Progress in Brain Research, Vol. 144
ISSN 0079-6123

CHAPTER 6

Consciousness absent and present: a neurophysiological exploration

Edmund T. Rolls*

University of Oxford, Department of Experimental Psychology, Oxford OX1 3UD, UK

Abstract: Backward masking was used to investigate the amount of neuronal activity that occurs in the macaque inferior temporal visual cortex when faces can just be identified. It is shown that the effect of the pattern mask is to interrupt neuronal activity in the inferior temporal visual cortex. This reduces the number of action potentials that occur to a given stimulus, and decreases even more the information that is available about which stimulus was shown because the variance of the spike counts is increased. When the onset of the mask follows the onset of the test stimulus by 20 ms, each neuron fires for approximately 30 ms, provides on average 0.06 bits of information, and human observers perform at approximately 50% better than chance in forced choice psychophysics, yet say that they are guessing, and frequently report that they are unable to consciously see the face and identify which face it is. At a longer Stimulus Onset Asynchrony of 40 ms, the neurons fire for approximately 50 ms, the amount of information carried by a single neuron is 0.14 bits, and human observers are much more likely to report conscious identification of which face was shown. The results quantify the amount of neuronal firing and information that is present when stimuli can be discriminated but not reported on consciously, and the additional amount of neuronal firing and information that is required for humans observers to consciously identify the faces. It is suggested that the threshold for conscious visual perception may be set to be higher than the level at which small but significant information is present in neuronal firing, so that the systems in the brain that implement the type of information processing involved in conscious thoughts are not interrupted by small signals that could be noise in sensory pathways.

Introduction

Damage to the primary (striate) visual cortex can result in blindsight, in which patients report that they do not see stimuli consciously, yet when making forced choices can discriminate some properties of the stimuli such as motion, position, some aspects of form, and even face expression (Weiskrantz et al., 1974; Stoerig and Cowey, 1997; Weiskrantz, 1997, 1998; De Gelder et al., 1999). In normal human subjects, backward masking of visual stimuli, in which another visual stimulus closely follows the short presentation of a test stimulus, reduces the visual perception of the test visual stimulus, and this paradigm has been widely used in psychophysics (Humphreys and Bruce, 1989). In this chapter the author considers how much information is present in neuronal firing in the part of the visual system that represents faces and objects, the inferior temporal visual cortex (Rolls and Deco, 2002), when human subjects can discriminate in forced choice, but cannot consciously perceive, face identity. The author also considers the implications that the neurophysiological findings have for consciousness.

The representation of faces and objects is in the inferior temporal visual cortex as shown by evidence that position, size and even for some neurons view, invariant representations of objects and faces are provided by neurons in the inferior temporal visual cortex (Rolls, 2000a; Rolls and

---

*Corresponding author. Tel.: +44-1865-271348;
Fax: +44-1865-310447; Web: www.cns.ox.ac.uk;
E-mail: edmund.rolls@psy.ox.ac.uk

DOI: 10.1016/S0079-6123(03)14400-6

Deco, 2002); that this is the last stage of unimodal visual processing in primates; and that lesions of what may be a homologous region in humans, the fusiform gyrus face and object areas (Ishai et al., 1999; Kanwisher et al., 1997), produce face and object identification deficits in the absence of low-level impairments of visual processing such as visual acuity (Rolls and Deco, 2002; Farah, 1990; Farah et al., 1995a,b). The inferior temporal visual cortex is, therefore, an appropriate stage of processing at which to relate quantitative aspects of neuronal processing to the visual perception of faces and objects. We have, therefore, studied the quantitative relationship between neuronal activity in the macaque inferior temporal visual cortex and visual perception (Rolls and Deco, 2002), and in this article the focus is on the relation between inferior temporal visual cortex and conscious visual perception, using the results from combined neurophysiological studies on the inferior temporal visual cortex and perceptual studies in humans with the paradigm of backward masking of visual stimuli (Rolls et al., 1994, 1999; Rolls and Tovée, 1994). A subsequent study by Kovacs et al. (1995) using a similar backward masking paradigm combined with primate electrophysiology confirmed the results.

Neurophysiology of the backward masking of visual stimuli

Rolls and Tovée (1994) and Rolls et al. (1994) measured the responses of single neurons in the macaque inferior temporal visual cortex during backward visual masking. Neurons that were selective for faces, using distributed encoding (Rolls and Tovée, 1995; Rolls et al., 1997; Treves et al., 1999; Rolls and Deco, 2002), were tested in a visual fixation task run as shown in Fig. 1. The visual

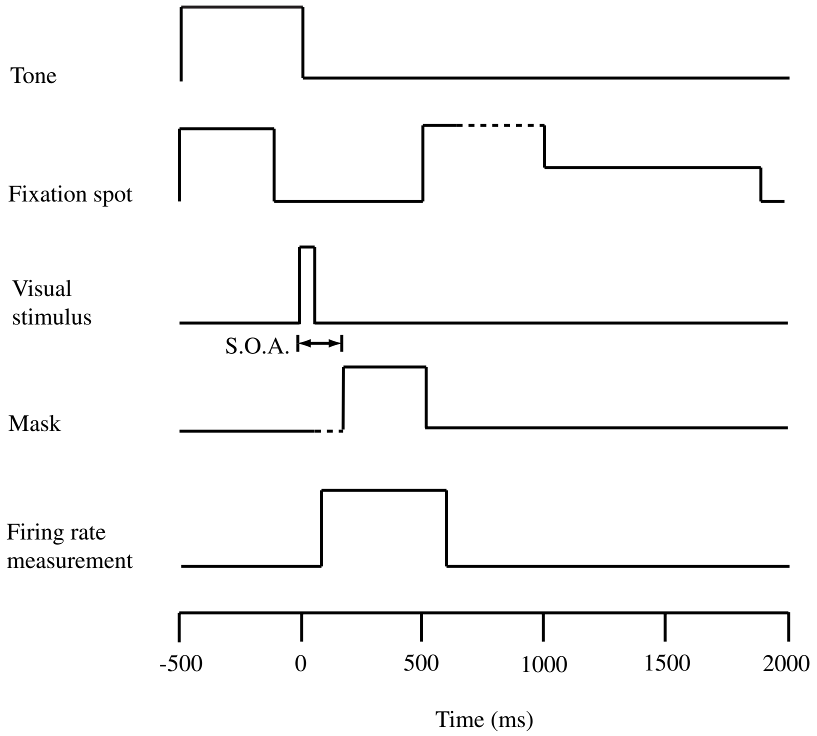

Fig. 1. The timing used in the backward masking visual fixation blink task. The Stimulus Onset Asynchrony is the time between the onset of the visual test stimulus and the onset of the pattern mask stimulus. The test stimulus duration was 16 ms. (After Rolls et al., 1994.)

fixation task was used to ensure that the monkey looked at the visual stimuli.

