

[Place, U. T. (1999d). Connectionism and the problem of consciousness. *Acta Analytica*, 14(22), 197-226. There are some additions added by the author after publication.]

ULLIN T. PLACE

CONNECTIONISM AND THE PROBLEM OF CONSCIOUSNESS¹

This paper falls into three parts. In Part 1 I give my reasons for rejecting two aspects of Horgan and Tienson's position as laid out in their book, the language of thought and belief-desire explanations of behaviour, while endorsing the connection they see between linguistic syntax and the syntax of a motor skill. In Part 2 I outline the theory that the brain consists of two input-output transformation systems *consciousness* whose function is (a) to categorise problematic inputs, (b) to select a response appropriate to such inputs once they have been categorised and (c) to initiate and monitor the execution of such response once selected, and the "*zombie-within*" whose function is (a) to identify and alert consciousness to any inputs that are problematic either because they are unexpected or because they are significant relative to the agent's current or perennial motivational concerns. In Part 3 I consider how far the properties of the two systems outlined in Part 2 can be understood in terms of the known properties of connectionist networks.

1.1 Disagreements and Agreements with Horgan & Tienson - Connectionism and the Digital Computer Model

To my way of thinking, the great virtue of the connectionist network as a model for the workings of the brain is that it allows us to discard once and for all and lock, stock and barrel the notion that the brain functions in the manner of a digital or von Neumann computer. The digital computer is a device constructed by human brains whose function is to perform quickly and efficiently tasks which, without this assistance, the best human brain performs slowly and inefficiently. Equally, tasks such as recognising the same underlying pattern in a variety of different contexts at which the brain excels are performed poorly and inefficiently, if at all, by the digital computer. In the light of these facts, the suggestion that we can understand the functioning of the brain on the model of a digital computer is grotesquely implausible. By contrast, not only does the connectionist network have an anatomical structure which, at the neuro-synaptic level, is modelled on that of the brain itself, its functional properties resemble those of the brain. Networks, unlike the digital computers which are used to implement them, are good at tasks such pattern-recognition which the brain is good at and the digital computer, unless specifically programmed to operate as a connectionist network, is poor at. Like the brain, networks are poor at the symbol-manipulation tasks at which the digital computer excels. Moreover, in contrast to the digital computer which stores and retrieves information in a way that is demonstrably quite unlike the way the brain remembers things, connectionist networks can learn and, when they do, they learn in a way that resembles the way living organisms learn both with respect to the conditions under which learning takes place (when the system is told when it is right and when wrong) and with respect to the microstructural changes involved (strengthening and weakening the "weights" of synaptic connections within the network).

In the light of these considerations it is surprising, not only that the idea that the digital computer could be a model for the brain should have taken hold as it did, but that having taken hold, it still exercises a surprising degree of influence over the thinking of connectionists most of whom, it has to be said, learnt their trade during the period when the digital computer was the dominant model. There are connectionists who subscribe to what I regard as the absurd doctrine of "implementational

¹ Presented at a Conference on T. Horgan & J. Tienson *Connectionism and the Philosophy of Psychology*, Ljubljana, Slovenia, 15th August 1997.

connectionism" (Pinker and Prince 1988) which holds that the brain is a connectionist network which *implements* a digital computer in much the same way that almost all artificial connectionist networks in current use are implemented on digital computers. Other connectionists, of whom Horgan and Tienson are typical see the connectionist network model replacing the digital computer model, but nevertheless, retain features from the latter which, to my way of thinking, they should have abandoned. I can perhaps see things like this rather more clearly than most, because I was never seduced by the digital computer model for the brain in the first place and am old enough to remember how things were before it appeared.

1.2 Retentions by Horgan & Tienson - The Language of Thought

There are two features from the digital computer model which Horgan and Tienson retain and which I think ought to be discarded (a) the language of thought hypothesis and (b) belief-desire type explanations of behaviour. That the language of thought hypothesis is a relic of the digital computer model is clear enough. Although it has a medieval anticipation in the shape of William of Ockham's "Mental" (Geach, 1957, p. 102) or "mental language" (Trentham 1970), as introduced by Fodor (1975) in his book of that name, "the language of thought" was part of a thoroughgoing computational theory of the functioning of the mind/brain. The notion was evidently inspired by the so-called "machine language" which any digital computer must be able to "understand" or, in other words, to which it must be able to respond appropriately, if it is to respond to the instructions and store in its working memory the data as they are presented to its reading head. In terms of Searle's (1980) Chinese Room analogy, whereas the operator need not understand the meaning of the characters he is putting up for display, he *must* be able to understand the instructions he is given as to which combination of characters he is to display at any moment of time. Fodor's language of thought is the language in which those instructions are written. The ability to "understand" them is something that must be "hard-wired" into the device at its construction. There is no evidence for the existence in the brain of anything like it or the tape-reading head it implies. Nor, given the availability of connectionist models which dispense with them entirely, is there any need to postulate such things.

1.3 Retentions by Horgan & Tienson - Belief-Desire Explanations

Since belief-desire explanations of behaviour were around long before the advent of digital computer model for the brain, the suggestion that Horgan and Tienson's retention of such explanations in their theory of behaviour is another hangover from the digital computer model for the brain may seem surprising. But it should be remembered that before the introduction of the digital computer model and the so-called "cognitive revolution" which it inspired, psychology, in the United States at least, was dominated by behaviourism which condemned belief-desire explanations of behaviour as unscientific mentalism. It was this ban on the use of belief-desire explanations which more than anything else contributed to the unpopularity of behaviourism, particularly amongst clinicians who need to use and respond to explanations of this type in their dealings with patients and clients. By the same token, it was the restoration of belief-desire explanations to scientific respectability through the fact that computer scientists were using these "mentalist" explanations to describe what their machines were doing which contributed mightily to the popularity of the cognitive alternative. With the advent of connectionism, the behaviourist repudiation of belief-desire explanations has returned. Only now it takes the form of the eliminativist materialist's (Churchland 1981) contention that what they call "folk-psychology" is a pre-scientific theory which is quite simply false and should be discarded in favour of a purely physiological or connectionist explanation of behaviour.

1.4 Objections to Belief-Desire Explanations in Psychological Science

It is my view that the behaviourists, their cognitive opponents and the eliminative materialists have all misunderstood the nature and problem of belief-desire explanations of behaviour. In the first place it is seldom appreciated that such explanations are of a type, very common in science, where two

dispositional properties combine to determine their joint manifestation. The explanation of the magnitude of current flow within a conductor in terms of a combination of the resistance of the conductor and the potential difference between its two ends - Ohm's Law in other words - is an example. The problem with belief-desire explanations is neither their dispositional character, nor the fact that you cannot predict behaviour from a knowledge of beliefs alone or desires alone, nor their supposed implicit reference to inner and objectively undetectable thought processes. It is just that they characterise the dispositional determinants of behaviour on the assumption that there is a consistent and rational connection between what agents say about the situation confronting them when they have no motive to dissemble and what they subsequently do when confronted with a situation of that kind.

Such explanations are unacceptable for scientific purposes for the following reasons:

- They describe a genuine causal relation within the organism's behaviour only in those cases where the way agents represent the situation to themselves in language does, indeed, determine how they behave. Consequently, such explanations have no legitimate application to the behaviour of pre-linguistic organisms, such as infants and animals. Although dolphins (Herman 1987; Herman, Kuczaj & Holder 1993) and sea lions (Schusterman & Krieger 1984; Schusterman & Gisiner 1988) have been taught to respond to simple sentences and although the African Grey parrot, Alex, (Pepperberg 1990) and the bonobo chimpanzee, Kanzi, (Greenfield & Savage-Rumbaugh 1990) have been taught to produce them, the sentences in question are either imperatives or questions asking for the names of things. Declaratives of the kind that express beliefs are conspicuous by their absence. There are also many aspects of human behaviour for which such explanations are scientifically inappropriate. A case in point is Kathy Wilkes' (1984, p. 228) example where the behaviour of tripping at the top of a flight of stairs is attributed to the belief that there was one more tread than there actually was.
- Even if one were to grant, which I would not, the existence of sentences in a language of thought controlling both the behaviour of animals and those aspects of human behaviour which are not controlled self-directed sentences in natural language, the whole purpose of this explanatory form, namely to allow us to predict the behaviour of others from what they say would be defeated, since no one, to my knowledge, has ever even claimed to have recorded and translated a sentence in such a language, let alone had that claim substantiated.
- There is reason to think that the social stigma that attaches to the insane reflects a system of social reinforcement and disinforcement practices within society whose function is to ensure the maintenance of a consistent and rational relation between what people say and what they do on which our ability to predict the behaviour of others depends. If this is correct, it implies that belief-desire explanations, so far from representing a naturally-occurring physiologically-based mechanism controlling human behaviour, depend for their predictive and explanatory success on a social convention maintained by a trial-and error-correcting learning process of a type which is observed in animals and which can be readily replicated in a connectionist network by means of the principle of back-propagation.

