

Neuronal firing rate, inter-neuron correlation and synchrony in area MT are correlated with directional choices during stimulus and reward expectation

A. Thiele · K.-P. Hoffmann

Received: 10 December 2007 / Accepted: 13 April 2008 / Published online: 29 April 2008
© Springer-Verlag 2008

Abstract Sensation, memories, and predictions contribute to choices in everyday life, and their relative impact should change with task constraints. To investigate how the impact from sensory cortex on decision making varies with task constraints we trained macaque monkeys in a direction discrimination task where they could maximize reward by waiting for sensory visual information early in a trial, while focusing on memory and reward prediction as a trial progressed. The task constraints caused animals to indicate decisions in complete absence of visual motion stimuli (stimulus independent decisions), as 25% of the trials were ‘no stimulus’ trials. On ‘no stimulus’ trials reward delivery depended on the current decision in relation to the decision history. Stimulus independent decisions occurred during an epoch when a stimulus could in principle have been presented, or afterwards when stimuli could not occur anymore. Stimulus independent decisions were significantly different during these two periods. Reward exploitation was more efficient late in the trial, but it was not associated with systematic activity changes in directionally selective neurons in area MT. Conversely, systematic changes of neuronal activity and firing rate correlation in directionally selective middle temporal area (MT) neurons were restricted to a short time period before

early decisions. Changing task constraints in the course of a single trial thus determines how neurons in sensory areas contribute to decision making.

Keywords Choice probability · Direction discrimination · Middle temporal · Motion processing

Introduction

The brain uses various sources of information to form decisions and guide behaviour appropriately. In order to do so it needs to gate the flow of information in relation to the subjects’ current goals. Under normal conditions sources of information that affect information processing and decision making are sensory cues (Britten et al. 1992), memory traces (Miller et al. 1993, 1996; Rainer et al. 1998), abstract response rules (Konishi et al. 1998; Wallis et al. 2001; Nakahara et al. 2002; Wallis and Miller 2003; Mansouri et al. 2006), attention (Desimone and Duncan 1995; Treue and Maunsell 1996; Roelfsema et al. 1998; Roberts et al. 2007) or predictions (Sugrue et al. 2004; Yang and Shadlen 2007), and neuronal signatures of these have been described in a variety of different areas and task constraints. An area that has been extensively studied in terms of its contribution to sensory based decision making is area MT of the macaque monkey (Newsome et al. 1990; Salzman et al. 1990; Britten et al. 1992, 1996; Thiele and Hoffmann 1996; Bradley et al. 1998; Thiele et al. 1999, 2001; Dodd et al. 2001). Neuronal activity in MT correlates with the choice when animals expect visual motion input, even when the input is not informative (Britten et al. 1996; Bradley et al. 1998), or entirely absent (Thiele and Hoffmann 1996). In the absence of informative cues small activity changes in

Electronic supplementary material The online version of this article (doi:10.1007/s00221-008-1391-z) contains supplementary material, which is available to authorized users.

A. Thiele · K.-P. Hoffmann
Allgemeine Zoologie & Neurobiologie,
Ruhr-Universität Bochum, Bochum, Germany

A. Thiele (✉)
Institute of Neuroscience, Henry Wellcome Building,
Newcastle University, Newcastle upon Tyne NE2 4HH, UK
e-mail: alex.thiele@ncl.ac.uk

area MT might be interpreted by downstream neurons as if they were associated with a specific direction of motion, thereby influencing decision making when monkey engage in a direction discrimination task. Electrical microstimulation experiments have shown that the influence of MT activity on downstream areas is most effective during specific epochs of a task; namely during epochs where decision relevant information was present in the external world, while it was less effective during externally cued task epochs where such information was absent (Seidemann et al. 1998). This was interpreted as evidence for selective gating of sensory signals in the cortex. It remains to be determined how the flow of activity is gated under more natural circumstances where changes of task constraints are gradual, not signaled by external cues (e.g. stimulus on and offset), and where activity in MT was not subject to artificial manipulation, but generated internally.

