Intellectica, 1992 1/2, 13-14, pp. 241-271

Simon J. THORPE

Psychology and Neuroscience : Towards a Common Language

1. Introduction

In recent years, scientists have been increasingly called upon to develop interdisciplinary research programs. For some, interdisciplinarity is essential if we are to make progress in the future. For others, it's a waste of time. Who is right? In this article, I would like to take the case of research in Psychology and Neuroscience, and ask to what extent interdisciplinary research is (a) feasible, and (b) desirable. I will conclude that, although for large parts of Psychology and Neuroscience interdisciplinary collaborations have little to offer, for many scientists working on particular topics, collaboration with other disciplines can be highly profitable.

There will no doubt be some readers who are already convinced that Neuroscience and Psychology can indeed cohabit - for them I will be preaching to the converted. The last decade has seen the development of a research domain entitled “Cognitive Neuroscience” whose very existence shows that the dividing lines between Psychology and Neuroscience are becoming increasingly blurred (Churchland and Sejnowski, 1988, LeDoux and Hirst, 1986, Requin, 1987, Scheibel and Wechsler, 1990, Thompson, 1990). In 1989 Michael Gazzaniga founded an interdisciplinary journal (“Journal of Cognitive Neuroscience”), and in

France, a number of laboratories have set up interdisciplinary research programs.

There are nevertheless a large number of researchers who see Psychology and Neurosciences as too different in approach to allow successful integration. It is to this audience that this paper is addressed.

Advantages of an interdisciplinary approach

One of the clearest advantages of using an interdisciplinary approach is that it often allows a problem to be seen from a completely different perspective. A good analogy is stereoscopic vision. If you look an a scene with one eye at a time, it always has a tendency to look flat. However, if you can use both eyes at the same time, stereoscopic mechanisms allow you to make out the true three-dimensional shape of the object. Indeed, an object so well camouflaged as to be invisible when seen with one eye, may become obvious when seen with two eyes. And so it is in science. If you can get a variety of views of the same problem, you are in a much better position to assess its overall form and to develop the best strategies for tackling it. Note that this does not mean that any particular view is better - that would be like trying to decide whether the left or right eye's views were better. What counts is the advantage gained by combining both views.

One way in which an interdisciplinary approach may help is by providing information that can constrain models. For example, it is perfectly possible for a psychologist to produce a model of some particular psychological processes using purely behavioural data. Often, several competing models may be available, and so the question arises of how the different models can be distinguished. It is at moments such as these when information coming from a different perspective can be particularly useful. We will see an example of this later when we discuss the speed of processing in the human visual system.

There is of course a danger when one confronts two disciplines as different as Psychology and Neuroscience that stems from the tendency towards reductionism. Some may well feel that if psychological phenomena can be "reduced" to neurophysiology, then the psychological phenomena themselves are no longer of interest and need not be studied further. I believe strongly that this is quite erroneous. Certainly, science should strive towards reductionist models when possible, but in the course of such an enterprise, it should never be forgotten that the aim is to understand the higher level processes. Blind analysis of lower level mechanisms without consideration of their overall function is unlikely to go far.

One way of reconciling the different points of view is to see interdisciplinary research in Psychology and Neuroscience not within the framework of reduction, but rather with a framework of synthesis. Valentino Braitenberg wrote a delightful book called "Vehicles: experiments in synthetic psychology"(Braitenberg, 1984), in which he imagined a series of simple artificial life-forms, equipped with simple sensory and motor capacities, and rudimentary nervous systems. He was able to show that even these simple "organisms" could demonstrate some remarkably sophisticated behaviours. Note that here the approach has been turned on its head - no longer are we trying to eliminate higher level properties by reducing them to phenomena at a lower level. Rather, the aim is to reproduce complex phenomena by putting together relatively simple elements. I believe that this "synthetic" approach re-establishes the correct balance between the different levels.

Interdisciplinary research: for whom?

In the preceding paragraphs, I have argued that interdisciplinary research offers (i) a richer perspective view, (ii) greater potential for constraining models and (iii) the promise of a "synthetic" rather than a purely eliminative reductionism. Does this mean that I believe all psychologists should suddenly establish active collaborations with the nearest neuroscientist and vice versa? The answer is very definitely no, and I will attempt to say why.

Both Psychology and Neuroscience are vast domains. Psychology in particular covers an immense area, stretching from industrial, occupational and social psychology to psychophysical studies of perceptual processes. Likewise, a brief look at the abstracts of any major Neuroscience meeting reveals the immense range of topics covered. You will find detailed studies of the anatomy and physiological of different brain structures, but also a vast amount of research on the molecular biology of neurotransmitters and their receptors. What percentage of work in these two great domains would actually benefit substantially from increased interdisciplinarity? I think the answer is perhaps quite low - perhaps as little as 10-20%, but although the percentage may be small, the need for interdisciplinarity is urgent.

What characterizes the 10-20% of Psychology and Neuroscience for which interdisciplinarity holds the greatest promise? I firmly believe that the critical point is whether or not there is a common object. To take an example which is close to my own heart, namely vision, it is clear that there are questions which are just as pertinent to someone working in Psychology as to someone working in Neuroscience. Questions such as "how long does it take to process an image?", "how many levels of processing are involved?", and "how does experience influence visual processing?" are all questions that are meaningful to people with backgrounds in Psychology or Neuroscience. It may well be that there are currently no satisfactory answers to such questions, but the mere fact that the questions have meaning to both communities encourages me in the belief that an interdisciplinary approach can be applied.

Vision is not the only area where the case for interdisciplinarity is clear. All areas of sensory processing merit such an approach, including auditory processing of sounds and speech, as well as tactile, gustatory and olfactory processing. One aspect of sensory processing that merits particular interest is the study of attention (Posner, et al., 1987, Wurtz, et al., 1982). Another key area of overlap concerns the study of memory mechanisms - there has long been the possibility that short and long-term memory mechanisms may be distinguished both psychologically and neurophysiologically. This is one reason why the "psychobiology" of memory is such a relatively advanced example of interdisciplinary research. Finally, the study of movement control is yet another area where an interdisciplinary approach is not only to be encouraged, but where it is almost a necessity.

Clearly, the study of perception, attention, memory and motor control does not exhaust the subject matter of either Psychology or Neuroscience, and so it is definitely not the case that everyone working in these areas should be pushed towards collaboration.

Interdisciplinarity: The obstacles

Nevertheless, even within those areas which are clearly suitable for interdisciplinary work, collaboration has been the exception rather than the rule, although this has clearly been changing in recent years. What are the reasons for this?

The first reason is related to the point I made in the last paragraph. The majority of researchers working in Psychology and Neuroscience really do have little overlap. As a result, if a psychologist was to try discussing with the first neuroscientist to hand, the chances are that nothing very interesting will happen. And of course, the same is true for the neuroscientist tempted by the possibility of interacting with someone trained in psychology. This, coupled with a natural tendency to lump every one in another discipline together ("I've tried talking to people in neuroscience - it just doesn't work"), is enough to explain why the initial results can be most discouraging. However, if the researchers in psychology and neuroscience have common interests (such as visual processing, for example), such early encounters can be very rewarding.

