrome. But the problems *are* the syndrome. A similar belief in the power of words is said (Dickenson, 1980) to be shown by traditional doctors in a Moslem sect in Nigeria, who write an Islamic text on a slate, wash the slate with water, and give the water to their patients as medicine, a curious example of the communication of meaning.

Achievement or promise?

The report of the Briefing Panel (Estes *et al.*, 1983) has been cited repeatedly because it has the confidence and enthusiasm typical of the field as a whole. But too many of its sections end with a mere promise. One subtheme is said to open 'a research vista that is more than just the *hope of increased power*' (p. 27), but another is said to be merely '*pregnant with hope*' (p. 21). One field is said to provide 'a fund of knowledge about *possible* mechanisms upon which theories of human cognition can be built' (p. 26). Another 'offers *promise* of a new burst of progress' (p. 29). Cognitive science is said to be able to 'provide a broadened theoretical framework within which *significant progress* can occur . . . ' (p. 36) (italics added). Perhaps we should not ask more from a Briefing Panel, but what has actually been done that could or would not have been done if the processing information revolution had never happened?

Step by step we may paraphrase the accomplishments of cognitive scientists in behavioural terms. They have not discovered 'the form in which information is represented in images'; they have discovered some of the ways in which stimuli enter into the control of behaviour. They have not discovered how 'all intelligence . . . arises from the ability to use symbols'; at best they have learned something about verbal behaviour. They have not studied rational choice; they have studied how people respond under conflicting contingencies of reinforcement or conflicting descriptions of contingencies. And so on. Many of them have made important discoveries and no doubt deserve credit even though they have misunderstood what they were doing. After all, we honour Columbus for discovering America, although he thought he was discovering India, and it scarcely matters that we still call the original inhabitants Indians. Many of the findings of cognitive science find a useful place in behavioural analyses.

The enthusiasm of cognitive scientists is not easily explained by looking at practical achievements. On the contrary, in reactivating the dream of the central initiating control of behaviour, cognitive science has drawn attention away from the *accessible* variables needed in an effective technology. Education is an example. When the first Sputnik went into orbit, Americans were stunned. Why had the Russians beaten us? Something must be wrong with the teaching of science and mathematics in American schools. The National Defense Act was soon passed, and educators went into action. From a conference of experts at Woods Hole, a cognitive theory of teaching emerged. Jerome Bruner's report of that conference, *The Process of Education*, became a sort of bible in schools of education. Students were to be taught to think – to grasp concepts, understand principles, and discover knowledge for themselves. Scientists and mathematicians spent millions of dollars preparing appropriate teaching materials. Twenty-five years have passed, and American students are no better in science and mathematics than they were before.

Several years before Sputnik, in experiments with teaching machines and a system of programmed instruction based upon a behavioural analysis of verbal behaviour, it was shown that what was then taught in American classrooms could be taught in half the time and with half the effort (Skinner, 1968). But classroom practices were not changed, largely because education remained (and, alas, still remains) committed to cognitive theories.

The Briefing Panel's references to teaching machines are strangely hesitant. We are told that 'an intelligent tutoring system that can provide genuine help in educating a student in some well-understood domain, such as mathematics or science, must provide several components. (1) A powerful model of the task domain, so it can itself solve problems in that domain. (2) A detailed model of the student's current level of competence, encompassing both partial and erroneous competence as well as perfect competence. (3) Principles for interpreting the student's behavior, so as to be able to infer the student's knowledge and difficulties. (4) Principles for interacting with the student, so as to lead the student to a higher level of competence' (Estes *et al.*, 1983, pp. 32–33). In many fewer words: to teach mathematics and science, we must (1) define our objectives, (2) find out what students already know, (3) present material to be learned in carefully designed steps, and (4) tell students immediately whether or not they have taken them successfully. Millions of students have suffered and are now suffering from the difference between those versions.

Summary

For many thousands of years people have learned how to change the world and have told others what they have learned. In doing so they have talked about causes and effects. For at least three thousand of those years they have talked about human behaviour in the same way, and for much of that time they looked for causes inside the behaving organism. As Onians (1951) has shown, the Homeric Greek attributed much of human behaviour to the *thumos* and the *phrenes*. The words referred primarily to the heart and lungs but also to what we call feelings and states of mind. (In a rather similar way, many cognitive scientists use 'brain' and 'mind' interchangeably, as if the Greeks had simply chosen the wrong organ.) Plato and, much later, Descartes argued, of course, that there were two kinds of organs and that one could speak of feelings and states of mind without alluding to the body.