1.5 Agreement with Horgan & Tienson - The Syntax of a Motor Skill

Having identified those aspects of Horgan and Tienson's book which I reject and explained why I do so, I must hasten to make clear that there is another aspect of their work which I warmly endorse, but which for reasons which should now be obvious, I would want to extract from the matrix in which it is currently embedded - the language of thought hypothesis. I refer to the notion that just as the syntax of natural language allows the speaker to put words and phrases together so as to construct indefinitely many sentences intelligible to the listener, the like of which neither party need have encountered before, so, they suggest, there is a syntax of action which allows an agent to combine movements into a coherent and successful action. Horgan & Tienson suggest that both forms of syntax are basically linguistic in nature, both are outgrowths of that hypothetical syntax that regulates the construction of sentences in the language of thought. I beg to differ. Since I reject the language of thought, there can be for me no underlying syntax of which the syntax of natural language and the

syntax of action can be outgrowth. On the other hand, since I accept the analogy between the syntax involved in sentence construction and that involved in the exercise of a motor skill, as in Horgan and Tienson's example of basket-ball playing, for me it is the syntax of skilled movement which is primary and the one that develops first in the evolutionary sequence. The syntax of natural language is a secondary outgrowth of the syntax that is needed to combine movements into a coherent and successful action.² That, after all, is what the Greek word [σύνταξις] originally meant, in the words of the Oxford English Dictionary, an "Orderly or systematic arrangement of parts or elements."

2 Consciousness and The Zombie-within

In the paper (Place, 1997) which I presented at the Inaugural Conference of the Association for the Scientific Study of Consciousness at Claremont, California in June 1997, I traced the evolution of the theory I shall present in what follows from two papers, 'The concept of heed' and 'Is consciousness a brain process?' which I published in the *British Journal of Psychology* in 1954 and 1956 respectively, through Donald Broadbent's (1958) *Perception and Communication* and *Decision and Stress* (1971), Nick Humphrey's (1974) 'Vision in a monkey without striate cortex', Larry Weiskrantz's (1986) *Blindsight*, Cowey and Stoerig's (1995; 1997; Stoerig & Cowey, 1997) demonstration of blindsight in monkeys, some recent work on the processes of selective attention by the philosopher Paul Snowdon (1995), the experimental neuropsychologist, Hal Pashler (1991; 1997), and the neuroscientist, Michael Posner (Posner & Petersen, 1990; Posner & Dehaene, 1994), finally the neurological work of Martha Farah (1990) and Milner and Goodale (1995) on the functions of the so-called "dorsal and ventral streams" in the visual system in the brain.

The picture that emerges from these various strands of evidence is of two parallel, but complementary and continuously interacting input-to-output transformation systems in the brain which I shall refer to respectively as 'consciousness' and 'the unconscious automatic pilot' or 'zombie-within'.

2.1 Consciousness as an Input-Output Transformation System

On this hypothesis, *consciousness* is a "*limited capacity channel*" (LCC) or rather a sequence of three such channels which has four sequentially ordered functions:

- (1) to *categorize* on the basis of what Broadbent (1971) calls the "evidence" and which I equate with conscious experience any input that is identified by the zombie-within as *problematic*, in that it is either *unexpected* or *motivationally significant*, i.e., significant relative to the individual's current or perennial motivational concerns (LCC 1),
- (2) to *react emotionally* to inputs which have been identified as problematic, both before ("physical" pleasure/pain) and after they have been categorized ("mental" pleasure/pain), thereby ensuring that the subsequent processes of response-selection and response-execution are brought into an adaptive relation to the individual's current and perennial motivational concerns,
- (3) to *select a response* appropriate both to the presence of a thing of that kind and to the individual's motivational concerns with respect to it (LCC 2), and
- (4) to initiate and monitor the *execution* of the response selected (LCC 3).³

² [Footnote added after publication.] Since these words were written, Rizzolatti and Arbib (1998) have provided formidable empirical evidence of the link between motor skill particularly that involving control and interpretation of hand movements, and the syntactic organisation of linguistic utterances. What has been shown by research at the Parma Institute of Human Physiology is that in an area of the monkey pre-motor cortex which is homologous with Broca's area in humans there are neurons which have been called "mirror neurons" which respond both when the individual is using visual feedback to its own movements and when it is watching the same movement being performed by another monkey or by a human being. This observation not only confirms the suggestion which is supported by a great deal of other evidence that vocal language was preceded in the course of evolution by a language of gesture, it also makes sense of the evidence which shows that Broca's area mediates both the production of syntactically organised sentences by the speaker and the interpretation of syntactically organised sentence utterances produced by others.

³ There is reason to think that in the human brain there are two more limited capacity channels, one which selects the name assigned to

The evidence suggests that although much of what goes on is unconscious (in the sense that the details are not available to be described or reported by the human subject), the whole of the cerebral cortex in mammals is devoted to the implementation of consciousness *in this functional sense*. In general it would seem that what the human subject reports are the *outcomes* of the processes of selective attention, categorization, emotional reaction, response-selection and response-execution, rather than the processes themselves. The exceptions to this rule are the process of *sensory conscious experience* which can, to a limited extent, be described independently of the way it is categorized, and the thoughts (images and sub-vocal speech) which contribute to, but do not exhaust, the process of response-selection, just as conscious experience of the feedback from the output as it develops contributes to, but does not exhaust, the process of response-execution.

2.2 The Unconscious ‘Automatic Pilot’ or ‘Zombie-within’

The functions of the *unconscious automatic pilot* or *zombie-within* are

- (1) to continuously *scan* the total current input so as to alert consciousness to any input it identifies as *problematic*,
- (2) to *protect consciousness from overload* either by ignoring those non-problematic inputs which require no response or by responding appropriately, but automatically and *without categorization*, to those for which there already exists a well practised skill or other ‘instinctive’ response pattern.

Like its namesake in popular mythology, the zombie-within is a creature of habit, routine and unquestioning conformity to the instructions it receives from consciousness. Anything out of the ordinary is immediately passed on for processing by consciousness. The one respect in which it differs from the traditional picture of its mythical namesake is in its capacity to learn from experience, limited though that is to the progressive shaping of minor variations in behaviour by their immediate consequences.

2.3 Interactions between the Two Systems

Although, as the blindsight phenomenon shows, there are other forms [of] interaction between the two systems, the three most important interactions between consciousness and the zombie-within are

- the action of the zombie in alerting consciousness in general and conscious experience in particular to problematic inputs,
- the gradual transfer to the zombie-within of stimulus-stimulus expectations and stimulus-response connections formed within consciousness as they become habitual (for PET scan evidence of this process, see Raichle, *et al.*, 1994), and
- the integration of the two systems in a well-developed motor skill where, as the syntactic organization of movement becomes increasingly automatised, i.e., gets taken over by the zombie, the easier it becomes for the mechanisms of selective attention to ensure that consciousness is focused on those aspects of the task that are crucial from the point of view of effective strategic decision-making and the timely initiation of the selected response.

2.4 Modules within Consciousness and the Zombie-within

the concept or category in question by the natural language spoken by the individual concerned and another which constructs a syntactically articulated sentence appropriate to the ‘thought’ which emerges from a further stage in the categorization process.

The multiplicity [of] functions identified within both consciousness and the zombie-within imply a multiplicity of modules within both systems. Figure 1 shows the arrangement of these modules as I currently construe it. As you will see, the diagram shows the output from the sense organs splitting into two streams, *consciousness* on the left, the *zombie* on the right.

The zombie is shown as consisting of four functionally defined modules:

- (1) the *problematic input detector (PID)* which separates inputs into problematic and non-problematic on the basis of relatively coarse criteria of unexpected/expected and motivationally significant/insignificant, and transmits the former *via*
- (2) the *involuntary attention-focuser (peripheral)* which mobilises and directs movements of the head eyes and body so as to bring the *source* of the problematic input within the range of all relevant sense organs, and
- (3) the *involuntary attention-focuser (central)* which attracts the focus of *conscious experience* to that part of the sensorium where the problematic input is located, while either ignoring or routing non-problematic inputs to output *via*
- (4) the *automatic pilot*.

Consciousness, as shown here, consists of three sequentially ordered limited capacity channels (LCCs), concerned respectively with *input-categorization (perception)*, *response-selection* and *response-execution*. The concept of the *limited capacity channel* comes from Broadbent (1958) whose 1971 book *Decision and Stress* restricts its application to the process of input categorization. The evidence that there is more than one such channel in the brain comes from Pashler (1991; 1997). In order to protect it from overloading, Broadbent's model requires that each LCC be provided with a system of subordinate modules, including

- (a) an "ingate" which *controls access* to the LCC,
- (b) a *buffer* or short term memory store in which inputs waiting to obtain access to the LCC are held until the ingate clears,
- (c) an *attention-focuser* which maintains the focus of attention on the task in hand until it is satisfactorily completed, and
- (d) an *outgate* which allows the attention-focuser to open the ingate to a new input once the previous information processing task has been satisfactorily completed.