We were additionally interested in the mechanisms by which decisions are formed. It is well established that systematic changes of firing rates in MT contribute to decision making in the absence of informative cues (for review see Krug 2004). A complimentary mechanisms which could benefit decision making is synchrony between neurons that are tuned for similar characteristics, as integration of excitatory post synaptic potentials (EPSPs) at downstream stages may benefit from synchronized inputs (Fries et al. 2001). Although a previous report failed to find a systematic relation between neuronal synchrony in MT and the animal's choice (Bair et al. 2001), a significant correlation between the speed of behavioural reports and gamma band synchronization in V4 during attentional selection has been demonstrated (Womelsdorf et al. 2006). If correlated activity contributes to decision making by increasing the impact of inputs in downstream neurons, it might be restricted to time periods where such inputs are relevant for the animal (i.e. when they wait for a visual stimulus), and might be reduced when directional choices cannot benefit from sensory information.

To investigate how firing rate and neuronal synchrony in area MT contribute to directional choice under changing task constraints, we trained monkeys to participate in discriminating the direction of moving stimuli or indicate decisions once the probability of stimulus occurrence had returned to zero for that particular trial. For the latter the probability of rewards depended on the animal's choice history. They thus had to change strategy and use memory traces if they wanted to maximize reward income. We found that firing rate, firing rate correlation and synchrony between simultaneously recorded neurons significantly correlated with the choice direction while animals waited for sensory stimuli. When an animal's decision was more based on an attempt to maximize the reward income, MT activity, firing rate correlation and synchrony were not

systematically related to the choice. Thus, the presence or absence of choice related activity pattern in area MT varies with task constraints within the time course of single trials.

Methods

All experiments were carried out in accordance with the European Communities Council Directive 1986 (86 909 EEC) and National Institutes of Health guidelines for care and use of animals for experimental procedures. Details regarding training, surgery, and post-operative care are given in (Thiele et al. 1999). Two male Macaca Mulatta participated in the study. Data from one of the two monkeys (AR) have been published previously in a different context, where error trials in relation to stimulus induced neuronal activity were investigated (Thiele et al. 1999). No aspects of any of the data from the second monkey (CS) have been published previously.

Stimuli and behavioural paradigm and controls

The monkeys were trained in a direction discrimination task (Fig. 1a). During the experiments, each animal was seated comfortably in a primate chair with its head restrained. We monitored eye movements using scleral search coils (optimum resolution of 1.5'), sampled at 500 Hz. Monkeys started a trial by clutching a central touch bar in front of its chest, upon which a fixation point (0.2° in diameter) was back projected by a light emitting diode onto a translucent tangent screen. The screen subtended 90° × 90° of visual angle at a viewing distance of 38 cm. The fixation point was always presented in the centre of the projection screen. The maximum fixation window was ±1°. Monkeys were required to fixate within 500 ms after the appearance of the fixation point.

Stimuli

During each experiment, the time of stimulus onset, contrast, and direction were varied randomly. The direction of motion was along one of the four cardinal directions. Stimuli were presented in the receptive field of the recorded neuron. When two or more single units were recorded simultaneously, the stimulus covered all of the units' receptive fields. Stimuli consisted of square wave gratings (0.3–0.5 cyc/deg) moving uni-directionally within a quadratic aperture. During the whole experiment, a stationary gaussian-filtered white noise stimulus was also back projected onto the screen (by a slide projector). This added stationary noise made the task more difficult, causing a higher percentage of error trials (Hoffmann and von Seelen 1980). The noisy stationary background was generated by

