Processing speed in the cerebral cortex and the neurophysiology of visual masking

EDMUND T. ROLLS AND MARTIN J. TOVEE†

University of Oxford, Department of Experimental Psychology, South Parks Road, Oxford OX1 3UD, U.K.

SUMMARY

In experiments to investigate the duration of the time for which cortical neurons respond when the identification of a visual stimulus is just possible, we presented a test face stimulus for 16 ms, and followed it at different intervals by a masking stimulus (either an N–O pattern or a face) while recording from single neurons in the temporal visual cortex of macaques. When there was no mask the cells responded to the 16 ms of the test stimulus for 200–300 ms, far longer than the presentation time. We suggest that this reflects the operation of a short-term memory system implemented in cortical circuitry. If the mask was a stimulus which did not stimulate the cells (either a non-face pattern or a face which was a non-effective stimulus for that cell), then, as the interval between the onset of the test stimulus and the onset of the mask stimulus (the stimulus onset asynchrony) was reduced, the length of time for which the cells fired in response to the test stimulus was reduced. It is suggested that this is due to the mask stimulating adjacent cells in the cortex which by lateral inhibition reduce the responses of the cells activated by the test stimulus. When the stimulus onset asynchrony was 20 ms, face-selective neurons in the inferior temporal cortex of macaques responded for a period of 20–30 ms before their firing was interrupted by the mask. With the same test-mask stimulus onset asynchrony of 20 ms, humans could just identify which of six faces was shown. These results provide evidence that a cortical area can perform its computation necessary for the recognition of a visual stimulus in 20–30 ms, and provide a fundamental constraint which must be accounted for in any theory of cortical computation.

1. INTRODUCTION

The amount of time needed by each cortical area of the brain for the computations it performs is of fundamental importance for understanding cortical function. In one approach to this, we recorded the responses of single neurons in the macaque inferior temporal visual cortex to a set of visual stimuli. The set of stimuli consisted of faces, as there are neurons in this visual cortical area for which faces are effective stimuli (Rolls 1992, 1994). Each neuron responded more to some of the stimuli than to others, and thus the response of the neuron reflected which stimulus had been shown. We quantified the amount of information a neuron reflected about each stimulus by using information theory measures (Tovee *et al.* 1993). We showed that, although more information is available if long epochs of the neuronal response (e.g. 400 ms) are taken, there is considerable information available in a 20–40 ms period of the spike train. In fact, if a period as short as 20 ms was taken near the start of the neuronal response to the stimulus, then the information available was 64.9% of that available in a 400 ms epoch. Although this experiment showed that a large proportion of the total information was available in a short period of the spike train, it did not address directly the issue of whether a period of neuronal firing of as little as 20–40 ms in a cortical area might be sufficient to allow correct perceptual identification of the stimulus.

One potential way to address the issue of how long a cortical area must be active to mediate object recognition is to use a visual backward masking paradigm. In this paradigm there is a brief presentation of a test stimulus which is rapidly followed (within 1–100 ms) by the presentation of a second stimulus (the mask), which impairs or masks the perception of the test stimulus. What these experiments leave unanswered is for how long visual neurons actually fired under the masking condition at which the subject could just identify an object. Although there has been a great deal of psychophysical investigation with the visual masking paradigm (Turvey 1973; Breitmeyer 1980; Humphreys & Bruce 1989), there is very little direct evidence on the effects of visual masking on neuronal activity. For example, it is possible that if a neuron is well tuned to one class of stimulus, such as faces, that a pattern mask which does not activate the neuron will leave the cell firing for some time after the onset of the pattern mask. To obtain direct neurophysiological evidence on the effects of backward masking of neuronal activity, to clarify both what happens with visual backward masking and to show how long neurons may respond in a cortical area when perception and identification are just possible, we performed the experiments described here.

The aim of the experiments was to investigate the

† Present address: Department of Psychology, University of Newcastle Upon Tyne, Ridley Building, Newcastle upon Tyne NE1 7RU, U.K.

