ULLIN T. PLACE

Connectionism and the Resurrection of Behaviourism[1]

The demise of behaviourism is traced to the advent of the serial-digital computer as a model for the functioning of the brain. With the advent of a new model in the shape of the parallel distributed processor (PDP) or connectionist network, the resurrection of behaviourism can be predicted. The relation between the two models is explained in terms of Skinner's (1966) distinction between "contingency-shaped" (modelled by the PDP) and "rule-governed" behaviour. Rule-governed behaviour in Skinner's sense is behaviour controlled by a verbal/symbolic "specification" of the relevant contingencies. The S-D computer is a device designed by a PDP (the human brain) to compensate for its own slowness and inefficiency in constructing and manipulating such symbolic specifications.

The serial-digital computer model of the brain and the demise of behaviourism

We are told that behaviourism is dead, or, if not actually dead, in the final throes of its death agony. But can we be sure? Behaviourism has suffered a serious reverse from the dominant position which it enjoyed in psychology, linguistics, the social sciences and even philosophy in the 1940's and 1950's; but, in my view, there is reason for thinking that the tides of fashion are at long last beginning to run once again in its favour.

In order to understand the basis for that judgment, we need to consider what it was that led to the demise of behaviourism in the first place. The *coup de grâce*

[1] Paper presented at the Christmas Meeting of the Experimental Analysis of Behaviour Group, Department of Psychology, University College, London, 17th December 1988 - revised August 1993.

was undoubtedly delivered by Noam Chomsky (1959) in his review of B.F. Skinner's (1957) book *Verbal Behavior* in *Language*; but it seems unlikely that Chomsky's review would have had the devastating effect which it undoubtedly did have, were it not for the fact that behaviourism in general and the behaviourist approach to language in particular was already being undermined by the serial digital computer as a model for the way thought, language and behaviour (or "action", as most cognitive scientists have been taught by the philosophers to say) are generated by the brain.

It is often supposed that what led psychologists and linguists to abandon behaviourism was its repudiation of private mental events. But this cannot be right. No behaviourist, to my knowledge, has ever denied the existence of private events. Some, the so-called "methodological behaviourists", have indeed denied the possibility of investigating such events scientifically. Skinner, on the other hand, has consistently repudiated methodological behaviourism. For the "radical behaviourist", as he calls himself, mental events exist and are available for scientific study through the verbal behaviour which they generate. He is inclined to deny that such events have any significant causal role in the control of behaviour generally; but even he has recently (Skinner 1987) had to concede that private events must at least have a causal role as discriminative stimuli relative to the verbal reports which describe them. For if they didn't, they wouldn't *be* reports of those events.[2]

But even if it were the case that behaviourists are denying the existence and causal role of private mental events, there is nothing, as far as I can see, in either Chomsky's review or in the serial-digital computer model of brain functioning which is readily construed as an argument for the view that such events exist and are causally effective with respect to human action.

As I see the matter there are four characteristic doctrines of behaviourism which are undermined by adoption of the serial-digital computer as model for the functioning of the brain:

(1) the behaviourist's repudiation of "mentalistic" language, in particular, the attribution to the behaving organism of what philosophers call "propositional attitudes", such as knowing or believing that so-and-so is the case,

[2] I have examined Skinner's claim that radical behaviourism accepts both the existence of private events and the possibility of studying them scientifically in my 'A radical behaviorist methodology for the empirical investigation of private events' (Place 1993).

(2) the behaviourist's contention that the principles of syntax and logic are abstractions constructed by grammarians and logicians reflecting the constraints imposed on learned patterns of sentence construction and inference behaviour by the demands of the verbal community for intelligibility in the one case and consistency in the other,

(3) the behaviourist's contention that, in the sense in which the principles of syntax and logic apply to it, thinking is just a matter of talking to oneself, of self-directed verbal behaviour on the part of the thinker, and that, consequently, as the late Sir Alfred Ayer put it in his Inaugural Lecture as Professor of Philosophy at University College, London, in 1947, "the process of thought can[not] be validly distinguished from the [linguistic] expression of it." (Ayer, 1947)[3]

(4) the view which is peculiar to the radical behaviourism of B.F. Skinner and which I myself believe to be mistaken according to which there is no place in psychology for explanations of behaviour in terms of the activity of the brain and central nervous system.