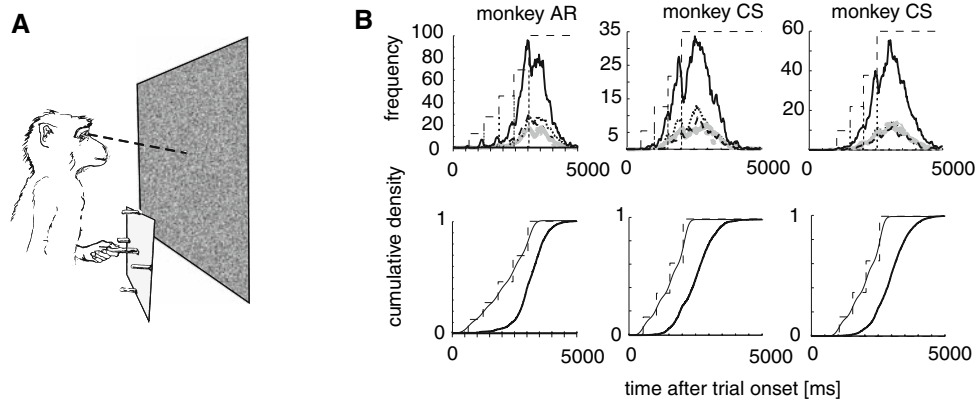


Fig. 1 **a** Experimental task. Monkeys faced a transparent screen onto which a stationary noise background was back-projected by a slide projector and onto which the stimulus was back-projected by a video projector. When visible stimuli (*bars* to the lower right of the monkey's fixation) were presented monkeys had to indicate the direction of motion with a hand movement to a corresponding touch bar. Stimulus motion direction, probability of stimulus onset time and visibility were randomized such that high levels of uncertainty about stimulus occurrence caused animal to indicate stimulus independent decisions (SIDs, "Methods"). In the current paper all our analyses exclusively deal with SID trials, i.e. when no visual stimulus was presented. **b** Distribution of SID timing (*black solid curve, upper*

row). SID timing for the four different directional choices are shown separately below (*grey and dotted lines*). SID times closely coincided (or followed) the possible stimulus presentation times, as evident by the small peaks in the distribution of SID timings. The step function shows the cumulative hazard function (normalized to 1). The lower row shows the cumulative hazard function (*dashed line*), the cumulative hazard function convolved with a Gaussian with 150 ms sigma (smooth line on top of the cumulative hazards function) and the normalized cumulative SIDs as a function of time after trial onset. In both animals the cumulative SID functions closely trailed the cumulative hazard function

putting an on-axis Fourier hologram of a piece of ground glass with Gaussian amplitude density distribution and constant power between 0 and 7 cyc/deg (von Seelen and Hoffmann 1976) into the beam of the slide projector. The mean intensity of the noise background was 1.7 cd/m² with a RMS contrast of 0.3053.

The stimulus contrast was varied from high contrast levels to invisible contrast (0%). Four contrast levels were tested for each neuron (17, 4, 2, and 0% contrast). Stimuli were back-projected with an EPS 4000 video projector (Electrohome) onto the translucent tangent screen, superimposed onto the stationary noise background. The stimulus and background intensity were measured using a photo-multiplier (EMI, 14 dinodes, 20S, aperture 0.04° of visual angle), and the linearity of measurements was ensured with 50% transmission neutral grey filters (Schott, Mainz).

Behavioural paradigm

The monkeys performed a reaction time task. As soon as they perceived (or believed they perceived) the direction of motion, they had to release the central bar and touch one of the four peripheral bars. These were positioned according to the directions of motion (see Fig. 1a). The release of the central touch bar was taken as the reaction time (RT). A 'go-signal' was never presented. After touching the peripheral touch bar, the monkey had to keep fixation for another 500 ms, during which the stimulus continued to

move through the receptive field (provided it was displayed).

Stimuli could move up, right, down, or leftwards within the aperture. Monkeys had to discriminate the direction of motion and indicate it by a corresponding hand movement to one of four touch-bars surrounding a central touch bar located in front of the animals. The animals could respond immediately when they saw a stimulus, i.e. they engaged in a reaction time task. Once a stimulus was presented they had 2,500 ms to indicate the direction of stimulus motion. On 'no stimulus' trials they had 5,000 ms to indicate a decision. Thereafter the trial was aborted without reward and the animal could start a new trial. On trials where they indicated the correct direction of stimulus motion they received a fixed quantity of 0.1 ml of juice for correct motion discrimination. Stimuli could occur 640, 1,240, 1,840, 2,440, and 3,040 ms after trial onset for monkey AR. Two different time series were used for monkey CS: (a) the stimulus could occur after 520, 1,020, 1,520, 2,020 ms, or (b) in a second set of experiments it could occur after 1,020, 1,520, 2,020, or 2,520 ms.