In the case of the input-categorizing LCC the limited capacity channel is shown as divided into two separate modules, *conscious experience* which, to use Broadbent's (1971) term, provides the "evidence" on which categorization is based and without which the blindsighted subject's judgments become "pure guesswork" and the *categorization* response itself.

Lying outside this system of three limited capacity channels are two modules, the *analyzer* and the *emotion servo*. The existence of the analyzer is suggested by the known functions of the visual cortical areas V1-V5. Recordings from individual cells in these areas have revealed what are known as "feature detectors", cells which fire in response to the presence within the current retinal input of various features and patterns which are relevant for the identification of objects and situations in the organism's visual environment. The features which are detected in this way become more and more abstract and involve responding to activity spread over a wider and wider area of the visual field the further they are

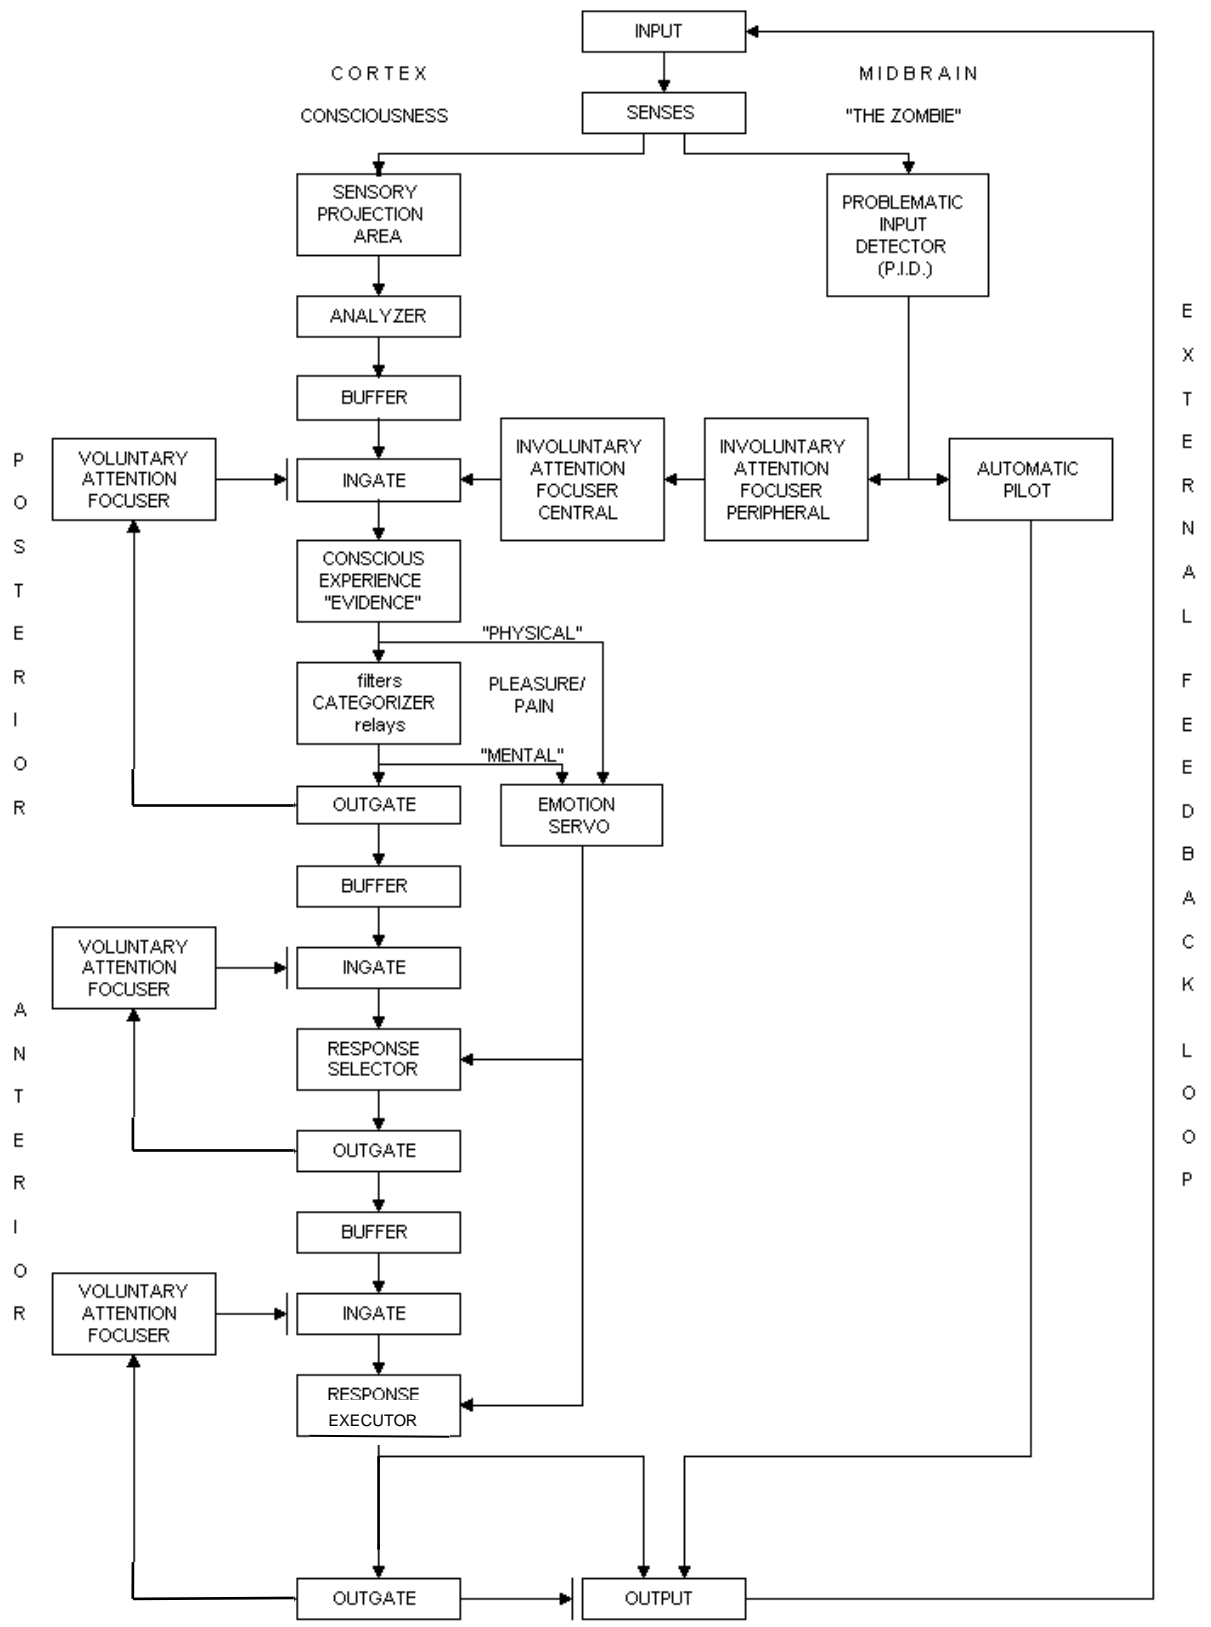


Figure 1: Consciousness and the Zombie-within - Suggested Layout of Modules

from V1. Recent research by Steve Luck and Nancy Beach (1996) confirming the Feature Integration Theory (FIT) of Ann Treisman and her colleagues (Treisman, 1988; Treisman & Gelade, 1980; Treisman & Gormican 1988) suggests that the effect of the module here identified as *conscious experience* is to "bind" the "information" provided by these individual feature detectors into a single *Gestalt*, thereby generating the "evidence" on which categorization or interpretation of the input is subsequently based.

Whereas the analyzer is construed as a module which precedes and prepares the ground for the process of input-categorization, the *emotion servo* is brought into play by the action of the input-categorization LCC. This activation occurs both before categorization in response to "raw" uninterpreted conscious experience, as in the case of "physical" pleasure and pain, and after categorization, as in the case of "mental" pleasure and pain. Its function is to provide motivation both for response-selection and for response-execution.