The Repudiation of mentalism

The behaviourist's repudiation of the use of mentalistic language in talking about the behaviour of living organisms for scientific purposes is made to appear rather silly in face of the fact that the practice of ascribing propositional attitudes to a digital computer, talking, in other words, about what the machine knows or has been programmed to think, is entirely natural and for certain scientific purposes unavoidable.

Syntax and logic as external constraints on learned verbal behaviour

The behaviourist notion that syntactic and logical principles are external normative constraints on the verbal behaviour of speakers conflicts with the generative role played by such principles in a digital computer and hence with the concept of the brain as a device constructed along these lines. For a digital computer cannot function

[3] I am indebted to Mr. L. Jonathan Cohen of Queen's College Oxford for reminding me of this quotation.

without a program. A program requires a programming language in which it is written. Moreover, the machine cannot respond to instructions written in the language, unless those instructions conform strictly to the syntactic principles which give the language its structure. Equally, a computer program will not function properly unless it conforms strictly to the principles of formal logic, as those principles are expressed in Boolean algebra. It follows that on the hypothesis that the brain is a digital computer, the principles of syntax and logic form an essential part of the very fabric of the software without which the device cannot begin to function. In the absence of any plausible counterpart for the process whereby programs are inscribed on magnetic tapes or disks which are inserted into and read by the device, it is much easier to construe the brain's counterpart of a program as "hardwired" or genetically pre-programmed than as abstracted *a posteriori* from the idiosyncratic practices of sentence construction and argumentation acquired through the vagaries of social reinforcement by the individual speaker, as the behaviourist would have us believe.

Language as communication precedes language as thought

Behaviourist doctrine holds that thinking is self-directed verbal behaviour, that verbal behaviour is primarily a system of interpersonal communication, one which is by and large exclusive to human beings, and, hence, that animals and pre-linguistic children cannot think in this way and, though they may perhaps be said to think in other ways, do not have the kind of thoughts to which the principles of syntax and logic apply. This doctrine conflicts with the digital computer model of the functioning of the brain which construes thinking as a computational process which *precedes* and is a prerequisite for the acquisition of linguistic competence by the speaker. It is true that the digital computer model requires an articulated programming language which Fodor (1975) calls "the language of thought" in which the controlling program is written; but such a language, assuming that it exists, is not and could not be a language in the sense in which that term is understood by the behaviourist, or for that matter the man-in-the-street, i.e., as a vehicle of interpersonal communication.

Skinner's repudiation of physiological psychology

Skinner's refusal to admit that there is a place within psychology for explanations of behaviour based on the workings of the brain and central nervous system generally (he allows such explanations as part of physiology), makes nonsense of the whole enterprise of artificial intelligence as a project within psychology and cognitive science, whether the model we use in developing our theory of how language, thought and behaviour are generated and controlled by the brain is a serial digital computer or some other device such as a parallel distributed processor. But, since Skinner's objection is not to the activity of explaining behaviour in terms of the brain, but rather to the disciplinary affiliations of those who perform that activity, there seems little reason either for taking it very seriously or for using it, as some are inclined to do, as grounds for ignoring everything else that behaviourists have to say about the environmental determinants of the behaviour of living organisms.

Skinner's distinction between "contingency-shaped" and "rule-governed" behaviour

Two developments have taken place, since the death of behaviourism and the advent of the cognitive revolution were generally proclaimed, the first of which has materially strengthened the case for the behaviourist position, while the second has substantially weakened the potency of the serial digital computer as a plausible model for the functioning of the brain.