A second reason stems from the insecurity that many researchers feel as soon as they are no longer on their home territory. For many, having to admit that they are just plain ignorant about a subject is too great a shock for the ego. I think that one should fight against such feelings - there is nothing wrong with admitting that there is a lot to learn. Furthermore, I know from experience that it is often the naive question from someone outside your own field which reveals the fundamental unresolved issue that you yourself have tended to push to one side, even subconsciously.

A third factor which often limits communication results from the fact that researchers in adjacent disciplines are often concerned with different types of measurements. Thus, in experimental psychology, results are typically expressed in terms of (a) reaction times, (b) performance (i.e., percentage correct), or (c) verbal responses. In neurophysiological studies, one typically determines the selectivity of neuronal responses. Surprisingly, the measurement of neuronal response times has, at least in the visual system, been remarkably rare (Robinson and Rugg, 1988), and this has severely limited one area where interdisciplinary collaboration should have been relatively straightforward. We will see later on how some progress has been made on this question.

Finally, I believe that perhaps the most significant obstacle to interdisciplinary research has been the lack of a common language. Psychologists and Neuroscientists have, at least in the past, simply not been on the same wavelength. Models in psychology and models in neuroscience have generally been completely different and so using data from one domain to validate models in another has been virtually impossible.

However, I would like to argue that all this is changing. The advent of connectionist models in psychology, essentially during the last ten years, means that we now have the possibility of generating a class of explanatory models which actually mean something to workers in both psychology and neuroscience (Bechtel and Abrahamsen, 1991, Memmi, 1990).

2. Connectionist models

In there simplest form, connectionist models can be defined as systems composed of a large number of interconnected processing units (Arbib, 1987). Each unit can be relatively simple, but processing power is obtained by allowing the network of units to operate in parallel. The state of the system is effectively defined by the state of activation of the units, and information is stored by varying the strengths of the connections between units. Although the history of connectionist modelling in psychology goes back a long way (Anderson and Rosenfeld, 1988), it has essentially been during the last ten years that they have started to really change the face of psychological theory. A major turning point was the publication of McClelland and Rumelharts' Interactive Activation model of word recognition in 1981 (McClelland and Rumelhart, 1981), but even more influential were the pair of volumes entitled "Parallel Distributed Processing: Explorations of the microstructure of cognition" published in 1986 (McClelland and Rumelhart, 1986, Rumelhart and McClelland, 1986). Since that time, connectionist models have been applied to many areas of psychology, including such diverse areas as the recognition of multiple objects (Mozer, 1991), face recognition (Burton, et al., 1990); attention (Phaf, et al., 1990, Sandon, 1990), language development (Plunkett and Marchman, 1991, Seidenberg and McClelland, 1989), the effects of brain damage (Hinton and Shallice, 1991), music perception (Bharucha, 1988), and conditioning (Kehoe, 1989) to name but a few.

There is not space here to describe in detail the pros and cons of connectionist models. The essential point to retain is that since the models themselves are at least to some extent "neurally-inspired", they at least have the potential of being interpretable in both the psychological and the neuroscience domains.

Despite this potential, many researchers often play down the biological realism of the models. Many connectionists will affirm that they make no claim whatsoever that the nodes in their models bear any resemblance to real neurones - they will often say that whether or not such an equivalence is possible is quite irrelevant to the explanatory potential of a connectionist model. While such a position is perfectly understandable and in many ways quite justified, it does limit the potential for interdisciplinary collaboration.

One reason why many connectionists see no direct relevance of neurophysiological findings for their models lies in the nature of the problems that they are addressing. Often, they are interested in high level cognitive functions such as language, and although such questions have been investigated with techniques in neuroscience such as evoked potentials (Kutas, et al., 1988, Paller, et al., 1987, Rugg, 1990) and positron emission tomography (Petersen, et al., 1988, Posner, et al., 1988), such studies have tended to provide information about anatomical and functional specializations in language processing but not about the response properties of individual neurones in the network. In fact, there have been some fascinating recent studies of neuronal activity in the human temporal lobe which have begun to shed some neurophysiological light on such problems (Creutzfeldt and Ojemann, 1989, Creutzfeldt, et al., 1989a, Creutzfeldt, et al., 1989b, Heit, et al., 1988)), but in general, detailed neurophysiological data is generally only available for more “peripheral” sensory and motor processing.

To increase the potential cross-fertilization between psychology and neuroscience that connectionist modelling can offer, it is necessary to choose an target area for which both psychological and neurophysiological data are available. The last few years have seen a small number of such attempts ((Lehky and Sejnowski, 1988, Lehky and Sejnowski, 1989, Lehky and Sejnowski, 1990, Zipser and Andersen, 1988), and I am confident that such attempts at integration will become increasingly common in the years to come.

In a previous section I mentioned the fact that in experimental psychology, experimental data usually refers to either the subjects performance (for instance as a percentage of correct responses) or to the subjects speed (measured as a reaction time). In some situations the experimenter may be interested in subjective reports, but in the majority of cases, it is either performance accuracy or reaction time which is of interest. In the next sections I will try to illustrate how connectionist models can link psychological with neurophysiological data domains for these two types of paradigm. In the first, I will outline one example based on work from my own group which illustrates how connectionist models can be used to bridge the conceptual gap between neurophysiological and psychological data in the area of orientation identification.

3. Connectionism and performance measures: The example of orientation identification.

It has been known since the early 1960s that the visual cortex contains neurons which selectively respond to contours with particular orientations (DeValois and DeValois, 1988, Henry, et al., 1974, Howard, 1982, Hubel and Wiesel, 1962, Orban, 1984). However, the way in which the activity of such neurons is used to code the orientation of a stimulus remains controversial. One popular view, based on Selfridge's "Pandemonium" model (Selfridge, 1959) is to think of these orientation-selective neurons as "feature-detectors" and to suppose that perceived orientation is determined by the ones which are most active. A clear example of this "peak-activity" coding position is given in a recent review article on computational maps in the brain in which Eric Knudsen and his co-workers proposed the following formulation: "Whether the desired information is the orientation or direction of motion of a visual stimulus ...., the answer is always represented as the location of a peak of activity within a population of neurons" ((Knudsen, et al., 1987), page 59). However, recent psychophysical and neurophysiological data cast doubt on such a formulation, at least in the case of orientation perception. It would appear that it is perhaps not the most active neurons which determine perceived orientation, but rather the relative levels of firing in neurons tuned to orientations offset to either side of the orientation to be judged.

In order to gain further insight into the relationship between the activity of orientation selective neurones and perceived orientation, Alexandre Pouget and I decided to use a connectionist model of orientation identification (Pouget and Thorpe, 1991, Thorpe and Pouget, 1989). The network itself was fairly straightforward (see figure 1). There were just two layers of units, an input layer and an output layer. The input layer was made up of a large number of units with properties similar to neurones in the visual cortex. They were orientation selective, in that they responded to stimuli having a certain range of orientations. They had a bell-shaped Gaussian tuning curve, with a bandwidth at half-height similar to real neurons (roughly 25°). Each unit preferred a particular orientation, so that the complete set of units covered the entire range of orientations. Importantly, the units in the input layer were noisy in that they did not always respond in the same way to a particular stimulus - instead they showed a mean level of firing to which was added a certain amount of gaussian noise. The output units were trained to classify the orientation of a presented stimulus, in much the same way as human subjects were trained in a parallel psychophysical study. During training, feedback was provided to allow the subject (or the network) to identify which of two possible orientations had been presented on any particular trial. When performance had reached asymptotic levels, feedback was eliminated and performance determined over a large number of trials (400 in the case of the human subjects). In separate experimental sessions, performance was assessed with stimuli varying in separation from as little as 0.35° to over 10°, and in this way, a psychometric function for orientation identification could be established.