Figure 2: The Ventral and Dorsal Streams (After Milner & Goodale, 1995, Figure 3.1, p. 68)

2.5 Locating these Modules within the Brain - The Ventral and Dorsal Pathway

Although this is a task which I approach with considerable trepidation in view of my limited expertise in the field of neuroscience, evidence from a variety of sources, neurology, electrophysiology, and, most recently, the newly discovered brain-imaging techniques, has made it possible to propose a tentative identification of the modules shown in Figure 1 with specific anatomically defined structures within the brain.

For the present I shall confine my remarks on this score to the "upstream" portion of Figure 1, that which precedes the transition within consciousness from conscious experience to categorization. In this area recent psychological research on visual agnosias (Farah, 1990; Milner and Goodale, 1995) has drawn attention to the functional significance of an anatomically identified bifurcation within the brain between two "streams", the *ventral* stream and the *dorsal* stream. As originally defined by Ungerleider and Mishkin (1982) [and as illustrated on Figure 2], these two streams bifurcate downstream of the striate or primary visual cortex (V1). The ventral stream travels *via* the extra-striate visual areas (V2-V5) to the *infero-temporal cortex*. The dorsal stream travels upwards to

"terminate" in the *posterior parietal cortex*. In other words, the bifurcation between the two streams lies entirely within the cerebral cortex.

Studies of the behaviour of patients with lesions restricted to one or other of these two streams show that lesions of the ventral stream result, depending on the site and extent of the lesion, in a variety of functional disorders involving the loss or disturbance of visual conscious experience associated with a loss or disturbance of the ability to recognise objects and the situations in which they occur, conditions such as *prosopagnosia* (loss of the ability to recognise faces) and *simultanagnosia* (loss of the ability to recognise the relations between multiple objects in a visually presented scene). Lesions of the dorsal stream, on the other hand, result in disturbances of the visual control of voluntary movement.

[As Milner and Goodale (*op.cit.*, p. 67) concede,] since the two streams bifurcate downstream of the striate cortex (V1), they cannot be invoked to explain the phenomena of 'blindsight', i.e., the visual functions that survive lesions of V1. What is needed to explain *that* is a pathway leading from the retina which by-passes V1 and leads to the posterior parietal cortex where the characteristic movements of reaching for an unseen object are generated and controlled. Just such a pathway is shown in the upper part of Figure 2. I call this pathway the "*sub-cortical (S-C) to dorsal⁴ pathway*". For although it is shown for convenience on Figure 2 above both the dorsal *and* ventral streams properly so-called, in fact, until it reaches its "destination" in the posterior parietal cortex, it is composed of structures (the *superior colliculus* and *pulvinar*) which lie *below* the cortex in the midbrain and thus below both dorsal *and* ventral streams.

As is apparent from Figure 2, the S-C to dorsal pathway branches off at the retina from what we may call the "*ventral pathway*" (to distinguish it from the ventral stream which forms part of it) consisting of the *lateral geniculate nucleus*, the *primary visual cortex* (V1), the *ventral stream* (V2-V5) and the *infero-temporal cortex*. The two pathways so-defined differ in two respects:

- Whereas, apart from the lateral geniculate nucleus, all the structures composing the ventral pathway are in the cortex, all the structures composing the S-C to dorsal pathway, apart from its final "destination", the posterior parietal cortex, are sub-cortical.
- Unlike the ventral pathway all of whose component structures apart from its "destination", the infero-temporal cortex, are concerned only with the processing of *visual* information, all the structures composing the S-C to dorsal pathway without exception process information from all sensory modalities.

Moreover, Michael Posner (Posner & Petersen, 1990; Posner & Dehaene, 1994) has adduced evidence that, in addition to the function which it shares with the intra-cortical dorsal stream of mediating the visual in the visual control of voluntary movement, the S-C to dorsal pathway also has an important role in the control of selective attention, constituting, as it does, Posner's (Posner & Dehaene, *op.cit.*, p. 76) "posterior attention system (posterior parietal cortex, pulvinar and superior colliculus)". This concatenation of evidence allows us, in the case of the visual modality, to identify the bifurcation between consciousness and the zombie-within as shown on Figure 1 with the bifurcation at the retina between a *ventral* pathway consisting of the *lateral geniculate nucleus*, the *primary visual* or *striate cortex* (V1), the *extra-striate visual areas* (V2-V5) and the *infero-temporal cortex* corresponding to the upstream portion of what I am calling "consciousness", and the *S-C to dorsal pathway* consisting of the *superior colliculus*, *pulvinar* and, more doubtfully in view of its cortical location, the *posterior parietal cortex* corresponding to the zombie. These relationships are shown Figure 3 in similar format to that of the upper part of Figure 1 with the ventral pathway on the left and the S-C to dorsal pathway on the right.

⁴ "Dorsal" here only in the sense that, like the dorsal stream properly so-called, it "terminates" in the posterior parietal cortex.

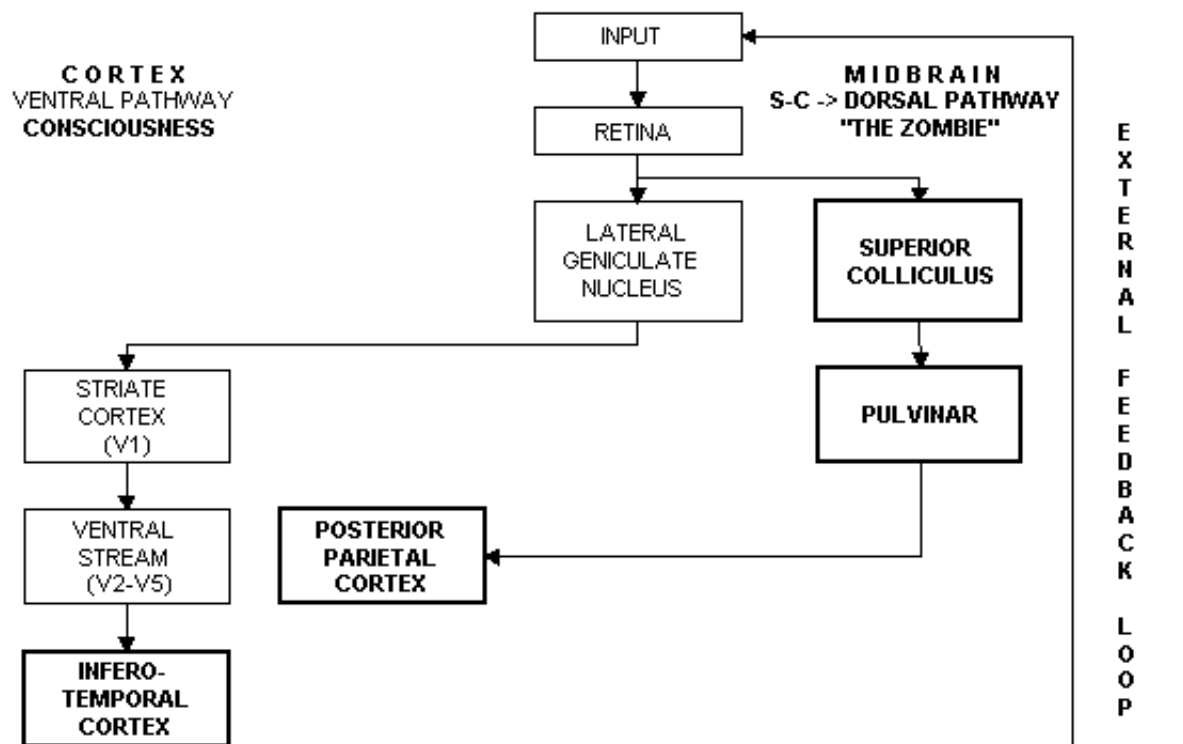


Figure 3: The Dorsal and Ventral Pathways - Multimodal Modules in Bold

Given the identification of the ventral pathway as the route whereby information is passed from the retina into consciousness and the S-C to dorsal pathway as the route into and through the zombie-within, what are we to say about the one exception, the posterior parietal cortex, which is the terminus of both the S-C to dorsal pathway and the intra-cortical dorsal stream? There would seem to be a connection here between the posterior parietal and the visual functions of reaching for objects and avoiding obstacles which are retained when visual conscious experience is abolished by lesions of the striate cortex (V1), thus occluding the ventral visual pathway and yielding the phenomenon of 'blindsight.' Although it is clear from the performance of human subjects with lesions of the striate cortex ("blindsight") that the functions of reaching for objects and avoiding obstacles do not require *visual* conscious experience of the relevant stimuli and are, to that extent, to be regarded on the present hypothesis as functions mediated by the zombie-within, the fact that the subject in such cases is induced to reach for an object he does not "see" by an appropriate instruction to *guess* where it was shows that conscious experience of *some* kind is involved in the production of such behaviour. A plausible hypothesis would be that reaching and obstacle-avoiding behaviour, though it does not require *visual* conscious experience, *does* require the integration of visual information supplied by the zombie along the S-C to dorsal pathway with conscious experience of the *somaesthetic feedback* from the movements involved as they develop, and that this integration is the contribution to these functions made by the posterior parietal cortex. However, some recent evidence (Rossetti, Rode and Boisson 1995; Rossetti 1997) on the somaesthetic counterpart of blindsight which the authors refer to as "numb-sense" shows that a patient (J.A.) with this condition can use his "normal" left hand to point accurately at the location of stimuli applied to the "numb" right hand which he cannot consciously feel. This shows that, provided the blindsight or numb-sense subject can be induced to guess at the location of the target object by pointing at it, successful voluntary movement does not require conscious experience, whether visual or somaesthetic, of the target towards which the movement is directed. What the evidence does not show is that such voluntary movement is possible without conscious experience of the feedback from the movement itself, whether visual, somaesthetic or both.