Proc. R. Soc. Lond. B (1994) **257**, 9–15
Printed in Great Britain

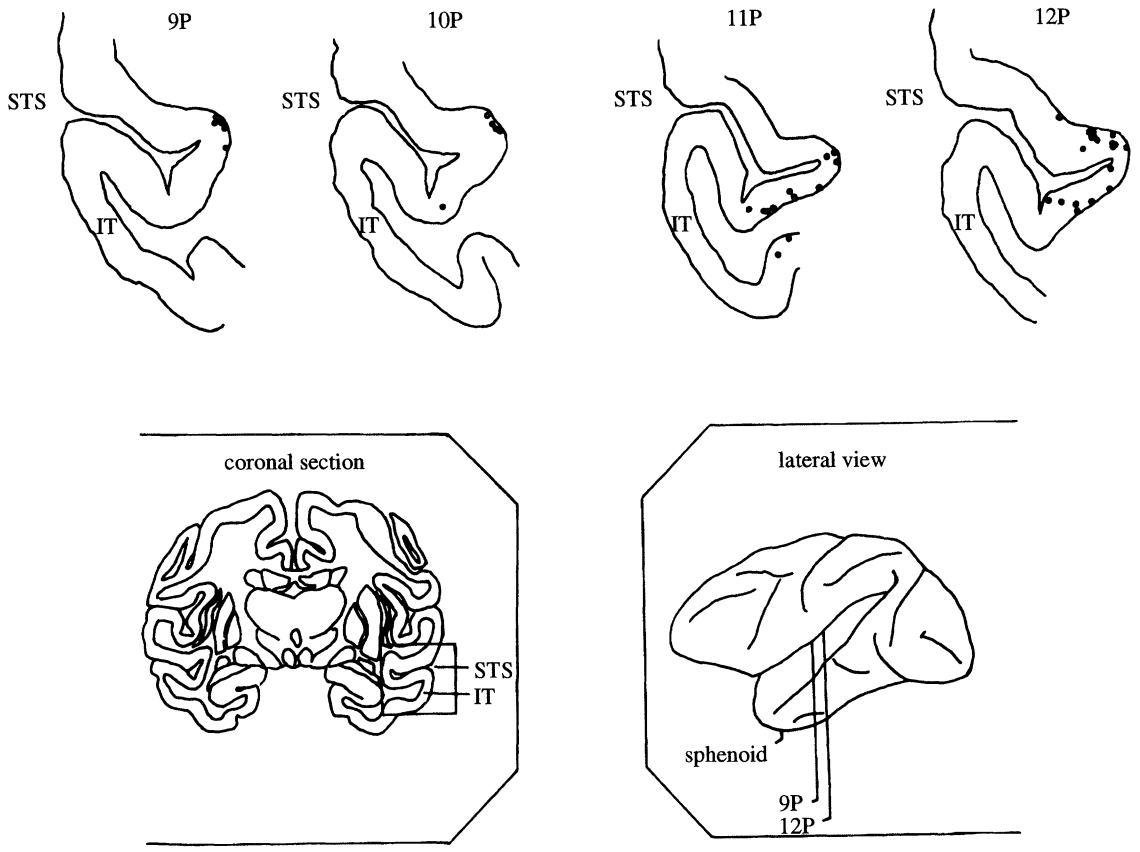

Figure 1. The recording sites shown on coronal sections of the neurons included in this study. The positions of the coronal sections are shown on a lateral view of the macaque brain. The distances refer to millimetres posterior to the sphenoid reference plane (see text). IT, inferior temporal cortex; STS, superior temporal sulcus.

effects of the backward masking paradigm on the responses of inferior temporal cortex neurons with face-selective responses. This population of neurons was studied because it is a population which can be identified reliably when recordings are made in the temporal cortical visual areas, and because the responses of these cells in these higher order cortical areas appear to be closely related to visual perception (Rolls 1992). The masks used were either stimuli to which the face-selective cell did not respond (a pattern mask made of overlapping letters N and O (see Rolls *et al.* 1994) or a non-effective face stimulus for that cell), or a face stimulus which did activate the cell. This investigation is part of a series directed towards understanding the normal functioning of the cerebral cortex, and how disorders of its function lead to perceptual and cognitive dysfunctions (Rolls 1984, 1990, 1992; Tovee & Rolls 1992*a*; Tovee *et al.* 1993).

2. METHODS

The activity of single neurons was recorded with glass-insulated tungsten micro-electrodes in the anterior part of the superior temporal sulcus (STS) (see figure 1) in two alert macaque monkeys (*Macaca mulatta*, mass 3.0 kg) seated in a primate chair, using techniques that have been described elsewhere (Tovee *et al.* 1993). The preparative procedures were performed aseptically under sodium thiopentone anaesthesia (by using pre-treatment with ketamine and post-treatment with the analgesic buprenorphine (Temgesic) and the antibiotic amoxycillin (Cynulox)), and all procedures

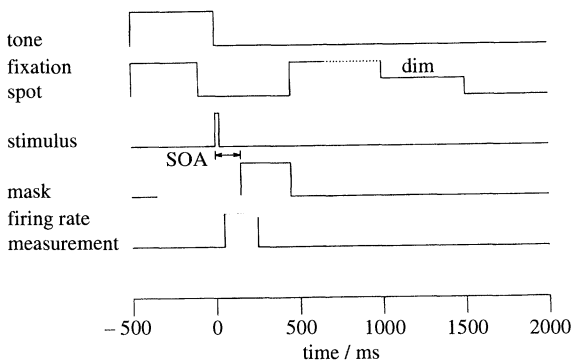

Figure 2. The timing used in the backward masking visual fixation blink task.