The probability of stimulus presentation, the stimulus luminance contrast, direction of motion, and its presentation time within a trial were all entirely randomized on a trial by trial basis, thus making the visibility of the stimulus, and its direction of motion entirely unpredictable, while its time of occurrence was only predictable in accordance with the hazard function (see below). The uncertainty of whether and when a stimulus would be

presented ensured that animals often indicated directional decisions in total absence of visual motion. Since these decisions occurred in complete absence of a moving visual stimulus we termed them stimulus independent decisions (SIDs). SIDs in the context of this paper could occur during ‘no stimulus trials’ as assigned by the computer. Additionally, SID trials within the context of this report were also trials where a stimulus should have been presented, but animals indicated a decision before the intended stimulus onset. If SIDs occurred before an intended stimulus presentation, the stimulus presentation was abandoned, i.e. the animals also indicate a decision in complete absence of a visual motion stimulus. These SIDs were caused by an ‘early’ (or anticipatory) choice of the animal, not by the computer generated trial selection. They were never rewarded, but are included in the paper as they are identical in terms of the physically display on the screen (the stationary structured background), and thus in terms of retinal stimulation. SIDs that occurred during computer generated ‘no stimulus’ trials were rewarded on 50% of the trials on average, but this probability of reward and the reward magnitude was adjusted according to the animal’s choice history (see below). This paper exclusively focuses on trials where SIDs of either type occurred.

Adjustment of reward probability and magnitude

Reward probability and magnitude were only adjusted for SIDs during computer generated ‘no stimulus’ trials, not on animal generated SID trials. The latter were never rewarded. Reward probability on computer generated ‘no stimulus’ SID trials was adjusted such that choice bias was reduced. At the same time it ensured that the animals could maximize their reward income on these SID trials by employing a strategy that kept track of previous choices and the associated income. Provided the animals learn these contingencies it would allow to determine whether directional choices are reflected differently in neuronal activity when choices are made while animals wait for a sensory stimulus, compared to when the identical choice is made based on the reward history. On trials when the computer selected a ‘no stimulus’ condition animals were randomly rewarded on 50% of the trials (mean) for SIDs. Reward probability on computer generated SID trials was adjusted according to the animal’s past SIDs, i.e. the animal’s decision history. The decision history consisted of the last 1,000 SIDs that had occurred, irrespective of whether it was a SID in relation to a no stimulus trial, or a SID that preceded an intended stimulus presentation. On SID occurrence the decision history was renewed in a ‘first in–first out’ manner, i.e. the SID 1,000 trials ago was eliminated from the buffer, the SID in position 999 was moved to position 1,000, position 998 was moved to

position 999, etc.. The SID history was stored to the computer disk, such that on each day the last 1,000 SIDs that had occurred during the previous session(s) were initially uploaded. SIDs that had occurred more recently were given more weight in calculating the current reward probability. The last 10 SIDs (block 1) had a weight of 0.5, SIDs 10–100 (block 2) had a weight of 0.3, and SIDs 100–1000 (block 3) had a weight of 0.2. The exact probability of a reward for a certain SID was calculated according to:

$$p(R)_B = \begin{cases} 0.5 + 0.5(1 - d_p/(s_B/n_d)) & |d_p \leq s_B/n_d \\ 0.5((s_B - d_p)/(s_B - s_B/n_d)) & |d_p > s_B/n_d \end{cases}$$

where $p(R)_B$ is the reward probability for a certain SID given the current block (1–3), d_p is the number of SIDs in the current block matching the current SID direction, s_B is the size of the current block (10, 90, 900), and n_d is the number of different SIDs possible (i.e. 4). Given that the different blocks of past decisions had different weights, the total reward probability was calculated according to:

$$p(R)_{\text{total}} = p(R)_{B1}0.5 + p(R)_{B2}0.3 + p(R)_{B3}0.2$$

where the three different $p(R)$ s are the reward probabilities based on the three blocks multiplied by the weighting factor. In addition to adjusting the probability of a reward, we also adjusted the reward magnitude based on the monkeys’ SID history. This ensured that animals could in principle estimate the current reward probability of the choice from the reward magnitude (a large reward is equivalent to high probability), and thus adjust choice accordingly (i.e. resample if a large reward was given, avoid this choice for a while if a small reward was given). Reward magnitude was adjusted by changing the solenoid opening time, i.e. adjusting the duration of reward delivery. In addition to linear increases/decreases of reward magnitude depending on choice history a step non-linearity was introduced for SIDs occurring at the transition of <50% reward probability and $\geq 50\%$ reward probability. The solenoid opening time was calculated according to

$$OT(p) = \begin{cases} 100 + 100p(R)_{\text{total}} & |p(R)_{\text{total}} \geq 0.5 \\ 100 - 100(0.5 - p(R)_{\text{total}}) & |p(R)_{\text{total}} < 0.5 \end{cases}$$

where $OT(p)$ is the opening time as a function of the current reward probability ($p(R)_{\text{total}}$) given the animals’ decision. The amount of reward was linearly related to the opening time from 50 ms opening time onwards, whereby a 50 ms opening time yielded 0.05 ml per reward, and every 1 ms increase added 0.001 ml to this. Opening times below 50 ms resulted in non linear decreases (accelerating decreases), such that at 25 ms opening time a reward yielded 0.01 ml, and basically no reward was delivered at opening times of <10 ms due to the inertia of the solenoid.