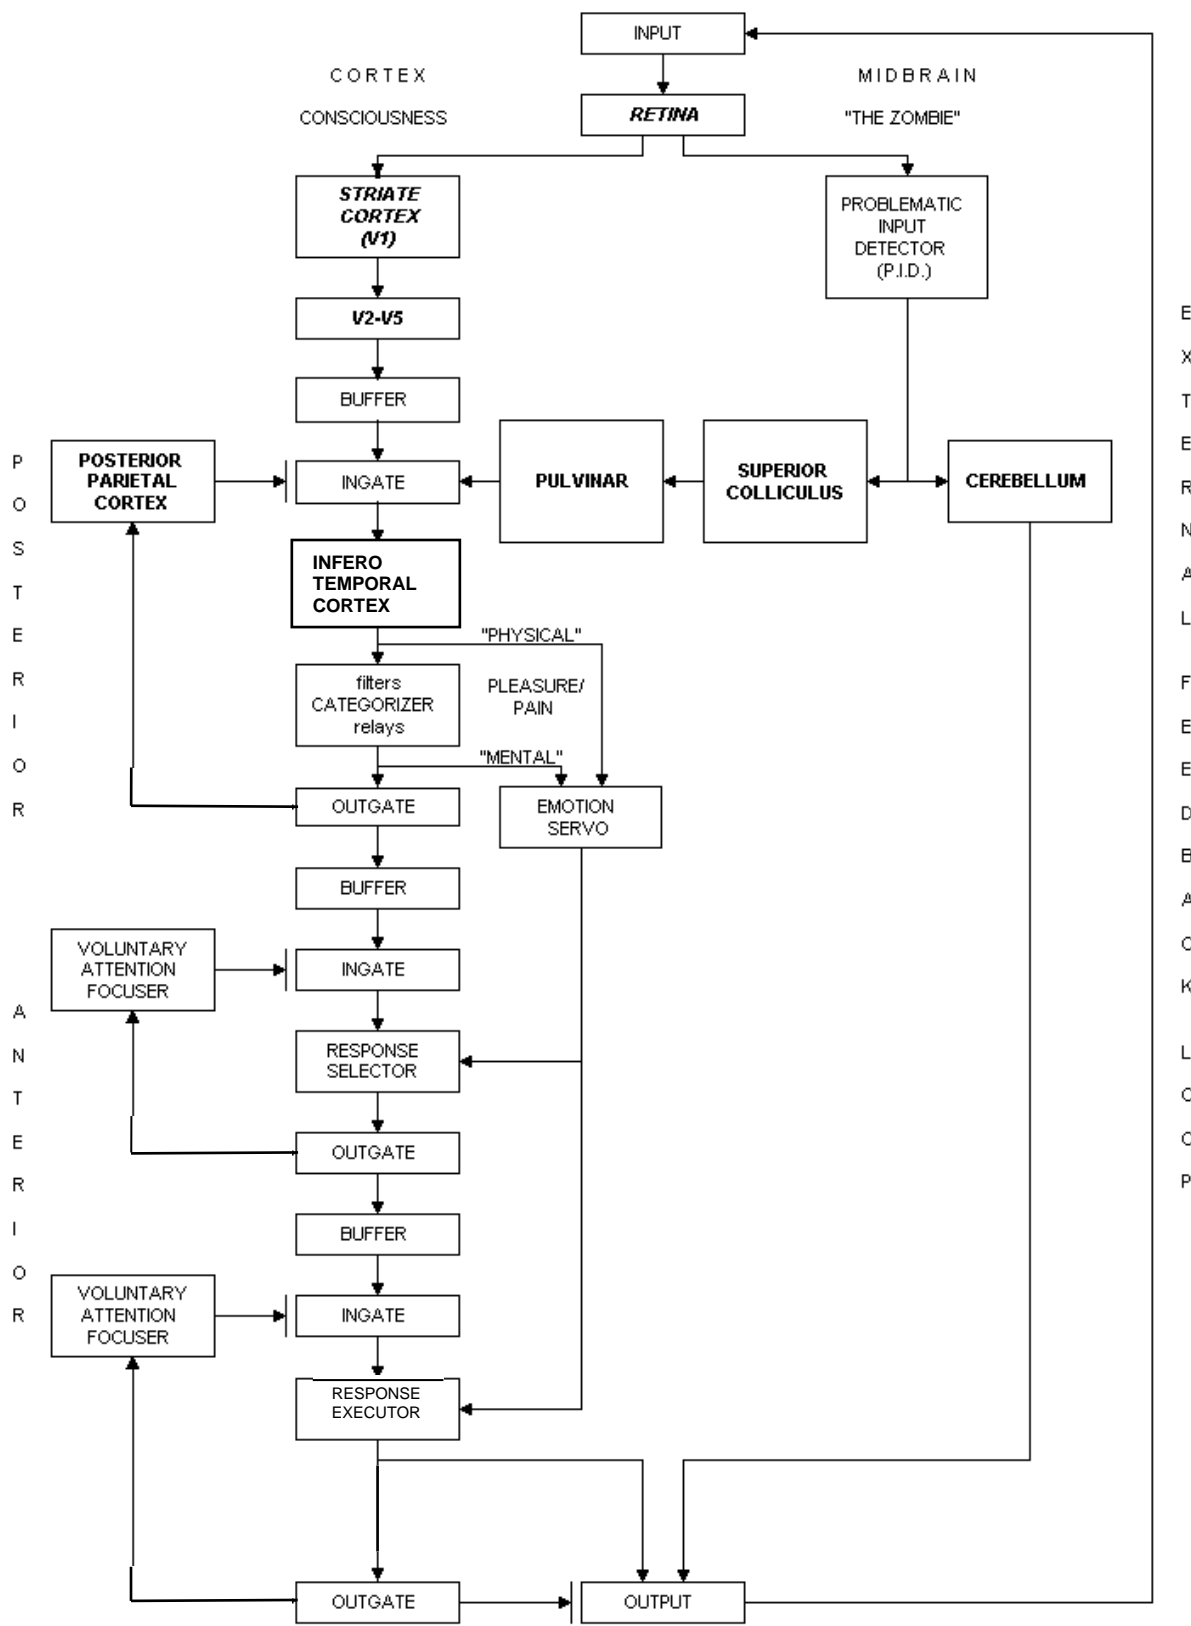


Figure 4: Consciousness and the Zombie-within – Tentative Identification in Bold

Although, as this evidence clearly demonstrates, the S-C to dorsal pathway has a *secondary* function in the visual control of voluntary movement, the fact that all its structures process

information from *all* sensory modalities, when combined with the brain-imaging and neurophysiological data reviewed by Posner and Dehaene (1994) and the evidence of disorders of attention, such as unilateral neglect, resulting from lesions of these structures, suggest that its *primary* function is to integrate the *involuntary* alerting of conscious attention to problematic inputs from all sensory modalities mediated by the two midbrain structures, the superior colliculus and pulvinar, with the *voluntary* maintenance of the focus of attention on such inputs until an adequate categorization of them is achieved mediated by the posterior parietal cortex. On this hypothesis the posterior parietal is construed as having two functions:

- a *general* function which is to *maintain* the focus of conscious attention within and between the different sensory modalities (acting on structures such as those in the ventral stream in the case of vision) on inputs to which the focus has been initially attracted by the "zombie" (in the shape of the superior colliculus and pulvinar) until such time as an adequate categorization of those inputs has been achieved,⁵ and
- a *specific* function which is to control voluntary movement by integrating, through the same mechanism of conscious attention focusing, the visual and somaesthetic feedback from such movements as they develop.⁶

Assuming that this analysis is approximately correct, we are in a position to make some tentative identifications of the modules shown on Figure 1 with some of the actual structures that have been identified anatomically within the brain as laid out on Figures 2 and 3. These tentative identifications are set out on Figure 4 which is a re-working of Figure 1 with the names of the neural structures substituted for the functional descriptions of the modules with which they have been provisionally identified in the preceding discussion in the special case of vision. Thus, in place of the SENSES we have the RETINA. In place of the SENSORY PROJECTION AREA we have the STRIATE CORTEX. In place of the ANALYZER we have V2-V5. In place of the AUTOMATIC PILOT we have the CEREBELLUM. In place of the INVOLUNTARY ATTENTION FOCUSER – PERIPHERAL we have the SUPERIOR COLLICULUS.⁷ In place of the INVOLUNTARY ATTENTION FOCUSER - CENTRAL we have the PULVINAR.⁸ Finally, in place of CONSCIOUS EXPERIENCE "EVIDENCE" we have, at least in the case of vision, the INFERO-TEMPORAL CORTEX.⁹ You will notice that Figure 4 omits the connection between the PULVINAR and the

⁵ Evidence confirming the suggestion that the posterior parietal cortex performs this function is provided by a recent study by Vanni, Revonsuo and Hari (1997) which shows that the magnetic alpha rhythm generated in the parieto-occipital sulcus (POS) is suppressed by object targets to a much greater degree than non-objects, a finding which is readily interpreted as showing the persistence of activity in this area when a stimulus is *not* readily categorized, as compared with the rapid shut down when it is.