The first of these developments is the introduction of the distinction drawn by Skinner in his 'An operant analysis of problem solving' (Skinner 1966) between what he calls "contingency-shaped" and "rule-governed" behaviour (Place 1988). This distinction needs to be understood in the light of Skinner's (1969) concept of the *three term contingency* which when stated in its more generalised form consists of

(A) the *antecedent* conditions which call for behaviour of some kind on the part of the organism in question,
(B) the *behaviour* to be performed, and
(C) the *consequences* of so behaving.

As conceived by Skinner, contingency-shaped behaviour is behaviour that is shaped and moulded by past encounters with the actual contingency, in other words, by personal experience of the *immediate* consequences in one's own case of doing one thing rather than another under the relevant antecedent conditions.

Rule-governed behaviour, on the other hand, is behaviour controlled by a verbal or other symbolic representation (or "specification", to use Skinner's term) of the contingency by which the organism believes it is confronted. Rules, in the sense in which Skinner uses that term, are of two kinds,

(a) *prescriptive rules*, like *If the baby cries, give it a bottle*, which specify an antecedent condition and the behaviour to be performed under that condition, and
(b) *descriptive rules*, like *If you give the baby a bottle, it will go back to sleep*, which specify the behaviour to be performed and the consequences that are likely to follow, whether immediately or at some later date, from doing one thing rather than another.

The importance of this distinction, as I see it, is that it re-asserts the traditional behaviourist view that the acquisition of competence as a receiver and speaker of a natural language of interpersonal communication precedes and is an essential pre-requisite for the control of behaviour by means of the process of representational thinking; but it does so in a way that draws attention to the difference between the way behaviour is controlled in organisms like animals and pre-linguistic children who lack the necessary linguistic skills and the way it is controlled when a linguistic or other symbolic representation is, as it were, interposed between the organism and its environment.

It suggests, moreover,

(a) that the only valid objection to the use of propositional attitude ascriptions or "mentalistic language", as the behaviourists call it, in the explanation behaviour is not that such language is intrinsically unscientific, but that it presupposes that the behaviour to be explained is "rule-governed" in Skinner's sense of that term and cannot, therefore, be legitimately applied to the explanation of behaviour that is contingency-shaped;
(b) that the reason why it makes good scientific sense to ascribe propositional attitudes to a serial digital computer is that it is a device which is designed *ab initio* to adjust to environmental contingencies in a rule-governed way,

> that is by constructing a symbolic representation of the environmental contingencies, rather than by reacting directly to them through the process of contingency-shaping as animals, including human beings, are designed by evolution to do.

There is now a substantial body of empirical evidence (Lowe 1979; 1983; Hayes 1989) which supports the distinction between the contingency-shaped behaviour of animals and pre-linguistic human infants on the one hand and the verbally controlled behaviour of older children and adults on the other. This evidence demands a radical revision both of the traditional behaviourist assumption that principles derived from the experimental study of animal behaviour can be applied directly to interpretation of human behaviour, and of the cognitivist's assumption that mentalist and computational principles which apply in the case of the "rational" behaviour of human adults can be extended to the behaviour of animals and pre-linguistic children and, *via* the hypothesis of the Unconscious Mind, to the automatisms and irrational behaviour of human adults.

Unfortunately, most psychologists had already given up reading the behaviourist literature by the time Skinner's 1966 paper and 1969 book were published. As a result, they were left with a stereotype of the behaviourist position which corresponds, in so far as it corresponds to anything, to a stage in the evolution of behaviourism which was already outdated long before Skinner introduced his distinction between contingency-shaped and rule-governed behaviour. In these circumstances, it is hardly surprising that Skinner's exposition of the distinction should have failed to make the impact it deserved to make outside the closed circle of his few remaining followers. That is why it has taken the second and much more recent development, the development of the parallel distributed processor as an alternative model for the way thought, language and behaviour are generated and controlled by the brain, to undermine the dominant position held by the serial digital computer model in the field of artificial intelligence and cognitive science, and thus open the door to a rehabilitation of the behaviourist standpoint.