were in accordance with the Policy Regarding the Care and Use of Animals approved by the Society for Neuroscience, and were licensed under the U.K. Animals (Scientific Procedures) Act 1986. Eye position was measured to an accuracy of 0.5° with the search coil technique. A visual fixation task ensured that the monkey looked steadily at the screen throughout the presentation of each stimulus. The task was a blink version of a visual fixation task in which the fixation spot was blinked off 100 ms before the target (otherwise called test) stimulus appeared. The stimuli were static visual stimuli subtending 8° in the visual field presented on a video monitor at a distance of 1.0 m. The fixation spot position was at the centre of the screen. The monitor was viewed binocularly, with the whole screen visible to both eyes.

The timing of the task is shown in figure 2. Each trial started at −500 ms (with respect to the onset of the test

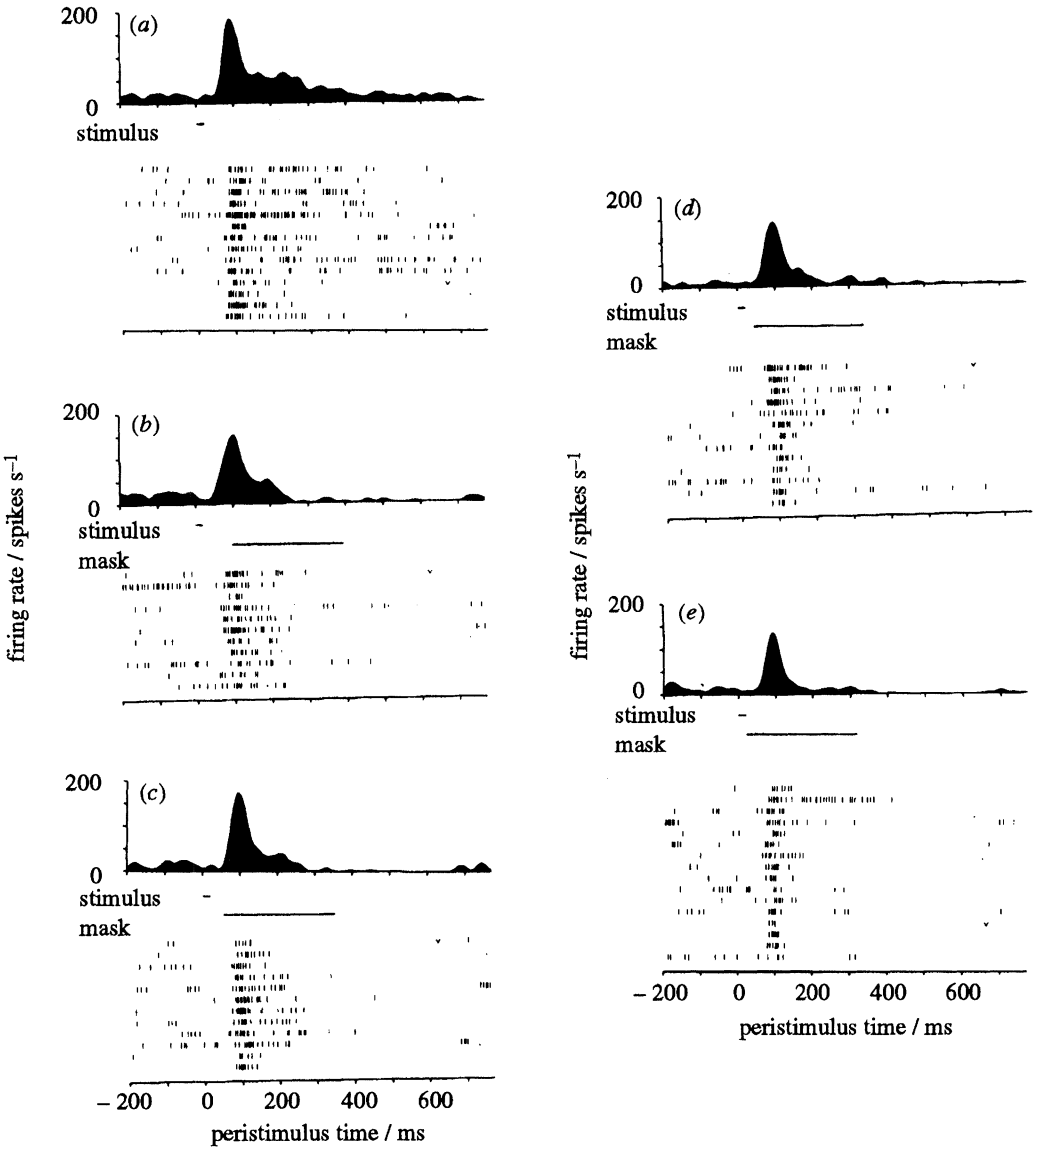

Figure 3. Peristimulus rastergrams and smoothed peristimulus spike density histograms based on responses in 8–16 trials to the test face alone (top raster, histogram pair), and to the test face followed by a masking stimulus (which was a face that was ineffective in activating the cell) with different SOA values. (SOA = stimulus onset asynchrony, the delay between the onset of the test stimulus and the onset of the mask stimulus.) 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 SOA, the mask stimulus was delayed until well after the end of the recording period shown.) SOA values are (*a*) 1000 ms, (*b*) 100 ms, (*c*) 60 ms, (*d*) 40 ms, (*e*) 20 ms.

image) with a 500 ms warning tone to allow fixation of the fixation point, which appeared at the same time. At −100 ms the fixation spot was blinked off so 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. At 0 ms, the tone was switched off and the test image was switched on for 16 ms. This 16 ms period was the frame duration of the video framestore with which the images were presented. The image was drawn on the monitor from the top to the bottom in the first 16 ms of the frame period by the framestore, with the remaining 4 ms of the frame period being the vertical blank interval. (The PAL video system was in use.) 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. The time between the start of the test stimulus and the start of the mask stimulus (the stimulus onset asynchrony (SOA)) was either 20, 40, 60, 100 or 1000 ms (chosen in a pseudo-random sequence by the computer). The 1000 ms condition was used to measure the response to the test stimulus alone (which was possible because the mask was delayed for so long). The duration of the masking stimulus was 300 ms. At the termination of the masking stimulus the fixation spot reappeared, and then after a random interval in the range 150–3350 ms it dimmed, to indicate that licking responses to a tube in front of the mouth would result in the delivery of a reward. The dimming period was 1000 ms, and after this the fixation spot was switched off, and reward availability terminated 500 ms later. The monkey was required to fixate the fixation spot, and if he licked at any time other than when the spot was dimmed, saline instead of fruit juice was delivered from the tube. If the eyes moved by more than 0.5° from time 0 until the start of the dimming period, the trial was aborted and the data for the trial were rejected. When a trial aborted, a high-frequency tone sounded for 0.5 s, no