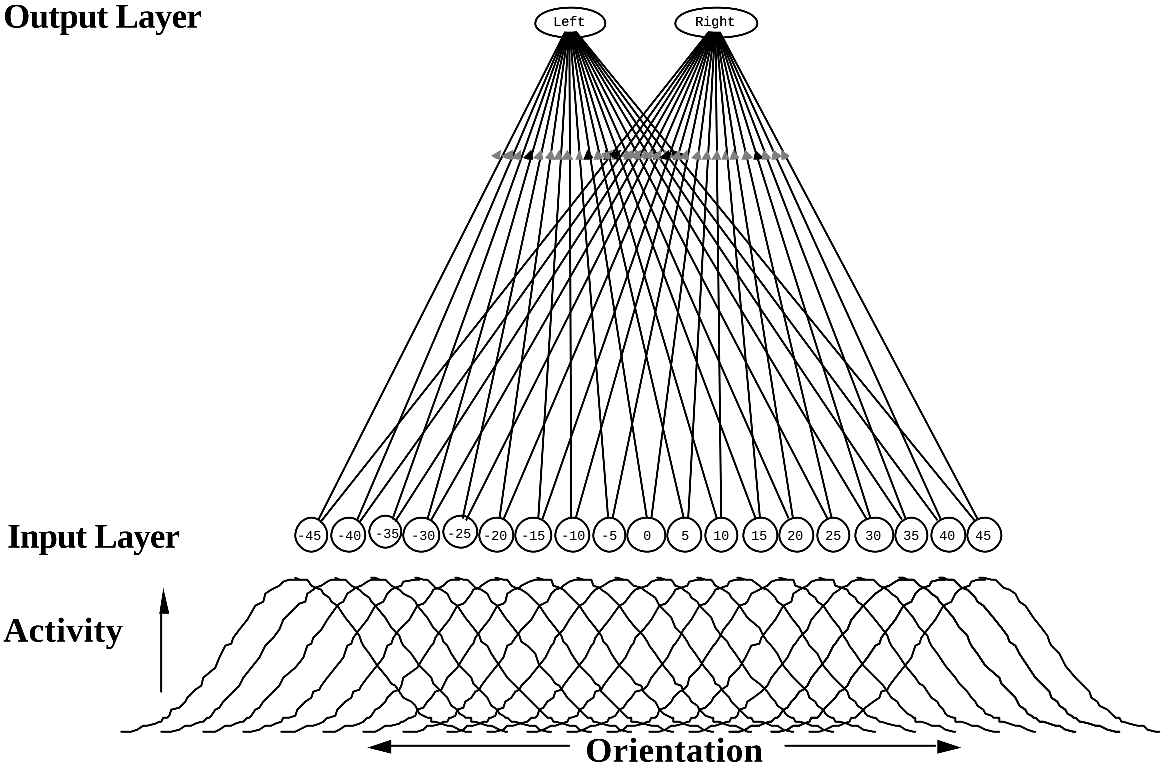

Figure 1. The connectionist model used by Pouget and Thorpe (1991) to model orientation identification. The input layer is composed of a series of orientation tuned units, each with a Gaussian filter shape. The output layer contains units that were trained to respond to different test stimuli.

The first result was to find that the network had performance curves that could accurately fit the human psychophysical data (see figure 2). In a task with two stimuli to identify, performance starts at close to 50% with the smallest separations, but increases rapidly to around 75% for a separation of 1.4° before reaching virtually perfect levels of performance at around 10°. It was therefore clear that we had a reasonably good model of the human psychophysical performance curves. But the second major finding was considerably more interesting. When we looked at the way in which the synaptic weights had been set up as a result of training we found that they highest weights were not associated with units tuned to the orientations used during training (see figure 3). Rather, it was units preferring orientations well to either side that were the most important. In effect this means that if you wish to decide whether a line is 1° to the left or to the right of vertical, your best strategy would not be to compare the activity of units tuned to +1° and -1°. A much better strategy is to compare the relative activity of units tuned to orientations 10-15° to either side of the vertical. In retrospect it is not too difficult to see why this is the case. Because of the bell-shaped tuning functions of units in the input layer, and their relatively large band-widths, there is hardly any change in the activity of a unit tuned to -1° when the stimulus changes in orientation from -1 to +1°. In contrast, the same small change in orientation produces a much more substantial change in output for units tuned to 10-15° to either side of the vertical.

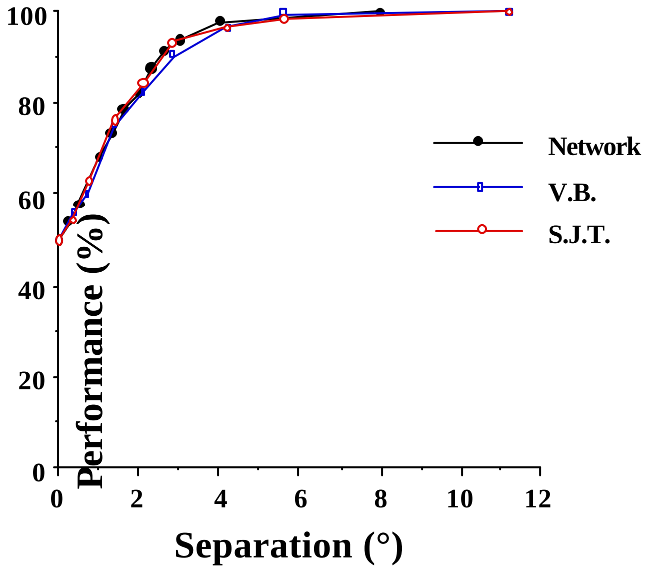

Figure 2: Comparison of the performance of the connectionist net with the performance of two human observers (SJT and VB) on an orientation identification task with two possibilities. In both cases performance increased from 50% correct (chance) with the smallest separation between the test stimuli, to 100% when the separation was more than about 6° (See Pouget and Thorpe, 1991).