The methods used are described by Rolls and Tovée (1994) and Rolls et al. (1994), and a few salient points follow. As shown in Fig. 1, at −100 ms, the fixation spot was blinked off so that there was no stimulus on the screen in the 100 ms period immediately preceding the test image. The screen in this period, and at all other times including the inter-stimulus interval and the interval between the test image and the mask, was set at the mean luminance of the test images and the mask, so that pattern discrimination with equally intense test and mask stimuli was investigated (see Bruce and Green, 1989). At 0 ms, the 500 ms warning cure tone was switched off and the test visual image was switched on for one 16 ms frame of a raster display image. The monitor had a persistence of less than 3 ms, so that no part of the test image was present at the start of the next frame. Stimulus Onset Asynchrony (S.O.A.) values of 20, 40, 60, 100 or 1000 ms (chosen in a random sequence by the computer) were used. (The Stimulus Onset Asynchrony is the time between the onset of the test stimulus and the onset of the mask.) The duration of the masking stimulus was 300 ms. The stimuli were static visual stimuli subtending 8 degrees in the visual field presented on a video monitor at a distance of 1.0 m. The faces used as test stimuli are illustrated in Fig. 2. The usual masking stimulus (to which the neuron being analysed did not respond) was made up of letters of the alphabet {N,O}, as shown in Fig. 2. The masking pattern consisted of overlapping letters, and this masking pattern was used because it is similar to the mask used in the previous psychophysical experiments (see Rolls et al., 1994). (In some cases the masking stimulus was a face

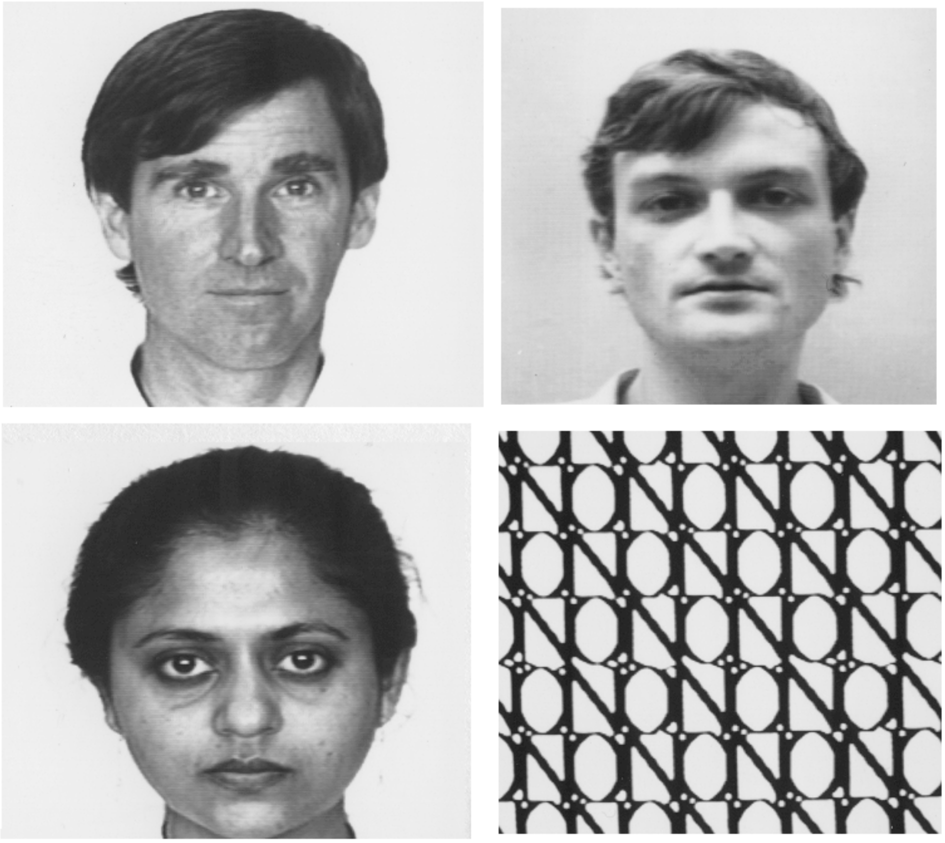

Fig. 2. Examples of the test images used. The mask is also shown. (After Rolls et al., 1994.)

stimulus that was ineffective for the neuron being recorded.)

Figure 3 shows examples of the effects of backward masking on the responses of a single inferior temporal cortex neuron in peristimulus rastergram and time histogram form. The top rastergram/spike density histogram pair shows the responses of the neuron to a single frame of the test stimulus (an effective face stimulus for that neuron). Relative to the prestimulus rate, there was an increase in the firing produced with a latency of approximately 75 ms, and this firing lasted for 200–300 ms, that is for much longer than the 16 ms presentation of the target stimulus. In the next pairs down, the effects of introducing a non-effective face as the masking stimulus with different S.O.A.s are shown. It is shown that the effect of the mask is to limit the duration of the firing produced by the target stimulus. Very similar masking was obtained with the standard N–O pattern mask. Similar experiments were repeated on 42 different cells (Rolls et al., 1994; Rolls and Tovée, 1995), and in all cases the