SID times and hazard function

For each monkey the SID times were determined as the time points when the hand disconnected from the central touch-bar. Thus, we refer to the decision time as the hand movement onset, although the decision itself will have occurred somewhat before the movement onset. The frequency of decisions as a function of the direction of the SID and as a function of time from trial onset was calculated. Additionally we determined the stimulus occurrence probability (the probability density function of stimulus occurrence), by calculating the frequency distribution of stimulus presentation as a function of time, and divided this distribution by the total number of trials (i.e. those when a stimulus was presented and those when no stimulus was presented). The relevant probability for the animal, of whether to indicate a decision if no stimulus has yet been detected is represented by the hazard function ($h(t)$). This is the probability that a stimulus will occur at any time given that no stimulus has been presented yet:

$$h(t) = \frac{p(t)}{1 - \int p(t)dt}$$

where $p(t)$ is the probability density function of stimulus occurrence, and the integral corresponds to its cumulative counterpart, the cumulative density function of stimulus occurrence (Smith 1995).

Behavioural controls

Throughout the trial (and until 500 ms after the indication of a decision) animals had to fixate the fixation spot within a window of 2° side length, i.e. $\pm 1^\circ$ from the fixation spot. Eye position was monitored by means of scleral search coils (optimum resolution of $1.5'$), sampled at 500 Hz.

Electrophysiology and assessment of basic neuronal properties

Recordings were performed using the ‘Eckhorn Matrix’ (7 channel, Uwe Thomas Recording, Marburg, Germany). Glass-insulated platinum–iridium electrodes (Uwe Thomas Recording, Marburg, impedance: 1.5–3 M Ω at 1 kHz) were advanced through guidetubes (outer diameter 305 μ m) into the brain. Up to seven guidetubes were inserted per recording session, the guide tubes were arranged in a circular array with one guidetube in the centre of six surrounding tubes (overall diameter was $915 \times 915 \mu$ m). Each guidetube was sharpened, such that all inserted guidetubes together formed a single tip. Amplified electrical activity from the cortex was band-pass filtered (0.3–10 kHz), and passed through oscilloscopes to Schmidt trigger thresholding (built in house) and spike

sorting devices (Alpha Omega, Nazareth, Israel). In monkey AR we focused on single unit recording (all cells reported here were subject to spike sorting), while in monkey CS we focused on single unit and multi unit recordings, whereby signals from some electrodes were subjected to spike sorting, while others were intentionally subjected to Schmitt-trigger thresholding, which resulted in multi-unit signals. The latter was done as synchrony between recording sites is more easily detected from multi- than single-unit recordings (Kreiter and Singer 1996; Roelfsema et al. 2004). For single unit recordings the quality of spike separation was controlled online by displaying the interspike interval distribution for each spike channel on the monitor of the recording personal computer.

Prior to recording, the receptive field location of each cell was mapped using a hand-held projector while the monkey fixated a central target on a dark background. Data were recorded only after the monkey had become well adapted to the background luminance ($\sim 0.5 \text{ cd/m}^2$) for approximately 20–30 min. In addition, the cell’s spikes (single and multi- unit activity) had to be well isolated (stable) for at least 5 min while the monkey performed the task before data acquisition began.