Some defects of the serial digital computer as a model for the brain

There are a number of reasons for thinking that the serial-digital computer with which we are all familiar is not, in fact, a very good model for the functioning of the human or, for that matter, the animal brain:

(1) The serial-digital computer is designed to carry out quickly and efficiently computational tasks which the human brain performs slowly and inefficiently, if at all.

(2) Trained human and animal intelligence is characterised by its intuitive grasp of complex issues, such as those involved in visual space perception, which, as far as we know, does not depend on any kind of searching through lists of alternative possibilities. It is also much quicker and more efficient in performance of such tasks than a device, however powerful, which has to rely on this kind of systematic searching through lists of alternatives.

(3) The time taken by the activity in one neuron in the brain to excite another neuron adjacent to it is much too long for the brain to be able to run through the number of sequential steps it would need to run through in order to compute the solution to the kinds of problem it is able to solve in the time it takes to solve them, if it did in fact operate in the step-by-step manner that a serial-digital computer operates.

(4) The model of the brain as a serial-digital computer requires that data (information) be stored in one or more spatially located stores from which the data is retrieved as and when stipulated by the controlling program. No evidence for the existence of such a localised memory store in the brain has ever been forthcoming from studies of the way in which brain functioning is affected by lesions at different sites within the brain. Phenomena such as *retrograde amnesia* in which loss of memory for past events as a consequence of brain injury or damage is greatest for the most recent events in the individual's past history, with progressively less effect the further back in time the recollection extends and the more often the event in question has been recollected in the past, make it tolerably certain that the individual's ability to remember both facts and past events is a matter of "stamping in" connections widely distributed through the brain, rather than storing information in a localised memory store.

These considerations, however, failed to make any impact on the infectious enthusiasm with which the serial-digital computer model was pursued, so long as there was no alternative model in sight. It is this that has now changed.

The connectionist model of brain functioning

The term "connectionism" has recently been introduced to describe a theory of artificial intelligence which proposes that the correct model for understanding the way the mind-brain functions is not the serial digital computer with which we are all familiar, but rather the device known as "a parallel distributed processor" or PDP. Although the term "parallel distributed processor" was not then used to describe them, the first devices of this kind were constructed more than forty years ago in the very early stages of research in artificial intelligence before the serial digital computer had come into its own. In the 1940's and 1950's the object of the exercise was to construct an electronic device in which valves or later transistors are wired up in the way that the neurons of the brain are wired up in the form of a network through which a pattern of excitation is transmitted from input to output. Each unit fires or fails to fire depending on the input it receives or fails to receive from two or more units behind it in the network. If the neuron in question fires as a result of this excitation, it will in turn contribute either to the excitation or to the inhibition of two or more units in front of it in the network.

Now let us suppose that the properties of the units in such a network are such that each time the activity in a particular unit is excited or inhibited by the output from units anterior to it in the network, its susceptibility to that effect in the future is enhanced. It now turns out that a network of units arranged in this way displays properties which are remarkably like those of a living organism whose behaviour is controlled by a brain. In the 1940's and 1950's such devices were constructed simply in order to see how far systems of this kind could reproduce simple behavioural functions such as those of classical conditioning as studied by Pavlov (1927).

These early studies were rapidly overtaken from the 1960's by the development of the serial-digital computer model of brain functioning. The serial-digital computer model superseded the early neural network models because it seemed to offer a way of explaining much more sophisticated mental processes than anything of which the neural networks at that time appeared capable. The recent revival of interest in models based on the neural network principle has been motivated by the serious difficulties which were encountered in programming a serial-digital computer to perform what for a human being or an animal are relatively simple sensory discrimination or pattern recognition tasks, an ability which computers must acquire if robots are to take over the kind of routine inspection tasks currently performed by human operators. It turns out that if a parallel distributor processor is harnessed to the appropriate sensors, it can rapidly learn to recognise complex patterns of sensory information without having to check and without having to be specifically programmed to check a list of alternative possibilities in the way the serial-digital computer does. It can do this, moreover, without being defeated, in the way the serial-digital computer invariably is, by a familiar pattern presented in a way which the system has never previously encountered and with which it has not been specifically programmed to deal, in other words, problems such as that presented by a familiar scene viewed from an unfamiliar vantage point, or at a very different season of the year from that which obtained when it was previously encountered.