reinforcement was available for that trial, and the inter-trial interval was lengthened from 8 s to 11 s.

The criterion for the face-selective neurons analysed in this study was that the response to the optimal face stimulus should be more than twice as large as to the optimal non-face stimulus, and that this difference should be significant (Rolls 1984, 1992; Baylis *et al.* 1985). If the neuron satisfied the criterion then it was tested with one of the effective face stimuli for that neuron. The default periods for which the firing rate was calculated were a 150 ms and a 300 ms period starting 50 ms after the onset of the target stimulus. (The period was chosen to start at 50 ms because none of the neurons started to respond with a latency shorter than this, and all started to respond strongly from some time in the post-stimulus period of 50–100 ms.) The mean firing rate over the 16–24 trials, together with its standard error, was calculated by the computer for each of the stimulus conditions for graphical presentation. In addition, an analysis of variance was done on the data which showed whether there was a significant effect of SOA.

3. RESULTS

Examples of the effects of backward masking on the responses of a single neuron are shown in peristimulus rastergram and time histogram form in figure 3. The rastergram–spike density histogram pair in figure 3*a* show the responses of the neuron to a single frame of the test stimulus (an effective face stimulus for that neuron). Relative to the pre-stimulus rate, there was an increase in the firing produced with a latency of approximately 75 ms, and this firing lasted for 200–300 ms, i.e. for much longer than the 16 ms presentation of the target stimulus. In figure 3*b*, *c* the effects of introducing a non-effective face as the masking stimulus with different SOAs 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, and in all cases the temporal aspects of the masking were similar to those shown in figure 3.

The data from several cells under conditions of backward masking are shown in figure 4. The mean firing rate (in a 150 ms period starting 50 ms after stimulus onset) is shown, together with the s.e.m. (calculated over 16–24 trials). It is clear that the effect of the mask was significant, and depended on the SOA, and this was confirmed by the fact that there was a highly significant effect ($p < 0.001$) of SOA in the ANOVA for each of the neurons shown in figure 4. For the three cells included in the data shown in figure 4, one was tested with the standard pattern mask, and two with a non-effective face used as the mask. The data had a similar form, although the firing rates were a little lower, if the firing rates were measured over a 300 ms period starting at 50 ms post-stimulus (rather than the 150 ms period for which data are shown in figure 4).

Similar effects to those illustrated for single cells in figures 3 and 4 were evident for the whole population of cells analysed completely, as shown in figure 5. In particular, the mask became monotonically more effective as the SOA was decreased to 20 ms.