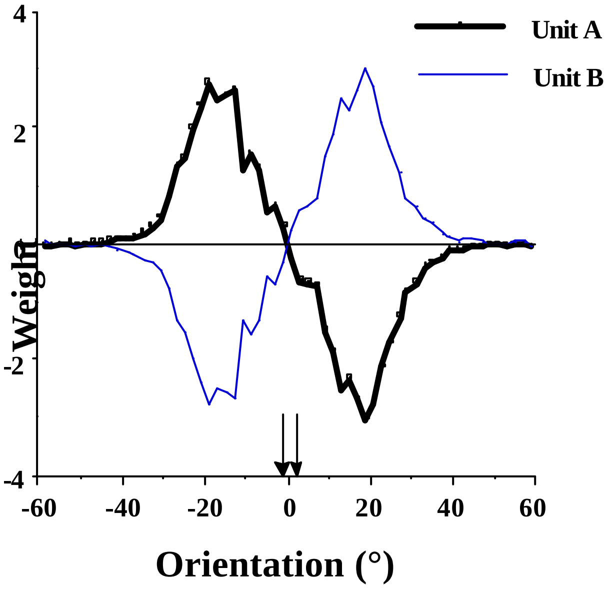

Figure 3: Connection strengths between the input layer units and the two output units A and B after training with two test orientations separated by 2°. Note that the highest weight strengths are not found for the most strongly activated units (i.e. those centered on the test stimuli), but rather the ones whose peaks are offset to the right and left by 15-20° (Pouget and Thorpe, 1991).

In presenting this example of a connectionist model for orientation identification, I hope to have shown how connectionist models can provide a way of linking neurophysiological and psychological levels of analysis. It is important to stress that until recently, few people could claim to understand how it was that humans could make orientation discriminations of less that 1°, despite the poor tuning of the neurones in the visual system. Connectionist models have provided a plausible way of bridging these two levels, but more importantly, they have provided insights into the mechanisms of coding, casting doubt on the widely accepted notion of "feature-detectors", and arguing for a more distributed coding scheme, at least in the case of orientation.

4. Connectionism and reaction time date

As I mentioned earlier, a great deal of work in experimental psychology deals not with performance measures such as the percentage of correct responses, but rather with reaction time measures. It may been shown, for example, that the time taken to decide whether a string of letters is a word or not depends on a variety of factors, such as the words frequency, what words had been presented just before, and whether the word in question had been seen earlier or not. How can connectionist models provide insights into such effects?

One of the most popular connectionist models ever developed is the interactive activation model, first presented by McClelland and Rumelhart in 1981 (McClelland and Rumelhart, 1981, Rumelhart and McClelland, 1982). The basic idea is to have a system of interconnected units in which units send outputs to other units which in turn send signals back. The fact that information is continuously flowing back and forth between units means that the whole system will have quite complex dynamics. Typically, in response to a particular input, the system with gradually settle into a stable configuration over a number of cycles. This type of constraint-satisfaction process is actually typical of many connectionist and neural network systems, including the associative memory nets described by Hopfield (Hopfield, 1982, Hopfield, 1984), Grossberg's Adaptive Resonance Theory or ART (Carpenter and Grossberg, 1990, Grossberg, 1987), along with many others.

In recent years, such interactive models have been used to model reaction time data in a number of areas. Typically, an interactive activation type of model is used, and the number of iterations required for the network to reach a certain configuration determined as a function of the initial conditions. Thus, it might take 50 cycles to reach criterion in one situation, and 54 cycles in another. If one assumes a certain value for the time of each iteration (for example 10 msec), one can account for differences in reaction time.

The use of interactive activation models is certainly a very important advance in our ability to understand reaction time data. However, there are problems. These problems stem from the fact that the human visual system appears to behave more like a feed-forward network than one in which the entire system gradually settles towards a final stable state. At this point it is worth noting that all neural networks can be roughly divided into two basic types. In the first type, units are interconnected, and so each unit is able to influence its targets and vice versa. Such networks form the basis of the interactive activation models we have just been talking about. In the second type of network, the flow of information is essential unidirectional, and the architecture can be described as a feed-forward net. This is characteristic of a variety of networks including perceptrons, multi-layer perceptrons and networks trained by standard learning procedures such as back-propagation. The two-layer network that we saw in the previous section on orientation identification is of this type, because there are no connections from the output layer to the input units.

Feed-forward nets have a number of characteristics that set them apart from more highly interconnected nets of the type used in interactive activation models. They test to be quick, because the output of the system is decided on the basis of a single forward pass. However, at least in conventional connectionist models, the response time of a feed-forward net will be fixed, and as a result, they fail to provide a means of integrating reaction time data.

Is visual processing feed-forward?

The main reason for believing that human visual processing relies mainly on feed-forward mechanisms comes from the speed with which processing can be completed. Although we often fail to appreciate the fact, it is truly remarkable just how quickly we can recognize images flashed before us. In experiments performed in my own group in collaboration with Eva Bonda, we have investigated the ability of normal human subjects to identify briefly presented natural images (famous faces, household objects, well-known locations, logos, etc.)(Thorpe, 1988). Even when presented for only 20 msec, subjects are able to name such stimuli without difficulty. This is true despite the fact that the subjects have no idea about the image to be presented - each subject sees each stimulus only once. Furthermore, identification is still remarkably good when the image is flashed in peripheral vision, i.e. the subject is not looking directly at it.

Such experiments illustrate the remarkable processing power of the human visual system, but unfortunately tell us little about the amount of time taken up by the visual processing itself. Such questions can be tackled more directly by looking at neurophysiological data on the responses of single neurones in the visual system of the monkey, as well as evoked potential recordings in the normal human subject. Perhaps the most striking data concerning the speed of visual processing comes from studies on neuronal responses selective for faces. Over the last decade, a large number of neurophysiological studies have confirmed the existence of neurones in parts of the monkey temporal lobe which have responses that are selective for faces. Such neurones can respond briskly whenever a face is shown to the animal, yet show little or no response to any other class of stimuli (Hasselmo, et al., 1989, Perrett, et al., 1987, Perrett, et al., 1982, Rolls, 1984, Yamane, et al., 1988) . These neurones have a number of fascinating properties, but in the present context, the important thing to note is that they typically start to fire between 100 and 140 ms after the onset of the visual stimulus. In some cases they have even been known to fire after only 80 ms. Other studies using evoked potentials in humans confirm the rapidity of this face selective response; both Jeffreys and Grusser have reported face-selective evoked potentials that start 120-160 ms after stimulus onset (Bötzel and Grusser, 1989, Jeffreys, 1989).

Such neurophysiological data shows that within about one tenth of a second of the onset of a visual stimulus, the visual system has already "decided" whether the image contains a face or not. Anatomical studies show that the part of the temporal lobe where such face-selective neurones are found is probably at least 10 steps away from the photoreceptors of the retina (see figure 4). In order to activate the face-selective neurones, visual information has to go through all ten layers in only 100 or so milliseconds, allowing only 10 milliseconds per layer. It is known that neurones send information in the form of a series of impulsions or action potentials. However, since the firing rate of cortical neurons rarely exceeds one hundred or so spikes per second, this means that in the 10 milliseconds allowed per layer, each neuron is very unlikely to generate more than one action potential. Thus, even if the flow of information through the visual system is entirely feedforward, each neurone can only generate one spike. Any use of feedback loops will add around 10 milliseconds to the processing time, and thus it simply difficult to see how extensive use of feed-back loops can be compatible with the rapidity of visual processing.