Direction selectivity

Initially we determined for each neuron whether it was directionally selective by calculating the direction index (DI). This was based on the neuronal activity following the period of 50–300 ms after a 17% contrast stimulus had been presented and the animal made a correct decision. For the calculation of the direction index, the direction of motion that yielded the largest activity was termed the preferred direction (PD), the direction of motion opposite to this direction was termed the anti-preferred direction (APD). The direction index (DI) was calculated according to:

$$DI = 1 - \frac{\text{act}_{\text{APD}}}{\text{act}_{\text{PD}}}$$

where act_{APD} was the neuronal activity associated with anti-preferred direction of motion, act_{PD} the neuronal activity associated with preferred direction of motion after subtraction of background activity. If the DI exceeded 0.5, the neuron was taken as direction selective. Only direction selective neurons are included in the current paper. For direction selective cells we then calculated the ‘vector average’ preferred direction. This preferred direction was the vector average across the four directions of stimulus motion at 17% luminance contrast as described previously (Thiele and Hoffmann 1996). It was calculated to analyse whether cells with more similar preferred direction showed stronger synchrony and rate correlation (see below). Due to

these two measures we will use two slightly different references to ‘preferred direction’ in this paper. We will refer to them as ‘preferred direction’ and ‘vector average preferred direction’, respectively, throughout the paper.

Definition of early and late SIDs

Early and late decisions were defined in the following way: We used the point where the experimental normalized cumulative hazard function reaches the value of 1 as a reference point. This point will be referred to as the point ‘where the cumulative hazard function reached its maximum’ in the remainder of the text. The point where the cumulative density function reached its maximum (Fig. 1) plus 500 ms was used as the cut-off between early and late SIDs. This resulted in a cut-off at 3,540 ms for monkey AR. In monkey CS early SIDs were defined as those that occurred within 2,520 ms after trial onset for sessions where the cumulative hazard function reached its maximum at 2,020 ms, and SIDs that occurred within 3,020 ms after trial onset for the sessions where the cumulative hazard function reached its maximum at 2,520 ms (see “Results”; supplementary material for additional analyses for different cut-offs). We did not use the point where the cumulative hazard function reached its maximum itself, because the animal’s choice in relation to low luminance contrast stimuli (2%) trailed the stimulus presentation time by ~ 500 ms (Thiele et al. 1999). Since the last stimulus could occur at the time the cumulative hazard function reached its maximum, the animal’s internalized timing of choices would probably trail this point by a similar amount, i.e. by 500 ms.

Neuronal choice probability analysis

For each cell we calculated a time resolved choice probability (Britten et al. 1996; Thiele et al. 1999). Therefore choices were subdivided into early and late SIDs, and the choice probability was calculated from 1,200 ms before the SID until 100 ms after the SID in a sliding window of 200 ms width in steps of 50 ms. We calculated choice probabilities provided at least five early and five late SIDs in preferred direction and at least five early and five late SIDs in anti-preferred direction were available. Choice probabilities based on only five trials could be regarded as being slightly unreliable, therefore we performed the same analysis provided there were (a) at least seven trials for each of the four conditions available, and (b) at least nine trials for each of the four conditions available. From these time resolved choice probabilities we calculated the population choice probability for early and late SIDs. A two-factor repeated measurement Kruskal–Wallis ANOVA was used to calculate whether the population choice probabilities for early and

late SIDs were significantly different from 0.5. Significance of choice probability for individual neurons was performed by using a permutation test (Britten et al. 1996).

Neuronal synchrony

Neuronal synchrony between pairs of units was calculated provided both units were directionally selective (see above), they both had a similar vector average preferred direction (angular difference $<120^\circ$), and at least ten trials for each of the four types of SIDs had occurred (i.e. early and late SIDs in preferred direction and anti-preferred direction). The criterion in relation to similar vector average preferred direction meant that some pairs were included where the individual units preferred somewhat different directions of motion (e.g. one unit preferred mostly ‘up’ while the other preferred mostly ‘right’). For such a case, upward and rightward SIDs were combined as being SIDs in preferred direction, while leftward and downward SIDs were combined as being SIDs in anti-preferred direction. For other combinations of preferred direction the respective choice directions were combined.