So far, the effects considered have been interruption of the processing of an effective stimulus for the cell produced by a mask which did not itself activate the cell. The neurophysiological data provide direct evidence that one way in which backward masking

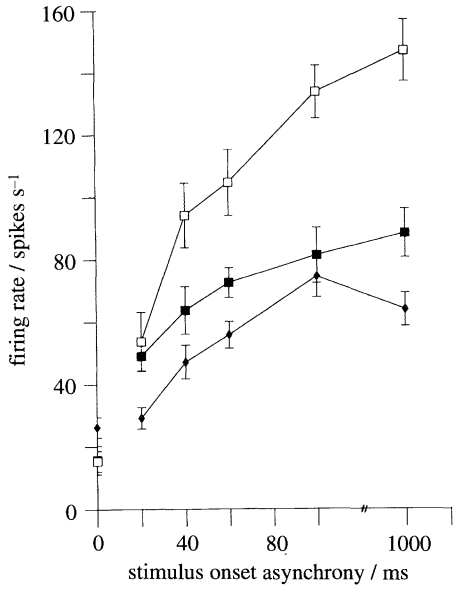

Figure 4. The neuronal responses of three different cells to a test face stimulus which was an effective stimulus for the cell, as a function of the SOA value when the standard mask was used. The mean and the standard error of the mean firing rate (taken over a 150 ms period starting 50 ms after the start of the test stimulus) are shown for each condition. AM 168 (filled squares), AM 166 (open squares), AM 126 (filled diamonds).

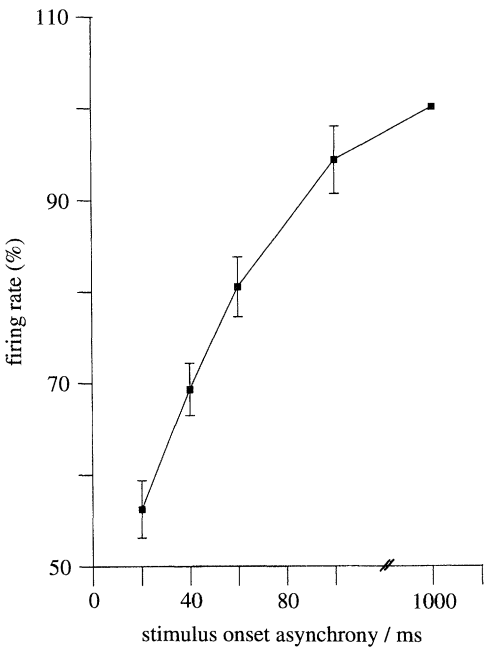

Figure 5. The average response of the population of 20 cells fully analysed as a function of SOA. The maximal response of each neuron was normalized to 100% without the mask stimulus before the mean percentage response was calculated. The mean ± the standard error of the mean is shown. The spontaneous rate of firing of the neurons was on average 16% of the maximal firing rate.

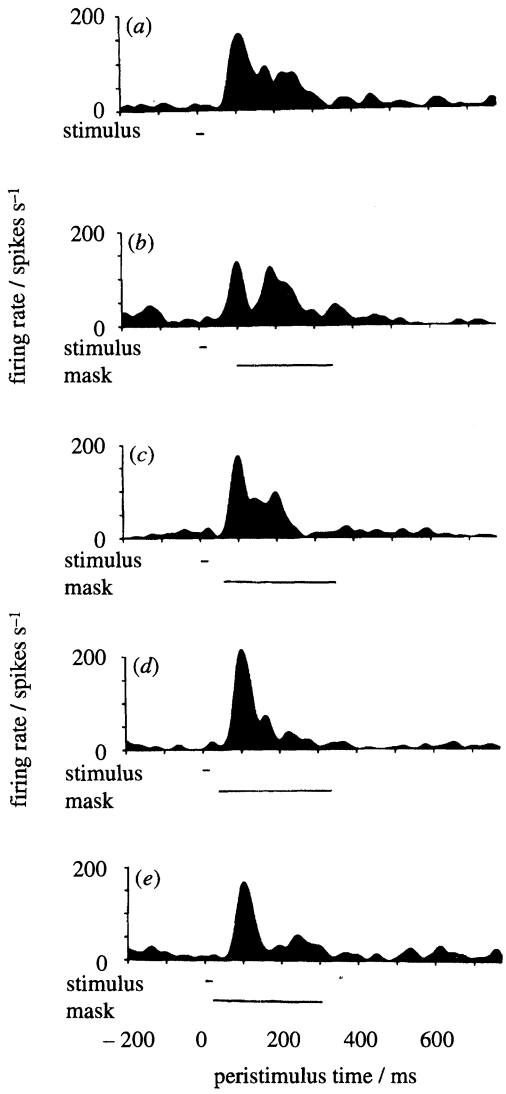

Figure 6. The neuronal responses of a cell to a target face stimulus which was an effective stimulus for the cell as a function of the SOA when the mask was a face which also activated the neuron from which the recording was being made. Peristimulus spike density histograms of the firing are shown for each stimulus condition. (The top trace shows the response to the target stimulus alone, in that, with this 1000 ms SOA, the mask stimulus was delayed until well after the end of the recording period shown.) SOA values are (*a*) 1000 ms, (*b*) 100 ms, (*c*) 60 ms, (*d*) 40 ms, (*e*) 20 ms.