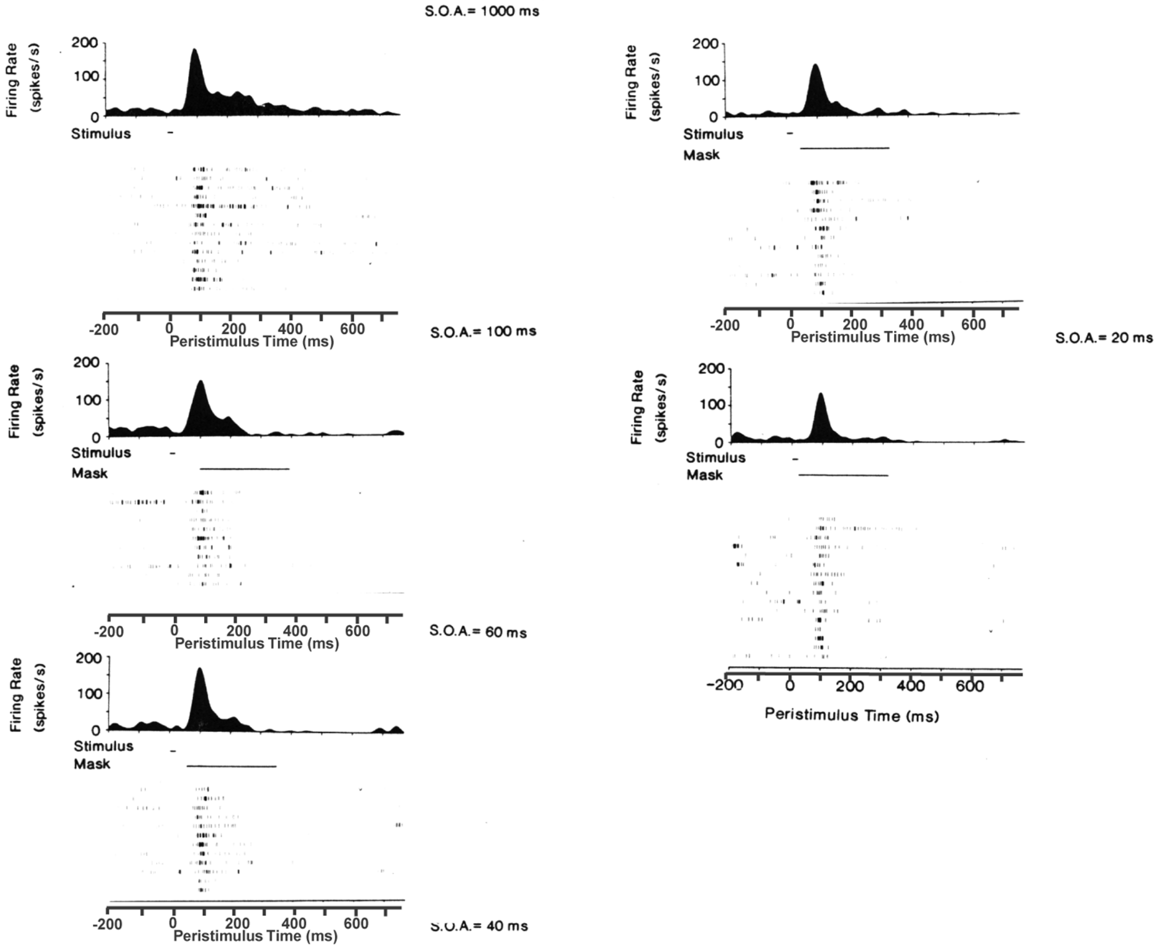

Fig. 3. Peristimulus rastergrams and smoothed peristimulus spike density histograms based on responses in 8–16 trials to the test face alone (top raster-histogram pair), and to the test face followed by a masking stimulus (which was a face that was ineffective in activating the cell) with different S.O.A. values. (S.O.A. = Stimulus Onset Asynchrony) The mask alone did not produce firing in the cell. The target stimulus was shown for 16 ms starting at time 0. (The top trace shows the response to the target stimulus alone, in that with this 1000 ms S.O.A., the mask stimulus was delayed until well after the end of the recording period shown.) (After Rolls and Tovée, 1994.)

temporal aspects of the masking were similar to those shown in Fig. 3.

One important conclusion from these results is that the effect of a backward masking stimulus on cortical visual information processing is to limit the duration of neuronal responses, by interrupting neuronal firing. The neuronal firing of inferior temporal cortex neurons often persisted for 200–300 ms after a 16 ms presentation of a stimulus. With a 20 ms Stimulus Onset Asynchrony, the neuronal firing was typically limited to 30 ms. With a 40 ms Stimulus Onset Asynchrony, the neuronal firing was typically limited to 50 ms. This persistence of cortical neuronal firing when a masking stimulus is not present is probably related to cortical recurrent collateral connections which could implement an autoassociative network with attractor and short-term memory properties (see Rolls and Treves, 1998; Rolls and Deco, 2002), because such continuing post-stimulus neuronal firing is not observed in the lateral geniculate nucleus (K. Martin, personal communication).

Information available in inferior temporal cortex visual neurons during backward masking

To fully understand quantitatively the responses of inferior temporal cortex neurons at the threshold for visual perception, Rolls et al. (1997) applied information theoretic methods (see Shannon, 1948; Rolls and Treves, 1998; Rolls and Deco, 2002) to the analysis of the neurophysiological data with backward masking obtained by Rolls et al. (1994) and Rolls and Tovée (1994). One advantage of this analysis is that it shows how well the neurons discriminate between the stimuli under different conditions, by taking into account not only the number of spikes, but also the variability from trial to trial in the number of spikes. Another advantage of this analysis is that it evaluates the extent to which the neurons discriminate between stimuli in bits, which can then be directly compared with evidence about discriminability obtained using different measures, such as human psychophysical performance. The analysis quantifies what can be determined about which of the set of faces was presented from a single trial of neuronal firing.

As a preliminary to the information theoretic analysis, the effect of the S.O.A. on the neuronal responses, averaged across the population of 15 neurons for which a sufficient number of trials was available, is shown in Fig. 4. The responses for the most (max) and the least (min) effective stimuli are shown for the period 0–200 ms with respect to stimulus onset. There was little effect (not significant) of the mask on the responses to the least effective stimulus in the set, for which the number of spikes was close to the spontaneous activity.

The transmitted information carried by neuronal firing rates about the stimuli was computed with the use of techniques that have been described previously (e.g., Rolls et al., 1997; Rolls and Treves, 1998; Rolls and Deco, 2002), and have been used previously to analyse the responses of inferior temporal cortex neurons (Optican and Richmond, 1987; Gawne and Richmond, 1993; Tovée et al., 1993; Tovée and Rolls, 1995; Rolls et al., 1997). In brief, the general procedure was as follows (Rolls et al., 1999). The response r of a neuron to the presentation of a particular stimulus s was computed by measuring the firing rate of the neuron in a fixed time window after