Connectionism and behaviour analysis

It is my belief that if only the behaviour analyst can be persuaded to give up what I personally regard as the absurd notion that we can study behaviour scientifically, without ever addressing the question of how that behaviour is generated by the brain,[4] we shall find that there is a natural affinity between connectionism and Skinnerian behaviour analysis. For if we look at the molar performance characteristics of a parallel distributed processor, we see that it is a device which, when given the appropriate sensors, outputs and feedback information (in other words the appropriate selective reinforcement), has the ability to learn an operant discrimination task in the manner described by Skinner (1938) in the case of living organisms like the pigeon and the rat.[5] The PDP, in other words, is a device which learns to adapt to the environment to which it is given access by its sensors by the process of "contingency-shaping" to which reference has already been made in which future behaviour is moulded by the immediate consequences of emitting similar behaviour in the past. The serial-digital computer, by contrast, is a device which is designed

[4] This is, perhaps, another case, alongside that of his arguments against mentalism (Place 1987), where Skinner's "gut intuition", as Dennett (1978) has called it, has led him to the right answer for the wrong reasons. His refusal to countenance cognitive psychology (Skinner 1977), on the grounds that speculation about what is going on in the brain belongs to physiology rather than psychology, has at least prevented behaviour analysts from following the crowd down the blind alley of the serial-digital computer model.

[5] A perfect example of operant discrimination learning in a connectionist network, for which I am indebted to Stephen T. Mills of the University of Ulster, is the case of the Gorman-Sejnowski network learning to discriminate between mines and rocks on the seabed by their distinctive sonar echo, as described by Paul Churchland in the 1988 edition of his *Matter and Consciousness*, pp.157-162.

from the outset to perform quickly and efficiently those computational tasks which the brain, because it operates as a PDP, doesn't do very well, but which it needs to do, or have done for it, in order to adapt to its environment by constructing what Skinner in his 1966 paper calls "a rule."

We have seen, that in that paper, Skinner defines "a rule" as a "contingency-specifying stimulus", in other words as a symbolic representation of the "contingencies" or behaviour-consequence relationships which are implicated in the situation by which the behaving organism is currently confronted. For Skinner, moreover "a symbolic representation," not that he uses that term, is to be understood as a representation in terms of the words and other symbols which go to make up a human natural language. It follows that, for him, the ability to construct symbolic[6] representations of environmental contingencies is an ability which, from an evolutionary perspective, has been acquired only very recently with the development of the capacity for linguistic communication, and is, therefore, an ability to which the human brain has not yet had the time to become fully adapted. In other words, the human brain is PDP device which has been finely tuned by millions of years of evolutionary history to generate behaviour in accordance with the principles of contingency-shaping. Its recently acquired ability to manipulate symbols and govern the behaviour of the organism by means of the resulting rules appears slow and cumbersome by comparison with a machine like the serial digital computer which is specially designed for this purpose.

Connectionism and the learning of syntactic and logical principles

The suggestion that the human brain is a PDP device which has only very recently acquired the ability to construct symbolic representations of its environmental contingencies, as a side-effect of acquiring the ability to communicate with others by means of language, has important implications for the critique of connectionism

[6] Symbolic' here as opposed to 'iconic'. There is evidence, particularly from the phenomenon of REM sleep, that mammals, at least, have the ability to construct iconic representations or *mental images* of events and states of affairs which are not currently impinging on their sense organs. However, a system of iconic representation which is not supported, as it is in the human case, by a symbolic system is severely restricted in its ability to represent environmental contingencies where there is an appreciable time lag, either between antecedent and required behaviour or between behaviour and its consequence. There is also no way, in absence of linguistic communication, in which such iconic representations can be enriched through information supplied by others.

by Fodor and Pylyshyn which appeared in the journal *Cognition* in 1988. While conceding that connectionism provides a better model for sense perception than does the serial computer, Fodor and Pylyshyn claim that the now traditional computational model provides a theory of the way language and thought are guided by syntactic and logical principles which is far superior to anything the connectionist model can currently supply or looks remotely like being able to supply in the future.