operates is by interruption, and that this interruption occurs even when the mask is a stimulus which does not stimulate a particular face-selective cell, such as an N–O pattern or a non-effective face. What happens if the mask is a stimulus which also happens to activate the cell? This is illustrated in figure 6, in which an effective stimulus for the cell was presented as a mask at different time intervals after an effective target. The response to the masking stimulus, a face which activated the neuron, was similar to the response of the neuron to the target stimulus, which is shown in figure 6*a*. The response to the target stimulus was shortened and reduced when it was followed by the effective mask stimulus at short SOAs (as can be seen at, for example, 20 ms and 40 ms in figure 6). Indeed, the neuron's response at these SOAs to an effective test face followed by an effective face stimulus used as a mask was smaller than to the test stimulus presented alone (compare figure 6*a*, when only the target stimulus was shown, with the traces for SOAs of 20 ms and 40 ms). At SOAs of 60 ms and 100 ms, there appeared to be two peaks in the neuronal response firing (see figure 6). This would be consistent with feedforward inhibition (produced by the mask) of the cell's response to the target stimulus, followed by a response to the mask stimulus. Thus, at SOAs of 60–100 ms, in the visual system, at least at the cellular level, there appear to be temporally separate responses to two successive effective stimuli for a cell. Similar effects were found in all five neurons tested in this way with an effective stimulus for the cell used in a backward masking paradigm.

4. DISCUSSION

The direct neurophysiological measurement of how neurons respond in visual backward masking thus shows that neurons fire for a very short time with short SOA stimulus conditions. In particular, with SOAs of 20 ms, it was found that the face-selective neurons responded for only 20–30 ms. This value can be seen most clearly from the rasters of the spikes that occurred on individual trials illustrated in figure 3. Moreover, it is clear that only a small number of spikes was elicited on each trial to an optimal stimulus for a cell when that stimulus was followed at an SOA of 20 ms by a mask.

The neurophysiological data described here can be compared directly with the effects of backward masking in human observers, studied for example in the same apparatus with the same stimuli. For the human observers, identification of which face from a set of six had been seen was 50 % correct with an SOA of 20 ms, and 97 % correct with an SOA of 40 ms (corrected for guessing) (Rolls *et al.* 1994). Comparing the human performance and the macaque neuronal responses under the same stimulus conditions leads to the conclusion that when it is just possible to identify which face has been seen, neurons in a given cortical area may be responding for only 20–30 ms. The implication is that 20–30 ms is sufficient time for a visual cortical area to perform sufficient computation to enable its output to be used for identification. This rapid processing places constraints on the type of processing that could be performed in a cortical area when it is performing sufficient computation for perceptual identification. Of course, perception is better if more time is allowed for cortical processing, and indeed, with an SOA of 20 ms, the subjective perceptual impression of which face has been seen is limited. However, the point is that sufficient processing has been performed to allow identification. We note that the cells described here parallel human perception in many ways, one of which is that these cells respond differently to different faces, and reflect in their responses information about the identity of a face seen (Tovee *et al.* 1993; Rolls 1994). In that these neurons reflect information about identity, their responses are likely to be relevant to a face identification task.

These neurophysiological findings provide direct evidence that the masking stimulus stops abruptly the

responses of the neurons that were activated by the target stimulus. From psychophysical investigations of masking (see, for example, Humphreys & Bruce 1989), the actual processing time available for the neurons activated by the test stimulus was not clearly identified. For example, it was possible that a neuron relatively selectively activated by a test stimulus might remain active for a period as long as 60–100 ms when a pattern mask which did not activate that neuron followed the effective stimulus with a short (20 ms) SOA. Even in experiments when the mask was a stimulus which had to be identified by a human observer, it was possible that this activated different neurons, and that therefore the set of neurons activated by the target might still continue responding. The present results show that this is not the case. The results also show that, at least for face neurons, a non-face pattern mask was very effective as a mask for neurons with face-selective responses. The masking efficacy was of the same order as another face. This implies lateral inhibition of neurons responding to faces by other neurons which are activated by non-face stimuli. The locus of this interaction could be in the temporal cortical areas themselves, or it could be at earlier stages of cortical visual processing. Thus, at the single cell level, interruption of the processing of the test stimulus can be produced by not only similar but also by dissimilar masks (cf. Humphreys & Bruce 1989).