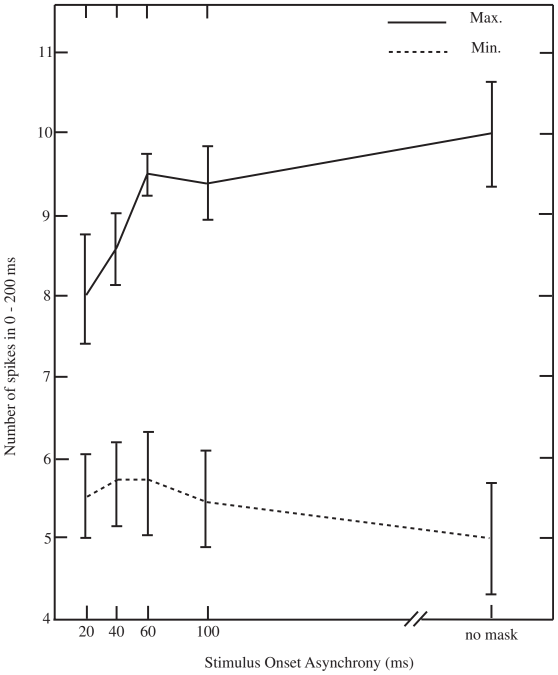

Fig. 4. The mean (±sem) across cells of the number of spikes produced by the most effective stimulus (max) and the least effective stimulus (min) as a function of Stimulus Onset Asynchrony (S.O.A.). (After Rolls et al., 1999.)

the stimulus presentation. The firing rates were then quantized into a smaller number of bins d than there were trials for each stimulus. After this response quantization, the experimental joint stimulus-response probability table $P(s, r)$ was computed from the data (where $P(r)$ and $P(s)$ are the experimental probability of occurrence of responses and of stimuli respectively), and the information $I(S, R)$ transmitted by the neurons averaged across the stimuli was calculated by using the Shannon formula (Shannon, 1948; Rolls and Deco, 2002):

$$I(S, R) = \sum_{s, r} P(s, r) \log_2 \frac{P(s, r)}{P(s)P(r)}$$

and then subtracting the finite sampling correction of Panzeri and Treves (1996), to obtain estimates unbiased for the limited sampling. This leads to the information available in the firing rates about the stimulus.

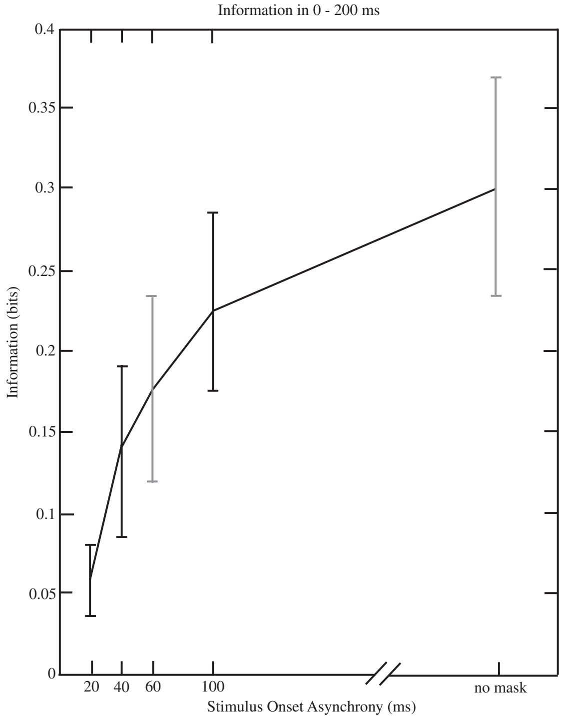

Fig. 5. The average (±sem) across the cells of the cumulated information available in a 200 ms period from stimulus onset from the responses of the cells as a function of the S.O.A. (After Rolls et al., 1999.)

Figure 5 shows the average across the cells of the cumulated information available in a 200 ms period from stimulus onset from the responses of the 15 neurons as a function of the S.O.A. This emphasizes how as the S.O.A. is reduced towards 20 ms the information does reduce rapidly, but that nevertheless at an S.O.A. of 20 ms there is still considerable information about which stimulus was shown. The reduction of the information at different S.O.A.s was highly significant (one way ANOVA) at $P < 0.001$. It was notable that the information reduced much more than the number of spikes on each trial as the S.O.A. was shortened. The explanation for this is that at short S.O.A.s the neuronal responses become noisy, as shown by Rolls et al. (1999). This emphasizes the value of measuring the information available, and not only the number of spikes (Rolls et al., 1999).

Human psychophysical performance with the same set of stimuli

Rolls et al. (1994) performed human psychophysical experiments with the same set of stimuli and with the same apparatus used for the neurophysiological experiments so that the neuronal responses could be closely related to the identification that was possible of which face was shown. The monitor provided maximum and minimum luminance of 6 and 0.13 footlamberts, and was adjusted internally for linearity, within an error of no more than 3%.

Five different faces were used as stimuli. All the faces were well known to each of the eight observers used in the experiment. In the forced choice paradigm, the observers specified whether the face was normal or rearranged, and identified whose face they thought had been presented. Even if the observers were unsure of their judgement they were instructed to respond with their best guess. The data were corrected for guessing to aid in comparison between classification and identification. This correction arranged that chance performance would be shown as 0% correct on the graphs, and perfect performance as 100% correct.

The mean proportion of correct responses for the identification task (and for the classification of normal versus rearranged) are shown in Fig. 6. The

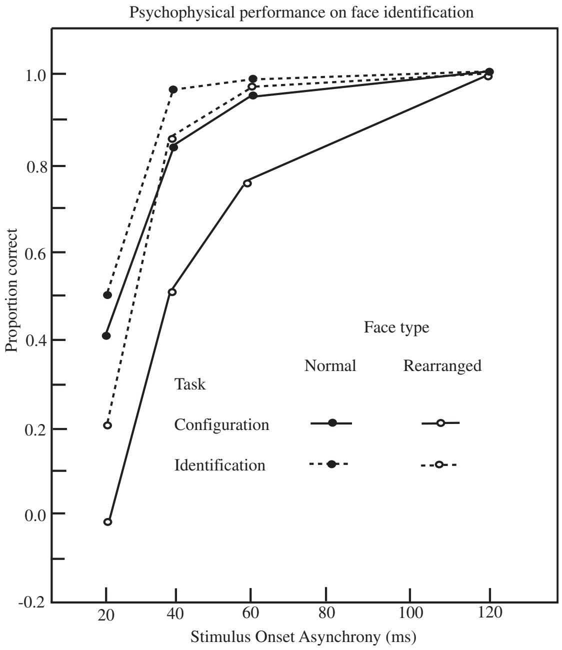

Fig. 6. Psychophysical performance of humans with the same stimuli as used in the neurophysiological experiments described here. The subjects were shown 'normal' faces or faces with the parts 'rearranged', and were asked to state which of 5 faces had been shown ('Identification'), and whether the face shown was in the normal or rearranged 'Configuration'. The plots show the proportion correct on the tasks of classification of spatial configuration (i.e., whether the face features were normal or rearranged), and of determination of the identity of the faces for faces in the Normal or Rearranged spatial configuration of face features, as a function of Stimulus Onset Asynchrony (S.O.A). The data have been corrected for guessing. The means of the proportions correct are shown. The test stimulus was presented for 16 ms. (After Rolls et al., 1994.)

proportion correct data was submitted to an arc sin transformation (to normalise the data) and a repeated measures ANOVA was performed. This analysis showed statistically significant effects of S.O.A. [$F(4,28) = 61.52$, $P < 0.0001$].