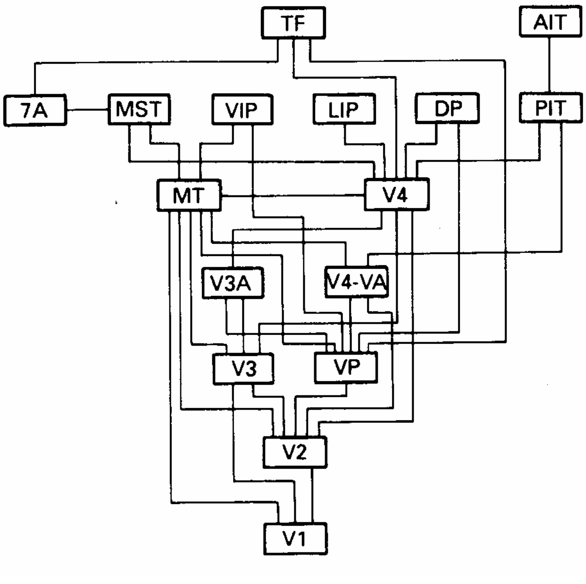

Figure 4: The anatomy of visual cortical areas is complex. In order to reach the face-selective neurones in the anterior inferotemporal lobe (AIT), information has to pass through areas V1, V2, V4 and the posterior inferotemporal areas (PIT) as well as the retina and thalamus (not shown).

Further evidence that feed-forward mechanisms predominate during visual processing comes from neurophysiological studies of the response dynamics of neurones in the visual cortex. We have recorded neurones in the primary visual cortex of the awake monkey to the presentation of visual stimuli varying in either orientation or stereoscopic depth (Celebrini, et al., 1990, Thorpe, et al., 1991, Thorpe, et al., 1989). In both cases, neurones were found to be fully selective at the moment they started to respond, which argues against the possibility that their responses depended on feedback from later stages.

There are thus a number of reasons for thinking that visual processing depends to a large extent on feed-forward mechanisms. But how can we reconcile this with the changes in reaction time that characterize much work in experimental psychology? In the next section I would like to suggest a way in which information could be encoded in neural networks. This coding scheme has the merit of allowing reaction time data to be accounted for in the context of a feed-forward system, that is, without resorting to the use of an interactive activation type architecture.

5. The coding problem

As I pointed out earlier, neurones are known to send information in the form of a series of action potentials. The pioneering experiments of Adrian in the 1920s were the first to show that the firing rate of sensory neurones is a function of stimulus intensity. Since that time, it has generally been thought that neurons send information about their state of activation in the form of a frequency code : the stronger they are activated, the higher is their firing frequency (see figure 5). However, the feasibility of using firing rate or indeed any scheme involving large numbers of spikes per neuron as a way of coding analog information is brought into question by the remarkably high speed of visual processing mentioned in the previous section. Given that the processing time per layer is on average only 10 msec, any neuron firing at less than 100 spikes per second (a situation fairly typical for cortical neurons), will only generate one spike before the next layer has to start responding, and this clearly limits the informational value of a frequency code. There are other ways of transmitting information which use the pattern of firing across a number of neurons (varying the proportion of active neurons or their degree of synchronization (Abeles, 1982, Gray, et al., 1989)), but there is another possibility, largely neglected in the past, based on the fact that when stimulated, the time taken for a neuron to reach threshold for generating a spike depends on the strength of the stimulus. Put simply, the stronger the stimulus, the faster the neuron depolarizes, and the sooner it generates a spike. Consider the effect of flashing an image on an array of receptors such as the one illustrated in figure 6 : a wave-front of spikes would be generated in the output fibres, with the leading spikes corresponding to the points where the intensity was highest, and thus a great deal of information would be contained in the relative arrival times of spikes in different fibres of the optic nerve. Although the example illustrates how a spatial intensity profile could be coded in a wavefront using only one spike per fibre, it should be realized that the same coding could also be used for other stimulus dimensions. Consider what would happen if the six neurons in figure 6 were sensitive not to the intensity in different spatial locations, but rather to contours at different orientations - in such a case, the wavefront would contain information about stimulus orientation.

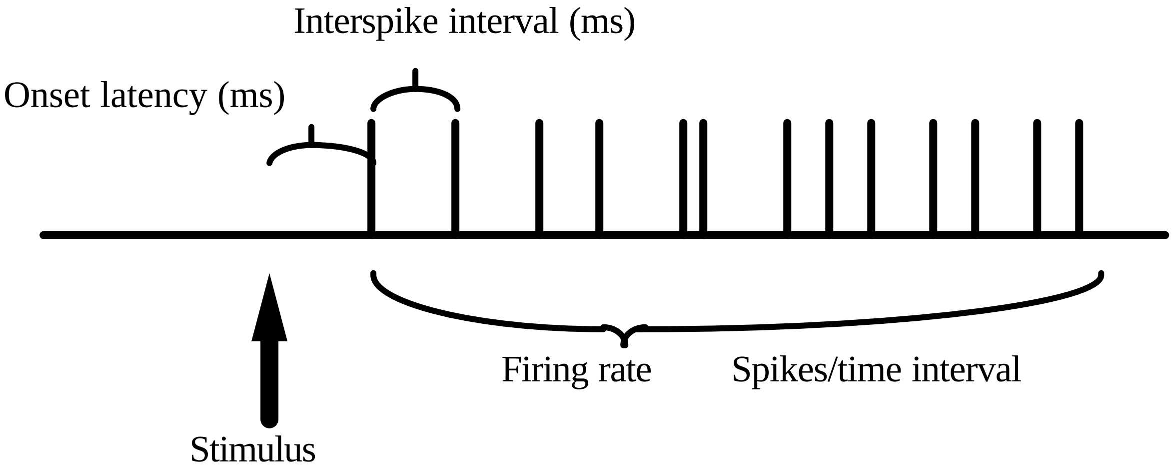

Figure 5: Possible neuronal codes based on the firing pattern of a single neuron. In response to the onset of a stimulus, the neuron can show changes in firing rate (number of spikes in a given time interval) which can also be viewed as a change in the interspike interval. In addition, the onset latency of the unit may contain information (the stronger the stimulus, the shorter the onset latency).

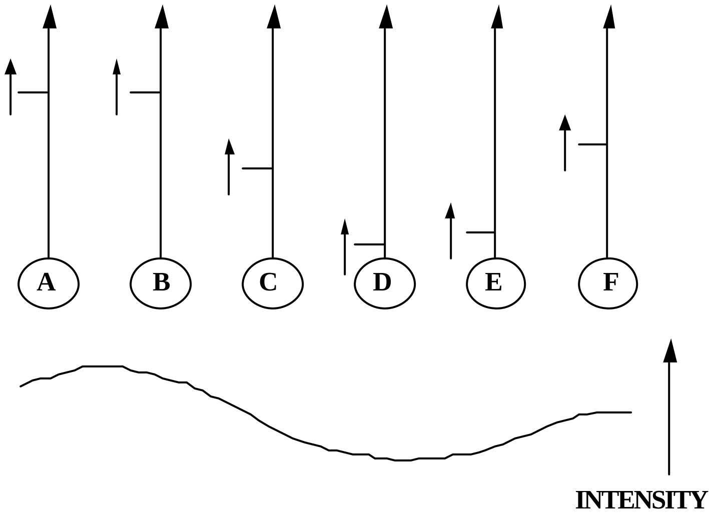

Figure 6: Six units can be used to code an intensity profile on the basis of only one spike per neuron, since the onset latency is a function of stimulus intensity.