The neurophysiological results also show that there is a population of neurons in the temporal cortical areas which respond to a single-frame (16 ms) presentation of the target stimulus, a face. Typically the neuronal response to a 16 ms stimulus lasts for longer than 16 ms, with some (although reduced, see figures 2 and 6) neuronal firing often being evident for 200–300 ms in response to the 16 ms stimulus. This effect suggests that there may be a short-term visual memory implemented by the continued firing of these neurons after a stimulus has disappeared. This memory appears to be implemented by cortical mechanisms, for such continuing firing after a brief (stroboscopic) illumination of an effective stimulus does not cause continuing firing in single neurons in the lateral geniculate nucleus, although it does in the primary visual cortex (K. A. C. Martin, personal communication). We suggest on the basis of these findings that the short-term visual memory is implemented by the connections made between nearby pyramidal cells in the cerebral cortex. These recurrent connections may function in part as an autoassociative memory, which not only enables cortical networks to show continued firing for a few hundred milliseconds after a briefly presented stimulus, but also confer some of the response specificity inherent in the responses of cortical neurons. (The other mechanism that would make a major contribution to the tuning of a neuron would be the feedforward synaptic strengths from the preceding stage of processing.) One effect which may be facilitated by this short-term visual memory is implementation of a trace learning rule in the visual cortex which may be used to learn invariant representations (Foldiak 1991; Rolls 1992, 1994; Wallis *et al.* 1993).

One factor which may contribute to the implementation of a trace learning rule in the cerebral cortex is that a memory trace of a stimulus remains for a short period after a stimulus has been shown in the form of continued firing. This means that the neurons activated by a stimulus are still in a state of activation (and postsynaptic depolarization) perhaps 500 ms later when the same stimulus is seen translated across the retina, viewed from a different angle, or at a different size, so that the active axons carrying the transformed representation can undergo hebbian associative synaptic modification onto just those neurons that remain in an activated state from the previous input produced by the same object. In this way, the invariant properties evident across short time epochs (of, for example, 0.5 s) of the inputs produced by objects may be learned by the visual system (see Wallis *et al.* 1993; Rolls 1994). Another mechanism that may contribute to the trace memory aspect of the associative synaptic modification rule that is needed for this type of invariance learning is prolonged (several hundred milliseconds) activation of the postsynaptic neuronal element produced as a result of NMDA receptor activation (Foldiak 1991; Rolls 1992).

The results of the masking experiments are consistent with previous work which suggests that very little time is required at each processing area for object recognition (Thorpe & Imbert 1989; Oram & Perrett 1992; Tovee *et al.* 1993). Tovee *et al.* (1993) used the techniques of information theory to analyse the responses of visual neurons in the IT and STS of awake, behaving macaques. Up to 64.9% of the information available in a 400 ms period is available in a 20 ms sample near the start of the spike train, and up to 87% is available in a 50 ms sample. The response latencies in different cortical areas also suggest a processing duration of 10–20 ms at each area (Thorpe & Imbert 1989; Oram & Perrett 1992). However, the results reported here are the first direct measurement of the time neurons need to be active to mediate object recognition, and the results described by Tovee *et al.* (1993) are the first to measure explicitly the information available in a range of short temporal epochs of neuronal spike trains. (Other groups that have measured the information available in neuronal spike trains have used much longer, fixed, epochs (see, for example Optican & Richmond 1987; Optican *et al.* 1991). Taken as a whole these results suggest very rapid processing in the visual cortex with individual neurons having to be active for only 20–30 ms.

This view of cortical processing provides constraints on how the cortex performs its processing. One view, which we believe is biologically plausible, is that in the 15–20 ms period there is sufficient time for feedforward activation from the preceding stage to activate the neurons, and for a sufficient proportion of the neurons to exchange information through the recurrent lateral excitatory connections for the system to settle into a stage determined by the feedforward activation and by the lateral recurrent connections within a cortical area. A recent analysis of the dynamics of recurrent processing suggests that, if the time constants of the cell membranes are incorporated into the model, the time

taken for the system to settle into an attractor state is in the order of the time constant of the synaptic effects on the cells, i.e. in the order of 10–20 ms (Treves 1993; Treves *et al.* 1994).