Forced choice discrimination of face identity was thus better than chance at an S.O.A. of 20 ms. However, at this S.O.A., the subjects were not conscious of seeing the face, or of the identity of the face, and felt that their guessing about which face had been shown was not correct. The subjects did know that something had changed on the screen (and this was not just brightness, as this was constant throughout a trial). Sometimes the subjects had some conscious feeling that a part of a face (such as a mouth) had been shown. However, the subjects were not conscious of seeing a whole face, or of seeing the face of a particular person. At an S.O.A. of 40 ms, the subjects' forced choice performance of face identification was close to 100% (see Fig. 6), and at this S.O.A., the subjects became much more consciously aware of the identity of which face had been shown (Rolls et al., 1994).

Discussion

The neurophysiological data (Rolls et al., 1994; Rolls and Tovée, 1994), and the results of the information theoretic analysis (Rolls et al., 1999), can now be compared directly with the effects of backward masking in human observers, studied in the same apparatus with the same stimuli (Rolls et al., 1994). For the human observers, identification of which face from a set of six had been seen was 50% correct (with 0% correct corresponding to chance performance) with an S.O.A. of 20 ms, and 97% correct with an S.O.A. of 40 ms (Rolls et al., 1994). Comparing the human performance purely with the changes in firing rate under the same stimulus conditions suggested that when it is just possible to identify which face has been seen, neurons in a given cortical area may be responding for only approximately 30 ms (Rolls and Tovée, 1994; Rolls et al., 1994). The implication is that 30 ms is enough time for a neuron to perform sufficient computation to enable its output to be used for identification. The results based on an analysis of the information encoded in the spike trains at different S.O.A.s support this hypothesis by showing that a significant proportion of information is available in these few spikes (see Fig. 5), with on average 0.06 bits available from each neuron at an S.O.A. of 20 ms. Thus when subjects feel that they are guessing, and are not conscious of seeing whose face has been shown, macaque inferior temporal cortex neurons provide small but significant amounts of information about which face has been shown. When the S.O.A. was increased to 40 ms, the inferior temporal cortex neurons responded for approximately 50 ms, and encoded approximately

0.14 bits of information (in a period of 200 ms, for the subset of face-selective neurons tested, see Rolls et al., 1999). At this S.O.A., not only was face identification 97% correct, but the subjects were much more likely to be able to report consciously seeing a face and/or whose face had been shown. One way in which the conscious perception of the faces was measured quantitatively was by asking subjects to rate the clarity of the faces. This was a subjective assessment and therefore reflected conscious processing, and was made using magnitude estimation. It is shown in Fig. 7 that the subjective clarity of the stimuli was low at 20 ms S.O.A., was higher at 40 ms S.O.A., and was almost complete by 60 ms S.O.A.

It is suggested that the threshold for conscious visual perception may be set to be higher than the level at which small but significant sensory information is present so that the systems in the brain that implement the type of information processing involved in conscious thoughts is not interrupted by small signals that could be noise in sensory pathways. Consideration of the nature of this processing, and the reason why it may be useful to not interrupt it unless there is a definite signal that may require the use of the type of processing that can be performed by the conscious processing system, is left to the end of the Discussion because the issues raised necessarily involve hypotheses that are not easy to test.

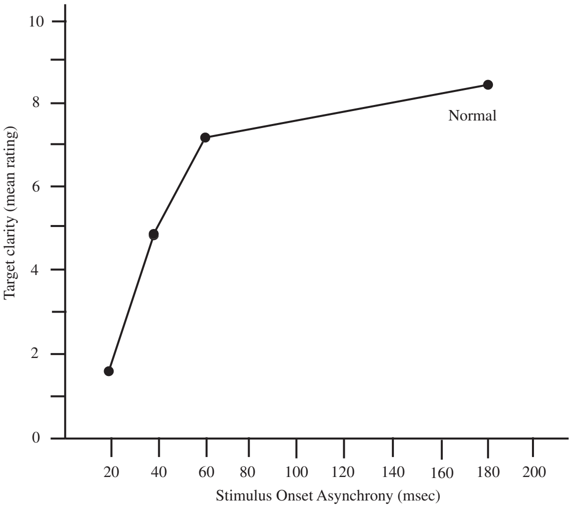

Fig. 7. Rating of the subjective clarity of faces as a function of stimulus onset asynchrony (S.O.A.). The mean of the magnitude estimation ratings is shown. (The clarity of the face, when presented without being followed by a masking stimulus, was assigned the number 10. If the observer was not able to see the features of the face, that was to be considered 0. The observers assigned a number from 0 to 10 to represent the perceived subjective clarity of the face. The effect of S.O.A. was statistically significant at $P < 0.001$. After Rolls et al., 1994.)

The results of the information analysis (Rolls et al., 1999) emphasise that very considerable information about which stimulus was shown is available in a short epoch of, for example, 50 ms. This confirms the findings of Tovée et al. (1993), Tovée and Rolls (1995) and Heller et al. (1995), and facilitates the rapid read-out of information from the inferior temporal visual cortex, and the use of whatever information is available in the limited period of firing under backward masking conditions. It was notable that the information in the no mask condition did outlast the end of the stimulus by as much as 200–300 ms, indicating some short term memory trace property of the neuronal circuitry. This continuing activity could be useful in the learning of invariant representations of objects (Rolls, 1992, 2000a; Wallis and Rolls, 1997; Rolls and Deco, 2002). The results also show that even at the shortest S.O.A. of 20 ms, the information available was on average 0.06 bits. This compares to 0.3 bits with the 16 ms stimulus shown without the mask (Rolls et al., 1999). It also compares to a typical value for such neurons of 0.35–0.5 bits with a 500 ms stimulus presentation (Tovée and Rolls, 1995; Rolls et al., 1997). The results thus show that considerable information (33% of that available without a mask, and approximately 22% of that with a 500 ms stimulus presentation) is available from neuronal responses even under backward masking conditions which allow the neurons to have their main response in 30 ms. Also, we note that the information available from a 16 ms unmasked stimulus (0.3 bits) is a large proportion (approximately 65–75%) of that available from a 500 ms stimulus.