What this argument overlooks, it seems to me, is that the principles of syntax and logic play a quite different role in the serial computer from that which they play in human thought and language. Because the serial computer is a device which is designed *ab initio* to construct and manipulate symbolic representations, the principles of syntax and logic are "hardwired" into the fabric of the machine language in which its basic programs are written. In such a device, syntactic and logical principles have a direct causal role in the generation of thought and language. The human brain, on the other hand, if we are right in thinking that it operates as a parallel distributed processor, only encounters syntactic and logical principles when it is faced, in the case of syntax, with the problem of constructing sentences in natural language, and in the case of logic, with the problem of making inferences from one feature of a symbolically represented situation to another in the absence of the cues provided by the presence of the situation in question in the organism's stimulus environment, inferring, for example, from the information that *A* is larger than *B* and that *B* is larger than *C* to the conclusion that *A* is larger than *C*, without needing to compare *A* and *C* directly.

Not only is there reason to think that there is no role for either syntax or logic before linguistic competence is acquired; there reason to think that when syntax and logic eventually appear, they are learned by the child, not as *internal* principles which govern the generation of thought and language, but as *external* constraints to which the brain must learn to conform (though without needing to formulate those constraints in words) in order to secure the appropriate reinforcement from the listener in the form of an acknowledgment of successful communication. The ubiquity of these acknowledgments of successful communication (*Mmhmm*'s, headnods, etc.) which are an invariable accompaniment of person to person vocal interaction is a phenomenon which only makes sense, so it seems to me, on the assumption that its function is to maintain conformity to linguistic convention by selectively reinforcing intelligible utterances and either withholding reinforcement or actively punishing (by compelling the speaker to repeat) unintelligible ones. This process, moreover, proceeds automatically with minimal awareness on the part of both of the listener in providing these reinforcers and of the speaker in receiving them. It is a classical case of the automatic contingency-shaping of behaviour by its immediate consequences.[7]

Like the human brain, the parallel distributed processor will never be able to emulate the serial computer as a device for constructing and manipulating formal symbolic representations. But if I am right in thinking that it functions as a contingency-shaped discrimination learning device, it seems likely that it will eventually provide us with a much more accurate model of the way such tasks are actually handled by the brain than the serial computer model can ever hope to do. As the connectionist model of the acquisition of syntactic, semantic and logical competence unfolds, I predict that the reliance of the model on the behaviour analytic principle of the contingency-shaping of verbal behaviour by the acknowledgment of successful communication on the part of the listener will become more and more apparent.

In conclusion, by courtesy of my cognitivist colleague at the University College of North Wales, Dr.Gordon Brown, I would like to quote from a newsflash which he received through an electronic bulletin board known as "Neuron Digest". The newsflash is dated the 4th of December 1988. It reads:

> **"Neural Network Learning Runs into Heavy Flak at McGill**
> While on a visit to Montreal, neural network theorist James McClelland recently presented modifications to neural network theory that suggest a network can be successfully used to learn syntactic and semantic relations which subsequently serve to discern the meaning of an utterance. This departure from "hardwired" syntax and semantics drew heavy criticism from the McGill audience of cognitive psychologists and linguists who accused him of Skinnerianism ("subject to every single one of Chomsky's critiques launched against Skinner back in the fifties"). However, McClelland stood his ground, arguing that the good recognition performance and the various appealing behaviors of the system fully vindicate the neural network approach. He did admit — publicly — that the exceedingly large number of learning trials (630,000) was a bit worrisome, and — privately — that the 32-bit-precision thresholds used to make semantic and syntactic decisions at the recognition stage was unrepresentative of probable functioning in the noisy neural context of the human brain."