Gray & Singer (1989) have reported oscillations in the neuronal responses of cells in the cat visual cortex. It has been suggested that these oscillations may form the neural substrate of a temporal binding mechanism (Gray & Singer 1989; Singer 1993). It was proposed that populations of neurons that are processing features belonging to the same object will have oscillations in their neuronal responses at the same frequency. The existence of oscillations in the primate visual cortex is a matter of some controversy (Tovee & Rolls 1992*a*, *b*; Young *et al.* 1992; Kreiter & Singer 1992). However, if one considers the oscillations which have been reported in the cat, these are not time locked to the onset of the stimulus, occur at variable time intervals after the onset of the stimulus, and occur in spindles which develop over tens of milliseconds and persist for 100–300 ms (Eckhorn *et al.* 1988; Gray & Singer 1989; Engel *et al.* 1990). In primates, such oscillations are unlikely to play a part in the rapid visual processing we have demonstrated here which is sufficient for the recognition of static complex patterns and objects.

This research was supported by Medical Research Council Grant PG8513790.

REFERENCES

Baylis, G. C., Rolls, E. T. & Leonard, C. M. 1985 Selectivity between faces in the responses of a population of neurons in the cortex in the superior temporal sulcus of the monkey. *Brain Res.* **342**, 91–102.

Breitmeyer, B. G. 1980 Unmasking visual masking: a look at the 'why' behind the veil of the 'how'. *Psychol. Rev.* **87**, 52–69.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M. & Reitbaeck, H. J. 1988 Coherent oscillations: A mechanism for feature linking in the visual cortex? *Biol. Cyber.* **60**, 121–130.

Engel, A. K., Konig, P., Gray, C. & Singer, W. 1990 Stimulus-dependent neural oscillations in cat visual cortex: Inter-columnar interaction as determined by cross-correlation analysis. *Eur. J. Neurosci.* **2**, 588–606.

Foldiak, P. 1991 Learning invariance from transformation sequences. *Neural Comp.* **3**, 194–200.

Gray, C. M. & Singer, W. 1989 Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. natn. Acad. Sci. U.S.A.* **69**, 1698–1702.

Humphreys, G. W. & Bruce, V. 1989 *Visual cognition.* Hove: Erlbaum.

Kreiter, A. K. & Singer, W. 1992 Oscillatory neuronal responses in the visual cortex of the awake macaque monkey. *Eur. J. Neurosci.* **4**, 369–375.

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

Optican, L. & 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**, 132–146.

Optican, L. M., Gawne, T. J., Richmond, B. J. & Joseph, P. J. 1991 Unbiased measures of transmitted information and channel capacity from multivariate neuronal data. *Biol. Cyber.* **65**, 305–310.

Rolls, E. T. 1984 Neurons in the cortex of the temporal lobe and in the amygdala of the monkey with responses selective for faces. *Hum. Neurobiol.* **3**, 209–222.

Rolls, E. T. 1990 A theory of emotion, and its application to understanding the neural basis of emotion. *Cogn. Emot.* **4**, 161–190.

Rolls, E. T. 1992 Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Phil. Trans. R. Soc. Lond.* B **335**, 11–21.

Rolls, E. T. 1994 Brain mechanisms for invariant visual recognition and learning. *Behav. Process.* (In the press.)

Rolls, E. T., Tovee, M. J., Purcell, D. G., Stewart, A L. & Azzopardi, P. 1994 The responses of neurons in the temporal cortex of primates, and face identification and detection. *Exp. Brain Res.* (In the press.)

Singer, W. 1993 Synchronization of cortical activity and its putative role in information processing and learning. *A. Rev. Physiol.* **55**, 349–374.

Thorpe, S. J. & Imbert, M. 1989 Biological constraints on connectionist models. In *Connectionism in perspective* (ed. R. Pfeifer, Z. Schreter & F. Fogelman-Soulie), pp. 63–92. Amsterdam: Elsevier.

Tovee, M. J. & Rolls, E. T. 1992*a* Oscillatory activity is not evident in the primate temporal visual cortex with static visual stimuli. *Neuroreport* **3**, 369–372.

Tovee, M. J. & Rolls, E. T. 1992*b* The functional nature of neuronal oscillations. *Trends Neurosci.* **15**, 387.

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

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

Treves, A., Rolls, E. T. & Tovee, M. J. 1994 On the time required for recurrent processing in the brain. *Proc. natn. Acad. Sci. U.S.A.* (In the press.)

Turvey, M. T. 1973 On the peripheral and central processes in vision: inferences from an information processing analysis of masking with patterned stimuli. *Psychol. Rev.* **80**, 1–52.

Wallis, G., Rolls, E. T. & Foldiak, P. 1993 Learning invariant responses to the natural transformations of objects. *International Joint Conference on Neural Networks* **2**, 1087–1090.

Young, M. P., Tanaka, K. & Yamane, S. 1992 On oscillating neuronal responses in the visual cortex of the monkey. *J. Neurophysiol.* **6**, 1464–1474.

Received 7 March 1994; accepted 8 April 1994