If we assume that what is felt are collateral products of the causes of behaviour, feelings can be a useful clue to those causes. We need a language of feelings and states of mind in our daily lives. It is the language of literature and most of philosophy. Clinical psychologists use it to learn many things about the histories of their clients that they could not discover in any other way. There are two languages in every field of knowledge, and it would be foolish to insist that the technical version always be used. But it must be used in *science*, and especially in a science of behaviour.

Behaviourism began by asking philosophers and psychologists for definitions. What were sensations? What was consciousness? What were the dimensions of an idea? The effect was inhibiting, and people who wanted to talk more freely about those things were held in check. Cognitive science has opened the floodgates. The report of the Briefing Panel uses the following words, *all undefined*: intelligence, mind, mental operations, imagination, reasoning, induction, understanding, thinking, imagery, symbolic behaviour, and intended meanings. Cognitive scientists are enjoying an intoxicating freedom, but we must ask whether it is a productive one.

The situation is serious, and perhaps a touch of rhetoric will be forgiven. I shall model my conclusion on Émile Zola's famous charge in the *Dreyfus Affair*: '*J'accuse* . . . '

I accuse cognitive scientists of misusing the metaphor of storage. The brain is not an encyclopaedia, library, or museum. People are changed by their experiences; they do not store copies of them as representations or rules.

I accuse cognitive scientists of speculating about internal processes with respect to which they have no appropriate means of observation. Cognitive science is often only premature neurology.

I accuse cognitive scientists of emasculating the experimental analysis of behaviour by substituting descriptions of settings for the settings themselves and reports of intentions and expectations for behaviour.

I accuse cognitive scientists of reviving a theory in which feelings and states of mind observed through introspection are taken as the causes of behaviour rather than as collateral effects of the causes.

I accuse cognitive scientists, as I would accuse psychoanalysts, of claiming to explore the depths of human behaviour, of inventing explanatory systems which are admired for a profundity which is more properly called inaccessibility.

I accuse cognitive scientists of relaxing standards of definition and logical thinking and releasing a flood of speculation characteristic of metaphysics, literature, and daily intercourse, perhaps suitable enough for such purposes but inimical to science.

Let us bring behaviourism back from the Devil's Island to which it was transported for a crime it never committed, and let psychology become once again a behavioural science.

Acknowledgement

The author has profited greatly from extensive discussions with Dr Pere Julia concerning the position.

References

Broadbent, D. E. (1965). Perceptual defence and the engineering psychologist. *Bulletin of The British Psychological Society*, **18**, no. 60, 1–15.
Dickenson, A. (1980). Psychiatrists and traditional doctors in Nigeria. *Bulletin of The British Psychological Society*, **33**, 237–240.
Estes, W. K. (1972). Reinforcement in human behavior. *American Scientist*, **60**, 723–729.
Estes, W. K., Newell, A., Anderson, J. R., Brown, J. S., Feigenbaum, E. A., Greeno, J., Hayes, P. J., Hunt, E., Kosslyn, S. M., Marcus, M. & Ullman, S. (1983). Report of the Research Briefing Panel on Cognitive Science and Artificial Intelligence. *Research Briefings 1983*. Washington, DC: National Academy Press.
Fodor, J. A. (1981). The mind–body problem. *Scientific American*, **244** (1), 124–132.
Kahneman, D. & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, **39**, 341–350.
Onians, R. D. (1951). *The Origins of European Thought*. Cambridge: Cambridge University Press.
Shevrin, H. & Dickman, S. (1980). The psychological unconscious: A necessary assumption for all psychological theory? *American Psychologist*, **35**, 421–434.
Simon, H. (1980). The behavioral and social sciences. *Science*, **209**, 72–78.
Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, **57**, 193–216.
Skinner, B. F. (1957). *Verbal Behavior*. Englewood Cliffs, NJ: Prentice-Hall.
Skinner, B. F. (1968). *The Technology of Teaching*. Englewood Cliffs, NJ: Prentice-Hall.
Skinner, B. F. (1984). The shame of American education. *American Psychologist*, **39**, 947–954.
Stent, G. S. (1975). Limits to the scientific understanding of man. *Science*, **187**, 1052–1057.
Sutherland, S. (1978). *The Times Literary Supplement*, 1 September.

Requests for reprints should be addressed to Professor B. F. Skinner, Department of Psychology and Social Relations, Harvard University, 33 Kirkland Street, Cambridge, MA 02138, USA.