Of course, the information contained in the relative timing of spikes in different neurons is only of use if subsequent processing can make use of these differences. There is good evidence that timing differences are indeed used during sensory processing in a variety of domains, including auditory localization (Knudsen, et al., 1987) and the detection of visual movement (Poggio and Koch, 1987). Furthermore, research on the neurobiology of echolocation in bats (Suga, 1990) and the processing of sensory information in electric fish (Kawasaki, et al., 1988) has shown that neurons can be sensitive to delays in the sub-millisecond range. However, in all these cases the timing differences result from the temporal characteristics of the physical stimuli themselves, but in principle, the nervous system could also make use of the temporal delays introduced by the analog to temporal transformation that is a characteristic feature of all neurons. The question therefore arises of whether significant changes in onset latency can also result from changes in stimulus parameters which are neither temporal in nature nor simply related to changes in intensity (there is already considerable evidence that response latencies typically decrease with increases in intensity). For this reason we have looked at the responses of neurons in area V1 of the visual cortex in the awake macaque monkey to stimuli differing in either (i) orientation, or (ii) stereoscopic disparity. In both cases we found evidence that changing either parameter could introduce delays of 10, 20 or even 40 milliseconds.

Could such stimulus-related changes in onset latency be of functional significance? We believe that the answer is yes. It has already been noted that the minimum response latencies for visual cortical neurones vary over a wide range, from less than 40ms to over 100 (Maunsell, 1987, Vogels and Orban, 1991). In principle these latency differences could result from differences in the cellular type of the neurones or their anatomical positions. But the present results show that the nature of the stimulus used is another factor crucial in determining response latency. Consider three structurally similar orientation selective neurones located in the same cortical layer of visual area V1. Because of their similarity, their minimum response latencies may well be identical. Nevertheless, if they are tuned to different orientations, their onset latencies in response to a stimulus with a given orientation may well differ quite substantially (our results suggest that latency shifts of 10 or more msec may be quite common). As a result, the sequence of activation in the three neurones will contain considerable information about stimulus orientation, even if only one spike per neurone is available.

The suggestion that the relative timing of the first spikes in response to a stimulus across an array of neurons may contain much usable information should not be taken to imply that subsequent spikes are not used or that the firing frequency is not an important characteristic of neuronal responses. It is clear that in most situations the two measures will be strongly linked. Nevertheless, the onset latency hypothesis does have the advantage of enabling a considerable amount of processing to be achieved using just the first spike. Furthermore, one very important aspect of using onset latency coding during feed-forward processing in a multi-layer system is that the first information to be sent for processing by later stages will come from the most active neurons. In effect, information will be forwarded for subsequent processing in order of priority, thus reducing the computational load on later stages and reducing the effects of noise.

6. Temporal coding and reaction times

How can the temporal coding model described in the previous section be related to psychological models? As I mentioned earlier, a large proportion of the database in experimental psychology consists of studies investigating the effects of various factors on response times (Luce, 1986). Although many mathematical models of response times have been generated over the years, it has only been relatively recently, with the development of connectionist models, that a neurophysiologically plausible framework has been available in which such effects can be modeled. However, most connectionist models that have attempted to integrate findings about response times have relied on some sort of constraint-satisfaction type process, of the type typical of interactive-activation models. But, as I have argued (Thorpe, 1989, Thorpe and Imbert, 1989), visual processing at least seems to rely on largely feedforward mechanisms, and such networks are generally considered to have a processing time that is fixed.

The situation would seem to be radically different if we take into account the way that neurones convert analog information into delays, as is the case in the temporal coding scheme that I have presented. If we consider that the visual system contains a large number of processing levels (probably at least ten), and that the time taken for visual information to traverse each layer is not fixed but depends on how effective the incoming signals are in activating the neurones, we have the basic form of a model which can account for a considerable range of psychological data.

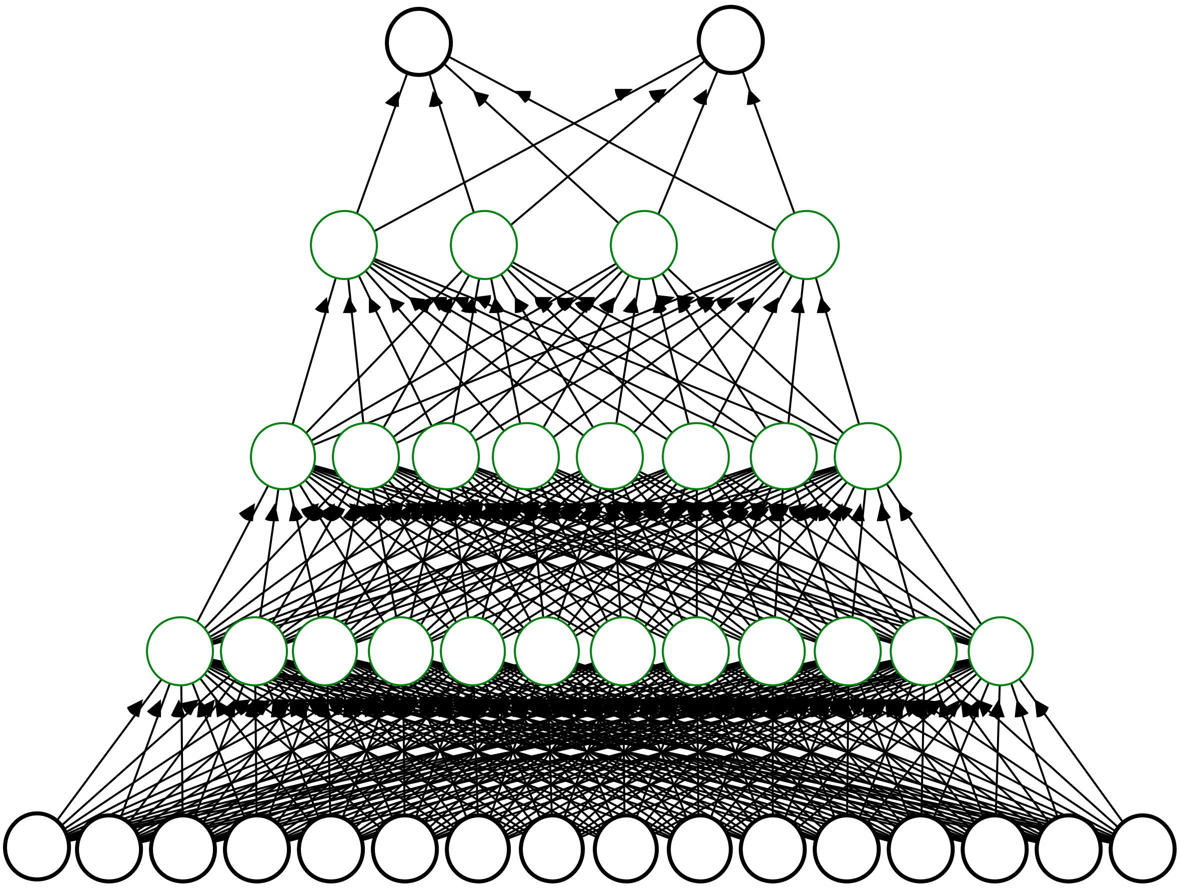

Figure 8: A highly schematic visual system with five layers of units and feed-forward connections. Because of the way in which neurons generate spikes (the delay introduced is inversely proportional to the strength of activation), the time taken to reach the highest layers will depend on many factors, including the nature of the input pattern, the contrast of the stimulus, expectancy, previous experience and so on.