⁶ It is worth pointing out that the initiation and control of voluntary movement which is the predominant function of the parietal cortex, considered as a whole, is the same function as that referred to on Figure 1 as RESPONSE EXECUTION. This shows that whereas there is a tendency, particularly when viewed from the standpoint of the visual sense modality, to think of the movement "downstream" in the sense of Figure 1 as a movement from the cerebral cortex to the front, the movement from RESPONSE SELECTION which is predominately a function of the frontal lobes to RESPONSE EXECUTION is a movement in the opposite direction from the frontal to parietal.

⁷ I was first alerted to the role of the superior colliculus in coordinating, at a pre-conscious level, the position and sensitivity of the different sense organs in relation to particular locations in environmental space by a seminar given in Oxford under the auspices of the McDonnell-Pew Centre for Cognitive Neuroscience by Professor Barry Stein of Wake Forest University in 1994 (Stein and Meredith 1995) More recently, my attention was drawn by Antti Revonsuo to evidence summarised by Rafal and Robertson (1995) of the role of the superior colliculus in controlling the reflexive orientating response whose absence on the affected side of the body is characteristic of the phenomenon of unilateral neglect.

⁸ A study by Vanni, Revonsuo and Hari (1997) provides evidence suggesting that the pulvinar is involved in modulating activity in the ventral stream (V2-V5), where object recognition or, as I would think, the preparation of the "evidence" for it occurs and that the effect of such modulation is to "select the next target for ventral processing". If this may be interpreted to mean that the pulvinar controls which parts of the total visual input are currently subject to "ventral processing" and hence in the focus of conscious attention, it supports the suggestion that the function of this structure is to regulate the involuntary attraction of the focus of conscious attention to problematic inputs by processes which are themselves necessarily pre-conscious, i. e., part of what I am calling the "zombie-within".

⁹ Although there is much evidence supporting the idea that activity in the infero-temporal cortex provides the "evidence" in

POSTERIOR PARIETAL CORTEX shown on Figures 2 and 3 and which is needed to explain the visual control of reaching for objects and obstacle avoidance when the relevant parts of V1 have been destroyed ("blindsight"). This has been done in order not to obscure the functionally much more important connection between the PULVINAR and the INGATE controlling access to the INFERO-TEMPORAL CORTEX *alias* CONSCIOUS EXPERIENCE. Other identifications in the case of modules further "downstream" will no doubt suggest themselves to those who have more familiarity with such matters than I have. Likewise the recent work on "numb-sense" mentioned above will doubtless soon make possible the identification of the somaesthetic counterparts of the purely visual structures shown on Figure 4. It may be that similar identifications can already be suggested for other sensory modalities. If not, future research will doubtless allow us to fill these gaps too.

But there, for the present, I shall let the matter rest. I hope I have said enough to persuade you that we are now in a position to answer the question which has remained unanswered since my (Place, 1956) paper 'Is consciousness a brain process?', namely, 'If consciousness *is* a brain process, which of the various processes in the brain that we now identify neuroanatomically is it?' It turns out in the light of what has been said above that that question is too simplistic. But complicated though it is, I hope I have persuaded you that the rudiments of an answer are within our grasp.

3 Towards a Connectionist Theory of Consciousness

In this final part of the paper I discuss the issue of how far the properties of the two systems described in Part 2 can be understood in terms of the known properties of connectionist networks. Since this is a topic on which I can claim even less expertise than I have in the field of neuroscience, I shall confine my remarks to a few disconnected observations.

3.1 Modularity and the Computational Model

In *The Modularity of the Mind* Jerry Fodor (1983) argues for a specific link between the conception of the mind or "mind-brain", as it is now fashionable to call it, as a system of discrete modules or faculties on the one hand and the computational model on the other. He contrasts this "Neo-Cartesian" perspective with the standpoint of Empiricism, by which, according to Chomsky (1980, quoted by Fodor, *op. cit.* p. 3) it is "assumed that development is uniform across (cognitive) domains, and that the intrinsic properties of the initial state (of the mind) are homogeneous and undifferentiated - an assumption found across a spectrum reaching from Skinner to Piaget (who differ on much else)." It goes without saying that in this debate Connectionism is on the side of Empiricism. Does this mean that in presenting a modular analysis of the mind-brain such as that proposed above, I am committed to a computational rather than a connectionist model? I sincerely hope not.

Broadbent's (1971) sense on which the categorization of problematic *visual* inputs is based, what is missing is convincing evidence that information from other sensory modalities is integrated with the visual "evidence" within this structure, as would be required if the unity of conscious experience which is demanded by functional considerations, as much as by the evidence of phenomenology, is to be explained by its concentration within an anatomically defined structure. There is some evidence (Gibson and Maunsell 1997), drawn to my attention by Antti Revonsuo, that there are cells in IT which respond to cross-modal associations between visual and auditory stimuli in a delayed match-to-sample memory task. But similar evidence in the case of the somaesthetic, olfactory and gustatory modalities appears to be missing. Moreover, as we have already seen, there is reason to think that the integration of the visual and somaesthetic feedback "evidence" which forms an integral part of the conscious control of voluntary movement, even when there is no conscious experience of the target towards which that movement is directed, takes place within the posterior parietal cortex. On the other hand, the fact that lesions of IT are associated with selective impairment of "semantic memory" may perhaps be explained on the hypothesis that such lesions disrupt the "evidence" on which the categorization of problematic sensory inputs in all sensory modalities are based. An alternative explanation might be that the boundary between conscious-experience/"evidence" and the categorization that is selected on the basis of it is located *within* rather than downstream of IT. For a further discussion of the nature of this boundary see the final section 3.5 of this paper.

It is true that in characterising many of the modules in my scheme I have used notions such as 'scanning', 'analyzer', 'categorisation', 'ingate' and 'outgate' which smack of the computational model. But there are other functions such as attention-focusing and figure-ground organization which, I would argue are more readily understood in terms of the behaviour of a connectionist network. On the one hand, there is no denying the fact that the neuroanatomically defined structures with which some modules are identified all have a network structure. On the other hand, what distinguishes one structure from another in most, if not all, cases are the distinctive character of the neurons which compose it and the distinctive pattern formed by their synaptic interconnections. These distinctive characteristics must have a functional significance; but, at present, as far as I am aware, we can only speculate as to what that functional significance might be. With the advent of connectionist networks with recurrent circuits (Jordan 1986) we have been made aware of the fact that structural features other than the number of layers may have an effect on the functional properties of a network. John Donahoe (Donahoe, Palmer & Burgos, 1997) has also reported using a network in which nodes fire spontaneously in the way neurons have long been known to do. But, as far as I am aware, there are as yet, no counterparts in connectionist networks for the different types of neuron that are found in the brain or for the different types of neuro-transmitter. All one can say is that future research aimed at making connectionist networks more like those of actual brains is more likely to throw light on the functional significance of such features than is any return to the computational model.

3.2 The Zombie as a Connectionist Network

A striking feature of the model I presented above is the contrast between the complexity of the processes involved in the "processing" by consciousness of problematic inputs and the relatively simple input-output transformations performed by the zombie-within. It is tempting when confronted by this contrast to say that, while consciousness in my sense resembles a digital computer, the zombie behaves like a relatively simple multi-layer network. The case for thinking that we can understand the operation of the zombie in terms of known connectionist principles is particularly strong in the case of what I refer to as the "automatic pilot", the more so, given the identification suggested above of the module with the cerebellum. In their now classic paper expounding the application of their "tensor network theory" to the functioning of the cerebellum Pellionisz and Llinás (1979) give what, but for its neuroscientific sophistication and the fact that it anticipated the advent of connectionism by some six years, could be aptly described as a connectionist theory.