These results provide evidence on how rapid the processing of visual information is in a cortical area, and provides a fundamental constraint for understanding how cortical information processing operates (see Rolls and Treves, 1998; Rolls and Deco, 2002). One direct implication of the 30 ms firing with the 20 ms S.O.A. is that this is sufficient time both for a cortical area to perform its computation, and for the information to be read out from a cortical area, given that psychophysical performance is 50% correct at this S.O.A. Another implication is that the recognition of visual stimuli can be performed using feedforward processing in the multi-stage hierarchically organized ventral visual system comprising at least V1–V2–V4-Inferior Temporal Visual Cortex, in that the typical shortest neuronal response latencies in macaque V1 are approximately 40 ms, and increase by approximately 15–17 ms per stage to produce a value of approximately 90 ms in the inferior temporal visual cortex (Rolls and Deco, 2002; Oram and Perrett, 1992; Dinse and Kruger, 1994; Raiguel et al., 1989; Vogels and Orban, 1994; Nowak and Bullier, 1997). (The fact that considerable information is available in short epochs, of for example 20 ms, of the firing of neurons provides part of the underlying basis for this rapid sequential activation of connected visual cortical areas (Tovée and Rolls, 1995; Rolls and Deco, 2002).) Given these timings, it would not be possible in the 20 ms S.O.A. condition for inferior temporal cortex neuronal responses to feedback to influence V1 neuronal responses to the test stimulus before the mask stimulus produced its effects on the V1 neurons. This shows that at least some recognition of visual stimuli is possible without top–down backprojection effects from the inferior temporal visual cortex to early cortical processing areas.

The processing time allowed for each cortical area to perform useful computation and for the information to be read out to the next stage of cortical processing is in the order of 15–17 ms as shown by the neuronal response latency increases from cortical area to cortical area noted above; and less than 30 ms as shown by the duration of the firing in the 20 ms S.O.A. condition when face identification was 50% better than chance. This is sufficient for recurrent collaterals to operate by feedback within a cortical area to allow them to implement attractor-based processing, as shown by analyses of the speed of settling of such networks provided that they are implemented with neurons with continuous dynamics (as implemented in models by integrate-and-fire neurons) and with spontaneous firing (Treves, 1993; Battaglia and Treves, 1998; Rolls and Treves, 1998). Indeed, the dynamics of a four-layer hierarchical network with an architecture like that of the ventral visual system are sufficiently rapid to allow recurrent feedback attractor operations to contribute usefully to information processing in the system (Panzeri et al., 2001; Rolls and Deco, 2002; contrast with Thorpe et al., 1996).

The inferior temporal visual cortex is an appropriate stage to analyse the identification of objects and faces, and to link to face identification in humans, because the ITC contains an invariant representation of faces and objects (Booth and Rolls, 1998; Rolls, 2000a; Rolls and Deco, 2002), and damage to corresponding areas in humans may produce face and object agnosias (Farah, 1990; Farah et al., 1995a,b).

The quantitative analyses of neuronal activity in an area of the ventral visual system involved in face and object identification described here which show that significant neuronal processing can occur that is sufficient to support forced choice but implicit (unconscious) discrimination in the absence of conscious awareness of the identity of the face is of interest in relation to studies of blindsight (Weiskrantz et al., 1974; Weiskrantz, 1997, 1998; Stoerig and Cowey, 1997; de Gelder et al., 1999). The issue in blindsight is that conscious reports that a stimulus has been seen cannot usually be made, and yet some forced choice performance is possible with respect to, for example, the motion, position and some aspects of the form of the visual stimuli. It has been argued that the results in blindsight are not due just to reduced visual processing, because some aspects of visual processing are less impaired than others (Weiskrantz, 1997, 1998, 2001; Azzopardi and Cowey, 1997). However, it is suggested that some of the visual capacities that do remain in blindsight reflect processing via visual pathways that are alternatives to the V1 processing stream (Weiskrantz, 1997, 1998, 2001). If some of those pathways are normally involved in implicit processing, this may help to give an account of why some implicit (unconscious) performance is possible in blindsight patients. Further, it has been suggested that ventral visual stream processing is especially involved in consciousness, because it is information about objects and faces that needs to enter a system to plan actions (Milner and Goodale, 1995; Rolls and Deco, 2002); and the planning of actions that involves the operation and correction of flexible one-time multiple-step plans may be closely related to conscious processing (Rolls, 1999; Rolls and Deco, 2002). In contrast, dorsal stream visual processing may be more closely related to executing an action on an object once the action has been selected, and the details of this action execution can take place implicitly (unconsciously) (Milner and Goodale, 1995; Rolls and Deco, 2002), perhaps because they do not require multiple step syntactic planning (Rolls, 1999).

The implication of this discussion is that in blindsight the dissociations between implicit and explicit processing may arise because different visual pathways, some involved in implicit and others in explicit processing, are differentially damaged. In contrast, in the experiments described here, the dissociation between implicit and explicit processing appears to arise from allowing differential processing (short information-poor versus longer) in the same visual processing stream and even area, the inferior temporal visual cortex. Thus the implications of the dissociations described here, and their underlying neuronal basis, may be particularly relevant in a different way to understanding visual information processing and consciousness.

One of the implications of blindsight thus seems to be that some visual pathways are more involved in implicit processing, and other pathways in explicit processing. In contrast, the results described here suggest that short and information-poor signals in a sensory system involved in conscious processing do not reach consciousness, and do not interrupt ongoing or engage conscious processing. This evidence described here thus provides interesting and direct evidence that there may be a threshold for activity in a sensory stream that must be exceeded in order to lead to consciousness, even when that activity is sufficient for some types of visual processing such as visual identification at well above chance in an implicit mode. The latter implicit mode processing can be revealed by forced choice tests and by direct measurements of neuronal responses. (Complementary evidence at the purely psychophysical level using backward masking has been obtained by Marcel (1983a,b) and discussed by Weiskrantz (1998, 2001).) Possible reasons for this relatively high threshold for consciousness are considered next.