Let us imagine a neural network corresponding roughly to the organization of the visual system (see figure 8). The lowest layer corresponds to the retina - the activity of the units in this layer directly reflect the pattern of intensities found in the original image. Progressively, as we move up through the hierarchy, the receptive field properties become progressively more complex, starting with circular "Mexican-hat" shaped receptive fields, then orientation and spatial frequency tuned units, and so on. If the system was being used for understanding written language, we might find units in even higher levels selective for letters, letter combinations and even words. How might such a system behave if we make the assumption that the units send information as impulsions, and that the time at which these impulsions reflects how strongly the unit is activated - the stronger the activation, the shorter the delay? If the input pattern corresponds to something that the network is set up to process well, activity will flow rapidly through the net, and the output units will become activated after a short delay. If on the other hand, the input pattern is noisy or random, activity will propagate through the net slowly if not at all. Many factors will influence the flow of information through the net. Depending on the way the network has been set up, these could include the familiarity of the input pattern, attention-like phenomena, priming effects and so on. In each case, the time taken to activate the output layers will vary, and such variations correspond naturally to variations in reaction time values found in psychological experiments.

Clearly, this is not the place to go into the details of how such a model could be applied. Nevertheless, I hope that it clear that the addition of extra neurophysiological facts to a connectionist model has considerably enhanced the ability of such models to cope with psychological data. Jacques Gautrais, one of my students, is currently developing a connectionist simulator which incorporates spiking neurons with delays which we hope will provide a way of modeling psychological data in the not too distant future.

7. Conclusions

This paper has tried to make a number of points.

- Interactions between researchers in psychology and neuroscience have been limited by a number of factors. One of the most important has been the absence of a common theoretical framework which can incorporate ideas from both disciplines.

- I think that connectionist models may well provide this common language. However, real interdisciplinarity will still only be possible between workers interested in the same basic objects (e.g. vision, audition, attention, memory, motor control…)
- One way of increasing overlap is to develop connectionist models which take into account realistic neurophysiological data. An example was given of this approach using orientation identification.
- It is important to stress that in such cases, we are not attempting to "reduce" psychology to neuroscience, but rather to "synthesize" psychological data from neurophysiological building blocks.
- One major to using connectionist models to model response time data is that generally, this has relied on using interactive activation type models. For a variety of reasons, visual processing is likely to make particular use of feed-forward mechanisms, which, at least in most connectionist models, do not show response time variations.
- A possible way of avoiding this problem is to introduce a more biologically realistic type of connectionist model in which the units generate impulses, and where analog signals get converted into delays. In a multi-layer feed-forward net using such a coding scheme, the response time (i.e. the time taken for information to pass through all the layers would vary according to conditions. This could well provide a way of modelling a wide range of psychological processes.

Simon J. THORPE
Institut des Neurosciences
Université Pierre et Marie Curie, Paris

Bibliography

Abeles, M. (1982). *Local cortical circuits*. Berlin: Springer.

Anderson, J. A. and Rosenfeld, E. (1988). *Neurocomputing : Foundations of research*. Cambridge: M.I.T. Press.

Arbib, M. A. (1987). *Brains, machines and mathematics. 2nd edition*. Berlin: Springer-Verlag.

Bechtel, W. and Abrahamsen, A. (1991). *Connectionism and the mind : An introduction to parallel processing in networks*. Oxford: Blackwells .

Bharucha, J. J. (1988). Music cognition and perceptual facilitation - A connectionist framework. *Music Perception, 5*, 1-30.

Bötzel, K. and Grusser, O. J. (1989). Electric brain potentials evoked by pictures of faces and non-faces - a search for face-specific EEG potentials. *Experimental Brain Research, 77*, 349-360.

Braitenberg, V. (1984). *Vehicles : Experiments in synthetic psychology*. Cambridge: MIT Press.

Burton, A. M., Bruce, V. and Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology, 81*, 361-380.

Carpenter, G. A. and Grossberg, S. (1990). ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks, 3*, 129-152.

Celebrini, S., Thorpe, S. J., Trotter, Y. and Imbert, M. (1990). Lack of masking in primate V1 neurons : Evidence for feedforward processing of orientation. *Perception, 19*, -.

Churchland, P. S. and Sejnowski, T. J. (1988). Perspectives on cognitive neuroscience. *Science, 242*, 741-745.

Creutzfeldt, O. and Ojemann, G. (1989). Neuronal activity in the human lateral temporal lobe. 3. Activity changes during music. *Experimental Brain Research, 77,* 490-498.

Creutzfeldt, O., Ojemann, G. and Lettich, E. (1989a). Neuronal activity in the human lateral temporal lobe. 1. Responses to speech. *Experimental Brain Research, 77,* 451-475.

Creutzfeldt, O., Ojemann, G. and Lettich, E. (1989b). Neuronal activity in the human lateral temporal lobe. 2. responses to the subjects own voice. *Experimental Brain Research, 77,* 576-489.

DeValois, R. L. and DeValois, K. K. (1988). *Spatial vision*. Oxford: Oxford University Press.

Gray, C. M., König, P., Engel, A. K. and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature, 338*, 334.

Grossberg, S. (1987). Competitive learning : From interactive activation to adaptive resonance. *Cognitive Sci., 11*, 23-64.

Hasselmo, M. E., Rolls, E. T. and Baylis, G. C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behavioural Brain Research, 32*, 203-218.

Heit, C., Smith, M. E. and Halgren, E. (1988). Neural coding of individual words and faces by the human hippocampus and amygdala. *Nature, 333*, 773.

Henry, G. H., Dreher, B. and Bishop, P. O. (1974). Orientation specificity of cells in cat striate cortex. *Journal of Neurophysiology, 37*, 1394-1409.

Hinton, G. E. and Shallice, T. (1991). Lesioning an attractor network - Investigations of acquired dyslexia. *Psychological Review, 98,* 74-95.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities *Proceedings of the National Academy of Sciences, 79,* 2554-2558.

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, 81*, 579-584.

Howard, I. P. (1982). *Human Visual Orientation*. Chichester: Wiley.

Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interactions and functional architecture in the cats visual cortex. *Journal of Physiology, 160*, 106-154.

Jeffreys, D. A. (1989). A face-responsive potential recorded from the human scalp. *Experimental Brain Research, 78*, 195-202.

Kawasaki, M., Rose, G. and Heiligenberg, W. (1988). Temporal hyperacuity in single neurons of electric fish. *Nature, 336*, 173-176.

Kehoe, E. J. (1989). Connectionist models of conditioning - A tutorial. *Journal of the Experimental Analysis of Behavior, 52*, 427-440.