3.3 "Evidence" for Computation in the Discrimination of Problematic Inputs

It is often suggested that the function of alerting consciousness to problematic inputs which must itself be an unconscious process and hence, on my view, a function of the zombie-within can only be explained by postulating a process of unconscious computation. According to Marcel (1983a) evidence for such processing is provided by the observation that written words whose tachistoscopic presentation is "masked" in such a way that the subject cannot consciously read them can, nevertheless, influence the way in which a sentence presented immediately afterwards in the unmasked condition is interpreted. He has also shown (Marcel 1983b) that the same is true of words presented in the blind field of subjects with lesions of the striate cortex ("blindsight"). It is argued that these and similar observations can only be explained by postulating unconscious computational processing in the part of the brain which I refer to as "the zombie within". I contend that this conclusion is unwarranted. With this exception, all the neurological evidence both in relation to the blindsight phenomena and in relation to other visual agnosias makes it unambiguously clear that the sensory discrimination abilities available to the zombie are insufficient for any kind of object-recognition. Moreover, even in these cases, there is no evidence of word recognition either in the sense that the subject can say what word was presented or in the sense that he or she is alerted by its presentation to expect an encounter with the kind of thing to which the word refers. This suggests that Marcel's observation can be explained entirely in terms of well known properties of connectionist networks

- by the ability of networks to discriminate visual patterns of which a written word is an example and
- by the associative links that will have formed between such visual patterns and the auditory patterns constituted by hearing the spoken version of the same word and between *those* and auditory patterns constituting the synonyms and antonyms of the word in question. These latter associations must develop *simply by contiguity*, given Chomsky's principle that sentences are typically constructed anew on each occasion of utterance, implying that a range of alternative ways of "saying the same thing" have to be dispositionally present whenever a sentence is uttered.




Figure 5: Network Required for the Association of Successive Stimulus Patterns (After Jordan 1986)

3.4 Stimulus-Stimulus Expectations and the Discrimination of the Unexpected

To show that there is no need to postulate unconscious computational processing to explain Marcel's observations is one thing, to show that it is possible to give a connectionist-associationist account of the way the zombie detects problematic inputs is another. Problematic inputs are of two kinds:

- those which are *unexpected* and
- those which, though they may be expected, are *motivationally significant* relative to the individual's current and perennial motivational concerns.

To explain how a connectionist network could discriminate motivationally significant inputs would take us too far afield. What we can do more easily is explain how it would discriminate the unexpected.

Given the principle of association by contiguity, it can be predicted that an organism will build up a vast number of *stimulus-stimulus expectations* based on observed regularities in the way an input of one type is invariably succeeded by an input of another type. Given a background of such expectations, a second input which differs from that expected on the basis of past experience, given the first input, is going to stand out like a sore thumb. That is a discrimination that is well within the known capacities of a connectionist network. But so also is the ability to learn stimulus-stimulus expectations.

Figure 5 is adapted from a diagram in a paper by Jordan (1986) and reproduced in McClelland and Rumelhart (1988, Fig. 11, p. 157). It shows the arrangement of a network which can learn just such stimulus-stimulus sequences with *reverberating circuits* to keep the S1 activity going until S2 appears or, as when the unexpected happens, fails to appear, and *recurrent circuits* which feed back, in this case, into the input nodes so as to generate an output only if both S1 and S2 activity are simultaneously impinging on the relevant input node. Figure 5 differs from that given by Jordan in that the recurrent circuits are shown as impinging on the so-called "hidden layer" between the input and output layers instead of on the input layer as in Jordan's diagram. This is to allow for the essentially dispositional character of such expectations. In the language of the philosopher, it allows for the fact that we are dealing here with *stimulus types* rather than stimulus tokens. Moreover, these are types within which there is room for considerable variation inside the limits imposed by the range of variation in the sequences previously experienced by the system. As we shall see in a moment, this generalization is achieved in a connectionist network by adding a further "hidden" layer downstream of that at which inputs arrive.




Figure 6: Two Layer Network with Stimulus Generalization Only




Figure 7: Three Layer Network with Abstract Feature Detectors in the Output Layer

3.5 Consciousness as a Connectionist Network: Feature Detectors and Conceptual Hierarchies

It is not only the zombie-within whose properties are explicable in terms of the known properties of a connectionist network. There are at least *some* properties of consciousness in the extended sense in which I am using that term which lend themselves to explanation in this way. That this is so is hardly surprising if, as I have argued, consciousness in this sense is to be identified with the activity of the cerebral cortex. For the micro-structure of the cortical "grey matter" is just as much that of a neural network as is that of structures such as the superior colliculus and cerebellum which we have assigned to the zombie. Furthermore, if the automatic pilot operates as a connectionist network, and if, as the evidence adduced by Raichle *et al.* (1994) suggests, input-output transformations worked out in consciousness are transferred in a streamlined form to the automatic pilot in the midbrain once they have become habitual, it is difficult to see how that transfer could take place unless consciousness itself is also a connectionist network or complex of such networks.

As I see the matter, the two features of consciousness that lend themselves most readily to interpretation in connectionist terms are

1. the hierarchy of *feature detectors* in the module identified as the ANALYZER on Figure 1 and as visual areas V2-V5 on Figure 4 which respond to more and more abstract features of the input as neural activity moves downstream, and
2. a similar hierarchy moving downstream in the opposite sense i.e. from the more abstract and general to the more specific, which must be postulated both within the CATEGORIZATION module, as illustrated by the principle underlying Aristotle's doctrine of the Categories and its application in the Tree of Porphyry, and within the RESPONSE SELECTION module as a vague intention is narrowed down to a specific response strategy.

A connectionist interpretation of the principle underlying the hierarchy of feature detectors in the Analyzer is illustrated on Figures 6 and 7. Figure 6 show a two layer network in which the same pattern wherever it occurs in the input layer triggers the same node, i.e., the same feature detector, in what in this case is the output layer. Figure 7 shows a three-layer network in which the function of the

intermediate or "hidden" layer is to detect patterns in the input layer and in turn trigger nodes in the output layer which respond to more abstract features of the patterns detected by the "hidden" layer (in this case whether or not the pattern contains a lower case letter alongside one or more upper case letters and whether it consists of two or three letters). It should be obvious that by piling layer on layer in this way the network can detect patterns in patterns in patterns, each layer responding to a more and more general or abstract feature of the input, exactly as is observed in the feature detector neurons as you progress from visual area V1 to visual area V5.

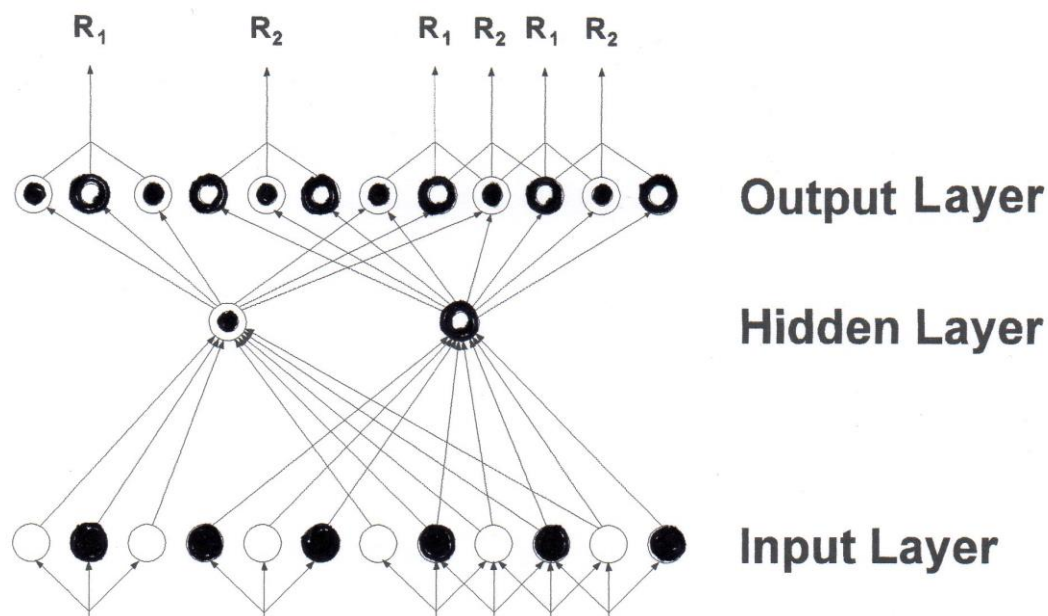


Figure 8: Three Layer Network with Stimulus and Response Generalization

Figure 8 shows the rudiments of the opposite process of increasing specificity as you move downstream on the output side. Here we have a three layer network in which firing in the feature detector nodes in the intermediate or "hidden" layer triggers the same pattern on the output side, but in different groups of effectors depending on the prevailing conditions in other respects. Likewise Figure 9 shows a three-layer network on the output side of a multi-layered network in which the eight nodes in the intermediate or "hidden" layer combine the one or other of each of the three abstract features specified by the six nodes in the "input" layer and are in turn expressed in terms of two pairs of values for the variables 'x' and 'y' in the output layer. Again one can see very easily that the effect of adding further layers on the output side is to move from the more general to the more specific, just as on the input side we move from the more specific to the more general.