It is suggested that the threshold for conscious visual perception may be set to be higher than the level at which small but significant sensory information is present so that the systems in the brain that implement the type of information processing involved in conscious thoughts are not interrupted by small signals that could be noise in sensory pathways, and that may not require use of the type of processing that can be performed by the conscious processing system. The exact nature of the information processing that is linked essentially to consciousness is the subject for great debate. My own theory is that phenomenal consciousness is the state that occurs when one is thinking about one's own syntactic (or more generally linguistic) thoughts (Rolls, 1999, 2000b, 1997). In that the theory is premised on thoughts about thoughts, it is a Higher Order Thought (HOT) theory of consciousness (see also Rosenthal, 1990, 1993), and in that the theory is about linguistic thoughts, the author has termed it a HOLT theory of consciousness (Rolls, 1999, 2000b). It may also be called a HOST theory, as the higher order thoughts are about syntactic thoughts. The computational argument that the author put to specify why the higher order thoughts are computationally useful if they are syntactic is as follows. If a multi-step plan involves a number of different symbols, and requires the relations (e.g., the conditional relations) between the symbols to be specified correctly in each step of the plan, then some form of syntax is needed; for otherwise, the symbols would not be bound together correctly in each step of the plan. If the plan produces an incorrect outcome, then a process that can reflect on and evaluate each step of the plan, and determine which may be the incorrect step, is a way to solve the credit assignment problem (Rolls, 1999). It is argued that it is implausible that such higher order syntactic thought processes would occur without it feeling like something to be the system that implements these processes, especially when the thoughts are grounded in the world (Rolls, 1999). This is only a plausibility argument. It should be noted that the higher order thoughts are linguistic in the sense that they require syntax, but not necessarily, of course, human verbal language. However, to the extent that some form of syntactic processing is closely related to consciousness, it is likely to be that a serial and time-consuming process is needed (with serial processing used to limit the binding problem to whatever syntax the brain can implement). Given the serial nature of this process, and its use for implementing long-term and/or multistep planning, it is suggested that it may be useful not to interrupt it unless the processing systems that can perform non-syntactic implicit processing to detect stimuli or stimulus change have sufficiently strong evidence that the signal is strong, and is of a type which may be behaviourally significant, such as a moving spot on the horizon, or an emotional expression change on a face.

Acknowledgements

This research was supported by Medical Research Council Programme Grant PG9826105. The author wishes to acknowledge the excellent contributions of many scientific colleagues to the work described here, including P. Azzopardi, D.G. Purcell, S. Panzeri, A.L. Stewart, A. Treves, and M.J. Tovée.

References

Azzopardi, P. and Cowey, A. (1997) Is blindsight like normal, near-threshold vision? Proc. Natl. Acad. Sci. U.S.A., 94: 14,190–14,194.

Battaglia, F. and Treves, A. (1998) Stable and rapid recurrent processing in realistic autoassociative memories. Neural Comput., 10: 431–450.

Booth, M.C.A. and Rolls, E.T. (1998) View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. Cereb. Cortex, 8: 510–523.

Bruce, V. and Green, P. (1989) Visual Perception: Physiology, Psychology and Ecology. Lawrence Erlbaum, London.

De Gelder, B., Vroomen, J., Pourtois, G. and Weiskrantz, L. (1999) Non-conscious recognition of affect in the absence of striate cortex. Neuroreport, 10: 3759–3763.

Dinse, H.R. and Kruger, K. (1994) The timing of processing along the visual pathway in the cat. Neuroreport, 5: 893–897.

Farah, M.J. (1990) Visual Agnosia. MIT Press, Cambridge, Mass.

Farah, M.J., Levinson, K.L. and Klein, K.L. (1995a) Face perception and within-category discrimination in prosopagnosia. Neuropsychologia, 33: 661–674.

Farah, M.J., Wilson, K.D., Drain, H.M. and Tanaka, J.R. (1995b) The inverted face inversion effect in prosopagnosia: evidence for mandatory, face-specific perceptual mechanisms. Vision Res., 35: 2089–2093.

Gawne, T.J. and Richmond, B.J. (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? J. Neurosci., 13: 2758–2771.

Heller, J., Hertz, J.A., Kjaer, T.W. and Richmond, B.J. (1995) Information flow and temporal coding in primate pattern vision. J. Comput. Neurosci., 2: 175–193.

Humphreys, G.W. and Bruce, V. (1989) Visual Cognition. Erlbaum, Hove.

Ishai, A., Ungerleider, L.G., Martin, A., Schouten, J.L. and Haxby, J.V. (1999) Distributed representation of objects in the human ventral visual pathway. Proc. Natl. Acad. Sci. U.S.A., 96: 9379–9384.

Kanwisher, N., McDermott, J. and Chun, M.M. (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. J. Neurosci., 17: 4302–4311.

Kovacs, G., Vogels, R. and Orban, G.A. (1995) Cortical correlates of pattern backward-masking. Proc. Natl. Acad. Sci. U.S.A., 92: 5587–5591.

Marcel, A.J. (1983a) Conscious and unconscious perception: experiments on visual masking and word recognition. Cog. Psychol., 15: 197–237.

Marcel, A.J. (1983b) Conscious and unconscious perception: an approach to the relations between phenomenal experience and perceptual processes. Cog. Psychol., 15: 238–300.

Milner, A.D. and Goodale, M.A. (1995) The Visual Brain in Action. Oxford University Press, Oxford.

Nowak, L.G. and Bullier, J. (1997) The timing of information transfer in the visual system. In: Kaas, J., Rockland, K. and Peters, A. (Eds.), Cerebral Cortex: Extrastriate Cortex in Primates. Plenum Press, New York, pp. 205–241.

Optican, L.M. and Richmond, B.J. (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: III. Information theoretic analysis. J. Neurophysiol., 57: 162–178.

Oram, M.W. and Perrett, D.I. (1992) Time course of neural responses discriminating different views of the face and head. J. Neurophysiol., 68: 70–84.

Panzeri, S. and Treves, A. (1996) Analytical estimates of limited sampling biases in different information measures. Network, 7: 87–107.

Panzeri, S., Rolls, E.T., Battaglia, F. and Lavis, R. (2001) Speed of information retrieval in multilayer networks of integrate-and-fire neurons. Network: Computation in Neural Systems, 12: 423–440.

Raiguel, S.E., Lagae, L., Gulyas, B. and Orban, G.A. (1989) Response latencies of visual cells in macaque areas V1, V2 and V5. Brain Res., 493: 155–159.

Rolls, E.T. (1992) Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. Philos. Trans. R. Soc. Lond., B. Biol. Sci., 335: 11–21.