Knudsen, E. I., du Lac, S. and Esterley, S. D. (1987). Computational maps in the brain. *Annual Review of Neuroscience, 10*, 41-66.

Kutas, M., Hillyard, S. A. and Gazzaniga, M. S. (1988). Processing of semantic anomaly by right and left hemispheres of commussurotomy patients. Evidence from event-related potentials. *Brain, 111*, 553-576.

LeDoux, J. E. and Hirst, W. (1986). *Mind and brain : Dialogues in cognitive neuroscience*. Cambridge: Cambridge University Press.

Lehky, S. R. and Sejnowski, T. (1988). Network model of shape-from-shading : neural function arises from both receptive and projective fields. *Nature, 333*, 452-.

Lehky, S. R. and Sejnowski, T. J. (1989). Simplifying network models of binocular rivalry and shape-from-shading. In C. Koch and I. Segev (Eds.), *Methods in neuronal modeling.* Boston: MIT Press.

Lehky, S. R. and Sejnowski, T. J. (1990). Neural models of stereoacuity and depth interpolation based on a distributed representation of stereo disparity *Journal of Neuroscience, 10*, 2281-2299.

Luce, R. D. (1986). *Response times : Their role in inferring elementary operations*. Oxford: Oxford University Press.

Maunsell, J. H. R. (1987). Physiological evidence for two visual subsystems. In L. M. Vaina (Eds.), *Matters of Intelligence* Dodrecht: D. Reidel.

McClelland, J. L. and Rumelhart, D. E. (1981). An interaction activation model of context effects in letter perception. Part 1. An account of the basic findings. *Psychological Review, 88*, 375-407.

McClelland, J. L. and Rumelhart, D. E. (1986). *Parallel distributed processing : Explorations in the microstructure of cognition. Vol. 2. Psychological and biological models*. Boston: MIT Press.

Memmi, D. (1990). Connexionnisme, intelligence artificielle, et modelisation cognitive. *Intellectica, 9*, 41-80.

Mozer, M. C. (1991). *The recognition of multiple objects.* Boston: MIT Press.

Orban, G. A. (1984). *Neuronal operations in the visual cortex*. Berlin: Springer-Verlag.

Paller, K. A., Kutas, M. and Mayes, A. R. (1987). Neural correlates of encoding in an incidental learning paradigm. *EEG. Clin. Neurophysiol., 67*, 360-371.

Perrett, D. I., Mistlin, A. J. and Chitty, A. J. (1987). Visual neurons responsive to faces. *Trends in Neuroscience, 10*, 358.

Perrett, D. I., Rolls, E. T. and Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47*, 329-342.

Petersen, S. E., Fox, B. T., Posner, M. I., Mintun, M. and Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single word processing. *Nature, 331*, 585-589.

Phaf, R. H., Van der Heijden, A. H. C. and Hudson, P. T. W. (1990). A connectionist model for attention in visual selection tasks. *Cognitive Psychology, 22*, 273-342.

Plunkett, K. and Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron - Implications for child language acquisition. *Cognition, 38*, 103.

Poggio, T. and Koch, C. (1987). Synapses that compute motion. *Scientific American, 256*, 42-48.

Posner, M. I., Inhoff, A. W., Friedrich, F. J. and Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology, 15*, 107-121.

Posner, M. I, Petersen, S. E., Fox, P. T. and Raichle, M. E. (1988). Localization of cognitive operations in the brain. *Science, 240*, 1627-1631.

Pouget, A. and Thorpe, S. J. (1991). Connectionist models of orientation identification. *Connection Science, 3*, 127-142.

Requin, J. (1987). Les neurosciences cognitives: au-delà du réductionisme, une science de synthèse? In J. Siguan (Eds.), *Comportement, Cognition, Conscience* (pp. 31-57). Paris: P.U.F.

Robinson, D. L. and Rugg, M. D. (1988). Latencies of visually responsive neurons in various regions of the rhesus monkey brain and their relation to human visual responses. *Biological Psychology, 26*, 111-116.

Rolls, E. T. (1984). Neurons in the cortex of the temporal lobe and in the amygdala of the monkey with responses selective for faces. *Human Neurobiology, 3,* 209-222.

Rugg, M. D. (1990). Event-related brain potentials dissociate repetition effects of high- and low-frequency words. *Memory Cognition, 18,* 367-379.

Rumelhart, D. E. and McClelland, J. L. (1982). An interaction activation model of context effects in letter perception. Part 2: The contextual enhancement effect and some tests and extensions of the model *Psychological Review, 89,* 60-94.

Rumelhart, D. E. and McClelland, J. L. (1986). *Parallel distributed processing : Explorations in the microstructure of cognition. Vol. 1. Foundations.* Boston: MIT Press.

Sandon, P. A. (1990). Simulating visual attention. *Journal of Cognitive Neuroscience, 2,* 213-232.

Scheibel, A. B. and Wechsler, A. F. (1990). *Neurobiology of higher cognitive function.* New York: Guilford Press.

Seidenberg, M. S. and McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review, 96,* 523-568.

Selfridge, O. (1959). Pandemonium: A paradigm for learning. In *The mechanisation of the thought processes* London: Her Majesties Stationary Office.

Suga, N. (1990). Biosonar and Neural Computation in Bats *Scientific American , 262 (6),* 34-41.

Thompson, R. F. (1990). Cognitive neuroscience in Europe. *Psychological Science, 1,* 287-293.

Thorpe, S. J. (1988). Identification of rapidly presented images by the human visual system. *Perception, 17,* A77.

Thorpe, S. J. (1989). Traitement d'images chez l'homme. *Techniques Sciences Informatiques, 7,* 517-525.

Thorpe, S. J., Celebrini, S., Trotter, Y. and Imbert, M. (1991). Dynamics of stereo processing in area V1 of the awake primate. *European Journal of Neuroscience Supplement, 4,* 83.

Thorpe, S. J., Celebrini, S., Trotter, Y., Pouget, A. and Imbert, M. (1989). Dynamic aspects of orientation coding in area V1 of the awake primate. *European Journal of Neuroscience, Suppl. 2,* 322.

Thorpe, S. J. and Imbert, M. (1989). Biological constraints on connectionist models. In R. Pfeifer, Z. Schreter, F. Fogelman-Soulié and L. Steels (Eds.), *Connectionism in Perspective.* Amsterdam: Elsevier.

Thorpe, S. J. and Pouget, A. (1989). Coding of orientation by the visual cortex : Neural network modelling. In R. Pfeifer, Z. Schereter, F. Fogelman-Soulié and L. Steels (Eds.), *Connectionism in Perspective.* Amsterdam: Elsevier.

Vogels, R. and Orban, G. A. (1991). Quantitative study of striate single unit responses in monkeys performing an orientation discrimination task. *Experimental Brain Research, 84*, 1-11.

Wurtz, R. H., Goldberg, M. E. and Robinson, D. L. (1982). Brain mechanisms of visual attention. *Scientific American, 246*, 124-135.

Yamane, S., Kawano, K., Kaji, S. and Komatsu, H. (1988). Neurons responsive to faces in the inferotemporal cortex of the monkey. *Neuroscience Research, S9*, S179.

Zipser, D. and Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature, 331*, 679-684.