Just one parting shot. On the scheme, as set out on Figure 1 the transition within consciousness from increasing generality on the input side to increasing specificity on the output side occurs at the interface between the CONSCIOUS EXPERIENCE/"EVIDENCE" and CATEGORIZATION modules, just where phenomenology, I suggest, would locate the ego.




Figure 9: Three Layer Network Moving from Abstract Features in the Input Layer to Specific Features in the Output Layer

References

- Broadbent, D. E. (1958) *Perception and Communication*. Oxford: Pergamon.
- Broadbent, D. E. (1971) *Decision and Stress*. London: Academic Press.
- Chomsky, N. (1980) Rules and representations. *The Behavioral and Brain Sciences* 3, 1-15.
- Churchland, P. M. (1981) Eliminative materialism and propositional attitudes. *Journal of Philosophy* 78, 67-90.
- Cowey, A. and Stoerig, P. (1995) Blindsight in monkeys. *Nature* 373 (6511), 247-9.
- Cowey, A. and Stoerig, P. (1997) Visual detection in monkeys with blindsight. *Neuropsychologia*, 35, 929-939.
- Donahoe, J. W., Palmer, D. C. and Burgos, J. E. (1997) The S-R issue: clarification of its status in Donahoe and Palmer's *Learning and Complex Behavior*. *Journal of the Experimental Analysis of Behavior*, 67, 193-212.
- Farah, M. J. (1990) *Visual Agnosia: Disorders of Object Recognition and what they tell us about Normal Vision*. Cambridge, MA: M.I.T. Press.
- Fodor, J. A. (1975) *The Language of Thought*. New York: Crowell.
- Fodor, J. A. (1983) *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Geach, P. T. (1957) *Mental Acts* London: Routledge and Kegan Paul.
- Gibson, J. R. and Maunsell, H. R. (1997) Sensory modality specificity of neural activity related to memory in visual cortex. *Journal of Neurophysiology*, 78 (3), 1263-1275.
- Greenfield, P. M. & Savage-Rumbaugh, E. S. (1990). Grammatical combinations in Pan paniscus: Processes of learning and invention in the evolution and development of language. In S. T. Parker & K. R. Gibson (Eds.), *"Language" and Intelligence in Monkeys and Apes: Comparative Developmental Perspectives* (pp. 540-578). New York: Cambridge University Press.
- Herman, L. M. (1987). Receptive competencies of language-trained animals. In J. S. Rosenblatt, C. Beer, M.-C. Busnel, & P. J. B. Slater (Eds.), *Advances in the Study of Behavior* (Vol. 17,

- pp.1-60). Orlando: Academic Press.
- Herman, L. M. Kuczaj, S. A. & Holder, M. D. (1993). Responses to anomalous gestural sequences by a language-trained dolphin: Evidence for processing of semantic relations and syntactic information. *Journal of Experimental Psychology: General*, 122(2), 184-194.
- Horgan, T. and Tienson, J. (1996) *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.
- Humphrey, N. K. (1974) Vision in a monkey without striate cortex: a case study. *Perception* 3, 241-255.
- Jordan, M. I. (1986) Attractor dynamics and parallelism in a connectionist sequential machine. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Luck, S. J. and Beach, N. J. (1996) Visual attention and the binding problem: a neurophysiological perspective. In R. Wright (ed.) *Visual Attention*. Oxford: Oxford University Press.
- Marcel, A. J. (1983a) Conscious and unconscious perception: experiments on visual masking and word recognition. *Cognitive Psychology*, 15, 197-237.
- Marcel, A. J. (1983b) Conscious and unconscious perception: an approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, 15, 238-300.
- McClelland, J. L. and Rumelhart, D. E. (1988) *Explorations in Parallel Distributed Processing*. Cambridge, MA: MIT Press.
- Milner, A. D. and Goodale, M. A. (1995) *The Visual Brain in Action*. Oxford: Oxford University Press.
- Moruzzi, G. and Magoun, H. W. (1949) Brain stem reticular formation and activation of the EEG. *Electroencephalography and Clinical Neurophysiology*. 1, 455-473.
- Pashler, H. E. (1991) Shifting visual attention and selecting motor responses: distinct attentional mechanisms. *Journal of Experimental Psychology: Human Perception and Performance* 17, 1023-1040.
- Pashler, H. E. (1997) *The Psychology of Attention*. Cambridge, MA: MIT Press.
- Pellionisz, A. and Llinás, R. (1979) Brain modeling by tensor network theory and computer simulation. The cerebellum: distributed processor for predictive coordination. *Neuroscience*, 4, 323-348.
- Pepperberg, I. M. (1990) Referential mapping: A technique for attaching functional significance to the innovative utterances of an African Grey parrot (*Psittachus erithacus*). *Applied Psycholinguistics*, 11, 23-44.
- Pinker, S. and Prince, A. (1988) On language and connectionism: analysis of a parallel distributed model of language acquisition. *Cognition*, 28, 73-193.
- Place, U. T. (1954) The concept of heed. *British Journal of Psychology* 45, 243-255.
- Place, U. T. (1956) Is consciousness a brain process? *British Journal of Psychology* 47, 44-50.
- Place, U. T. (1997) Consciousness and the Zombie-within. Paper presented at the Inaugural Conference of the Association for the Scientific Study of Consciousness, Claremont, CA, June 15th 1997.
- Posner, M. I. and Dehaene, S. (1994) Attentional networks. *Trends in Neuroscience* 17, 75-79.
- Posner, M. I. and Petersen, S.E. (1990) The attention system of the human brain. *Annual Review of Neuroscience* 13, 25-42.
- Rafal, R. and Robertson, L. (1995) The neurology of visual attention. In M.S. Gazzaniga (ed.) *The Cognitive Neurosciences*. Cambridge, MA: MIT Press, ch. 40, pp. 625-648.
- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A.-M. K., Pardo, J. V., Fox, P. T. & Petersen, S. E. (1994) Practice-related changes in human functional anatomy during non-motor learning. *Cerebral Cortex* 4, 8-26.
- Rizzolatti, G. & Arbib, M. A. (1998) Language within our grasp. *Trends in Neuroscience*, 21, 188-194.
- Rossetti, Y. (1997) Implicit perception in action: short-lived motor representations in space. P.G. Grossenbacher (ed.) *Advances in Consciousness Research*. Amsterdam: John Benjamins.

- Rossetti, Y., Rode, G. and Boisson, D. (1995) Implicit processing of somesthetic information: a dissociation between Where and How? *Neuroreport* 6 (3), 506-510.
- Schusterman, R. J. & Gisiner, R. C. (1988) Artificial language comprehension in dolphins and sea lions: The essential cognitive skills. *The Psychological Record*, 38, 311-348.
- Schusterman, R. J. & Krieger, K. (1984) California sea lions are capable of semantic comprehension. *The Psychological Record*. 34, 3-23.
- Searle, J. R. (1980) Minds, brains and programs. *The Behavioral and Brain Sciences* 3, 417-424.
- Snowdon, P. (1995) Perception and attention. Paper presented to a one-day conference on 'Attention and Consciousness: Psychological and Philosophical Issues', Department of Philosophy, University College London, 26th May 1995.
- Stein, B. E. and Meredith, M. A. (1993) *The Merging of the Senses*. Cambridge, MA: MIT Press.
- Stoerig, P. and Cowey, A. (1997) Blindsight in man and monkey. *Brain*, 120, 535-559.
- Treisman, A. (1988) Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, 40, 201-237.
- Treisman, A. and Gelade, G. (1980) A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Treisman, A. and Gormican, S. (1988) Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95, 15-48.
- Trentman, J. (1970) Ockham on mental. *Mind*, 79, 586-590.
- Ungerleider, L. G. and Mishkin, M. (1982) Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.) *Analysis of Visual Behavior*. Cambridge, MA: M.I.T. Press.
- Vanni, S., Revonsuo, A. and Hari, R. (1997) Modulation of the parieto-occipital alpha-rhythm during object-detection. *Journal of Neuroscience*, 17 (18), 7141-7147.
- Weiskrantz, L. (1986) *Blindsight*. Oxford: Clarendon Press.
- Wilkes, K. V. (1984) Is consciousness important? *British Journal of the Philosophy of Science*. 35, 223-243.