Rolls, E.T. (1997) Consciousness in Neural Networks? Neural Netw., 10: 1227–1240.

Rolls, E.T. (1999) The Brain and Emotion. Oxford University Press, Oxford.

Rolls, E.T. (2000a) Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. Neuron, 27: 205–218.

Rolls, E.T. (2000b) Précis of The Brain and Emotion. Behav. Brain Sci., 23: 177–233.

Rolls, E.T. and Tovée, M.J. (1994) Processing speed in the cerebral cortex, and the neurophysiology of backward masking. Proc. R. Soc. Lond., B. Biol. Sci., 257: 9–15.

Rolls, E.T. and Tovée, M.J. (1995) The sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. J. Neurophysiol., 73: 713–726.

Rolls, E.T., Tovée, M.J., Purcell, D.G., Stewart, A.L. and Azzopardi, P. (1994) The responses of neurons in the temporal cortex of primates and face identification and detection. Experimental Brain Research, 101: 473–484.

Rolls, E.T., Treves, A., Tovée, M. and Panzeri, S. (1997) Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. J. Comput. Neurosci., 4: 309–333.

Rolls, E.T. and Treves, A. (1998) Neural Networks and Brain Function. Oxford University Press, Oxford.

Rolls, E.T., Tovée, M.J. and Panzeri, S. (1999) The neurophysiology of backward visual masking: information analysis. J. Cogn. Neurosci., 11: 335–346.

Rolls, E.T. and Deco, G. (2002) Computational Neuroscience of Vision. Oxford University Press, Oxford.

Rosenthal, D. (1990) A theory of consciousness. ZIF Report No. 40, Zentrum fur Interdisziplinaire Forschung, Bielefeld, Germany.

Rosenthal, D.M. (1993) Thinking that one thinks. In: Davies, M. and Humphreys, G.W. (Eds.), Consciousness. Blackwell, Oxford, Ch. 10, pp. 197–223.

Shannon, C.E. (1948) A mathematical theory of communication. AT&T Bell Laboratories Technical Journal, 27: 379–423.

Stoerig, P. and Cowey, A. (1997) Blindsight in man and monkey. Brain, 120: 535–559.

Thorpe, S., Fize, D. and Mariot, C. (1996) Speed of processing in the human visual system. Nature, 381: 520–522.

Tovée, M.J., Rolls, E.T., Treves, A. and Bellis, R.P. (1993) Information encoding and the responses of single neurons in the primate temporal visual cortex. J. Neurophysiol., 70: 640–654.

Tovée, M.J. and Rolls, E.T. (1995) Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. Visual Cognition, 2: 35–58.

Treves, A. (1993) Mean-field analysis of neuronal spike dynamics. Network, 4: 259–284.

Treves, A., Panzeri, S., Rolls, E.T., Booth, M. and Wakeman, E.A. (1999) Firing rate distributions and efficiency of information transmission of inferior temporal cortex neurons to natural visual stimuli. Neural Comput., 11: 611–641.

Vogels, R. and Orban, G.A. (1994) Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. J. Neurophysiol., 71: 1428–1451.

Wallis, G. and Rolls, E.T. (1997) Invariant face and object recognition in the visual system. Progress in Neurobiology, 51: 167–194.

Weiskrantz, L. (1997) Consciousness Lost and Found. A Neuropsychological Exploration. Oxford University Press, Oxford.

Weiskrantz, L. (1998) Blindsight. A Case Study and Implications, 2nd Edition, Oxford University Press, Oxford.

Weiskrantz, L. (2001) Blindsight – putting beta (β) on the back burner. In: De Gelder, B., De Haan, E. and Heywood, C. (Eds.), Out of Mind: Varieties of Unconscious Processes. Oxford University Press, Oxford, Ch. 2, pp. 20–31.

Weiskrantz, L., Warrington, E.K., Sanders, M.D. and Marshall, J. (1974) Visual capacity in the hemianopic field following a restricted occipital ablation. Brain, 97: 709–728.

Progress in Brain Research, Vol. 144
ISSN 0079-6123

CHAPTER 7

Rapid serial visual presentation for the determination of neural selectivity in area STSa

Peter Földiák*, Dengke Xiao, Christian Keysers, Robin Edwards
and David Ian Perrett

School of Psychology, University of St. Andrews, St Andrews, Fife KY16 9JU, UK

Abstract: We show that rapid serial visual presentation (RSVP) in combination with a progressive reduction of the stimulus set is an efficient method for describing the selectivity properties of high-level cortical neurons in single-cell electrophysiological recording experiments. Rapid presentation allows the experimental testing of a significantly larger number of stimuli, which can reduce the subjectivity of the results due to stimulus selection and the lack of sufficient control stimuli. We prove the reliability of the rapid presentation and stimulus reduction methods by repeated experiments and the comparison of different testing conditions. Our results from neurons in area STSa of the macaque temporal cortex provide a well-controlled confirmation for the existence of a population of cells that respond selectively to stimuli containing faces. View tuning properties measured using this method also confirmed earlier results. In addition, we found a population of cells that respond reliably to complex non-face stimuli, though their tuning properties are not obvious.

Introduction

The visual cortex encodes information about the visual world by the activity pattern of a large number of neurons in the visual cortex. In spite of the number of neurons involved in this representation, understanding the response properties of individual single neurons is still a fundamental scientific problem, as it is the stimulus selectivity of the neurons making up the population that determine the encoding and the nature of the representation. The encoding of stimuli by single neurons is also interesting as the properties directly represented by single neurons have substantial influence on what tasks can be performd by the visual system efficiently (Gardner-Medwin and Barlow, 2001).

Individual neurons throughout the vertebrate visual cortex are tuned to several parameters or properties of the stimulus, and they respond only when each of these relevant parameters lie within some relatively narrow range specific to the neuron. Finding effective stimuli for a particular sensory neuron involves, on the one hand, identifying the parameters of the stimulus that are relevant to the neuron under investigation and, on the other hand, finding the combinations of these stimulus parameters that are effective at activating the neuron. This is a non-trivial task even in the relatively well-characterized early visual areas, where the stimulus parameters relevant to the neurons are often assumed to be known. For instance, in primary visual cortex, stimulus location, size, orientation, spatial and temporal frequency, color and stereoscopic disparity are often considered the relevant parameters. The problem, however, becomes much harder in higher visual areas, for instance, in higher areas of the

---

*Corresponding author. E-mail: peter.foldiak@st-andrews.ac.uk

DOI: 10.1016/S0079-6123(03)14400-7