With regard to the 630,000 trials needed for the PDP to acquire linguistic competence, I don't know if anyone has calculated how many reinforcements in

[7] Experimental evidence of the behaviour-shaping effect of these "response tokens" or "back channels," as they are called in the literature of conversation analysis, is provided by a well known study by Joel Greenspoon (1955) and, in a more natural conversational setting, by Bill Verplanck (1955).

the form of head-nods, *MmHmm*'s, etc., each speaker receives in the course of an average conversation; but the number must be very large,[8] and that's only maintaining an already existing competence, not acquiring a new one!

References

Ayer, A.J. (1947) *Thinking and Meaning: Inaugural Lecture*. London: H.K. Lewis.

Chomsky, N. (1959) Review of B.F. Skinner's *Verbal Behavior*, *Language*, 35, 26-58.

Churchland, P.M. (1988) *Matter and Consciousness*, Revised Edition, Cambridge, Massachusetts: MIT Press.

Dennett, D.C. (1978) *Brainstorms: Philosophical Essays on Mind and Psychology*. Montgomery, Vermont: Bradford Books.

Fodor, J.A. and Pylyshyn, Z.W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Fodor, J.A. (1975) *The Language of Thought*. New York: Crowell.

Greenspoon, J. (1955) The reinforcing effect of two spoken sounds on the frequency of two responses. *American Journal of Psychology*, 68, 409-416.

Hayes, S.C. (Ed.) (1989) *Rule-Governed Behavior: Cognition, Contingencies and Instructional Control*. New York: Plenum

Lowe, C.F. (1979) Determinants of human operant behaviour. In M.D. Zeiler and P. Harzem (eds.) *Reinforcement and the Organisation of Behaviour*. Chichester: Wiley, pp.159-192.

Lowe, C.F. (1983) Radical behaviourism and human psychology. In G.C.L. Davey (ed.) *Animal Models and Human Behaviour*. Chichester: Wiley, pp.71-93.

Pavlov, I.P. (1927) *Conditioned Reflexes*. English Translation by G.V. Anrep. London: Oxford University Press.

Place, U.T. (1987) Skinner re-skinned. In S. and C. Modgil (eds.) *B.F. Skinner: Consensus and Controversy* Lewes: Falmer Press, pp. 239-248.

Place, U.T. (1988) Skinner's distinction between rule-governed and contingency-shaped behaviour. *Philosophical Psychology 1*, 225-234.

Place, U.T. (1992) Eliminative connectionism and its implications for a return to an empiricist/behaviorist linguistics. *Behavior and Philosophy*, 20, 21-35.

Place, U.T. (1993) A radical behaviorist methodology for the empirical investigation of private events. *Behavior and Philosophy*, 20, 25-35.

Skinner, B.F. (1938) *Behavior of Organisms*. New York: Appleton-Century.

Skinner, B.F. (1957) *Verbal Behavior*. New York: Appleton-Century- Crofts.

Skinner, B.F. (1966) An operant analysis of problem solving. In B.Kleinmuntz (Ed.) *Problem Solving: Research, Method and Theory*. New York: Wiley.

Skinner, B.F. (1969) *Contingencies of Reinforcement*. New York: Appleton-Century-Crofts.

Skinner, B.F. (1977) Why I am not a cognitive psychologist. *Behaviorism*, 5, 1-10.

[8] Evidence cited in a recent paper (Place 1992) suggests that the listener tends to supply a reinforcer whenever the speaker completes a sentence.

Skinner, B.F. (1984) An operant analysis of problem solving. In A.C. Catania and S. Harnad (Eds.) *The Canonical Papers of B.F. Skinner. The Behavioral and Brain Sciences*, 7, 583-591.
Skinner, B.F. (1987) Outlining a science of feeling. *The Times Literary Supplement* 490, 501-2.
Verplanck, W.S. (1955) The control of the content of conversation: reinforcement of statements of opinion. *Journal of Abnormal and Social Psychology*, 51, 668-676.