Cross-correlograms were calculated for each SID type at 1 ms resolution from -100 to $+100$ ms averaged over at least ten trials. We used two different approaches to calculate neuronal synchrony. The first approach is identical to the method published by Bair et al. (2001), which has been developed to account for firing rate co-variation at short time scales. We refer the reader to that paper for details. The other method is a variant of this measure which quantifies correlation of precise spike timing (R_{CCG}^*) after contributions from firing rate co-variation has been removed (Roelfsema et al. 2004). For both approaches we calculated the raw correlograms and subtracted the PSTH predictor to correct for possible decision locked synchrony (Aertsen et al. 1989). Firing rate co-variation at short time scales was quantified by the neuronal correlation coefficient (NCC), the calculation of which is described in detail in Bair et al. (2001). The neuronal correlation coefficient was calculated over correlogram time intervals of ± 5 , ± 10 , ± 20 , and ± 30 ms relative to time zero. Our overall results were robust with respect to the choice of time intervals. Most pairs had a correlogram peak width of $\sim \pm 10$ – 30 ms centred at time zero. Accordingly, the data plotted in the figures and the corresponding P values were derived from ± 20 ms time intervals. For a neuronal correlation coefficient to be included into our data set, the cross-correlograms had to have at least 200 entries. This criterion was usually far exceeded. As for the NCC we quantified R_{CCG}^* by determining the correlation coefficient from the autocorrelogram normalized area under the cross-correlogram from time intervals of ± 5 , ± 10 , ± 20 , and ± 30 ms relative to time zero.

Coarse firing rate correlation

Coarse rate correlation was calculated for the period of 500–100 ms before the 4 different SID types (early and late in preferred direction and anti-preferred direction, respectively), provided at least ten SIDs for each SID type were available. For each trial we calculated the activity (spikes/s) that occurred within the analysis window 500 to 100 ms before the SID for unit 1 and the activity (spikes/s) that occurred within the analysis window in the same trial for unit 2. From these joint trial by trial firing rates we calculated Pearson's correlation coefficient as a function of SID type for all directionally selective recording pairs. We thus obtained four correlation coefficients for each neuronal pair: (a) for early SIDs in preferred direction of the units, (b) for early SIDs in anti-preferred direction of the units, (c) for late SIDs in preferred direction of the units, and (d) for late SIDs in anti-preferred direction of the units.

Results

Animals were trained to detect the direction of motion of visual stimuli, which could move into one of the four cardinal directions (four alternative forced choice task, Fig. 1). High levels of uncertainty regarding visibility, stimulus timing and direction of motion caused them to indicate directional choices in total absence of visual motion stimuli (stimulus independent decisions, SIDs). In this paper we exclusively focus on trials where no visual motion stimulus was shown, i.e. on SID trials. These SID trials include trials where no stimulus was intended to be shown, as well as trials where the animals indicated premature choices (choices before stimulus onset), and no stimulus was presented as a consequence. Continuous adjustments of reward probabilities on SID trials based on the animal's choice history forced the animals to make current choices in light of previous choices and associated rewards if they intended to maximize their reward income once the cumulative hazard function had maxed out. The adjustments of reward probabilities were such that repetitive sampling of the same SID direction resulted in reduced reward probability and magnitude. Within these task constraints we found that animals indicated two different types of SIDs. A subset of SIDs closely followed the hazard function. The latter is defined as the probability of stimulus occurrence given that it has not yet occurred (see "Methods"). SIDs that occurred before the maximum of the cumulative hazard function plus 500 ms were defined as early SIDs (see supplementary materials for additional information and alternative cut-offs between early and late SIDs). The remaining SIDs occurred when

the probability that the relevant stimulus might occur had dropped to zero, and the cumulative hazard function had reached its asymptote (late SIDs > maximum of cumulative hazard function +500 ms; Fig. 1). To maximize reward income animals should have attempted to extract visual directional information before the cumulative hazard function had reached its maximum (although no such information was present). During this period they might have indicated a choice (early SID) if a specific neuronal population increased its activity relative to all other populations, although no relevant sensory stimulus in the external world for such an increase occurred. Downstream areas, involved in decision making, might interpret a selective increase in neuronal activity as if a visual motion stimulus of low visibility has been presented. The ensuing choice direction would then correspond to the preferred direction of the neuronal population showing the increased activity. We further hypothesized that choices after the maximum of the cumulative hazard function +500 ms (late SIDs) might have been governed by a 'cognitive' decision to participate in the reward lottery, thereby increasing income. We would at least assume that they were triggered by a cognitive decision not to wait for a stimulus any longer, and thus disregard activity from sensory areas. These hypotheses predict specific differences in terms of the ability to exploit the reward structure of the task, in terms of overall choice allocation, and in terms of the activity distributions associated with specific choices in motion area MT. We will first analyse whether animals exploited the reward structure differently during early and late SIDs, and then analyse the neuronal activity in area MT.

Late SIDs and reward maximization

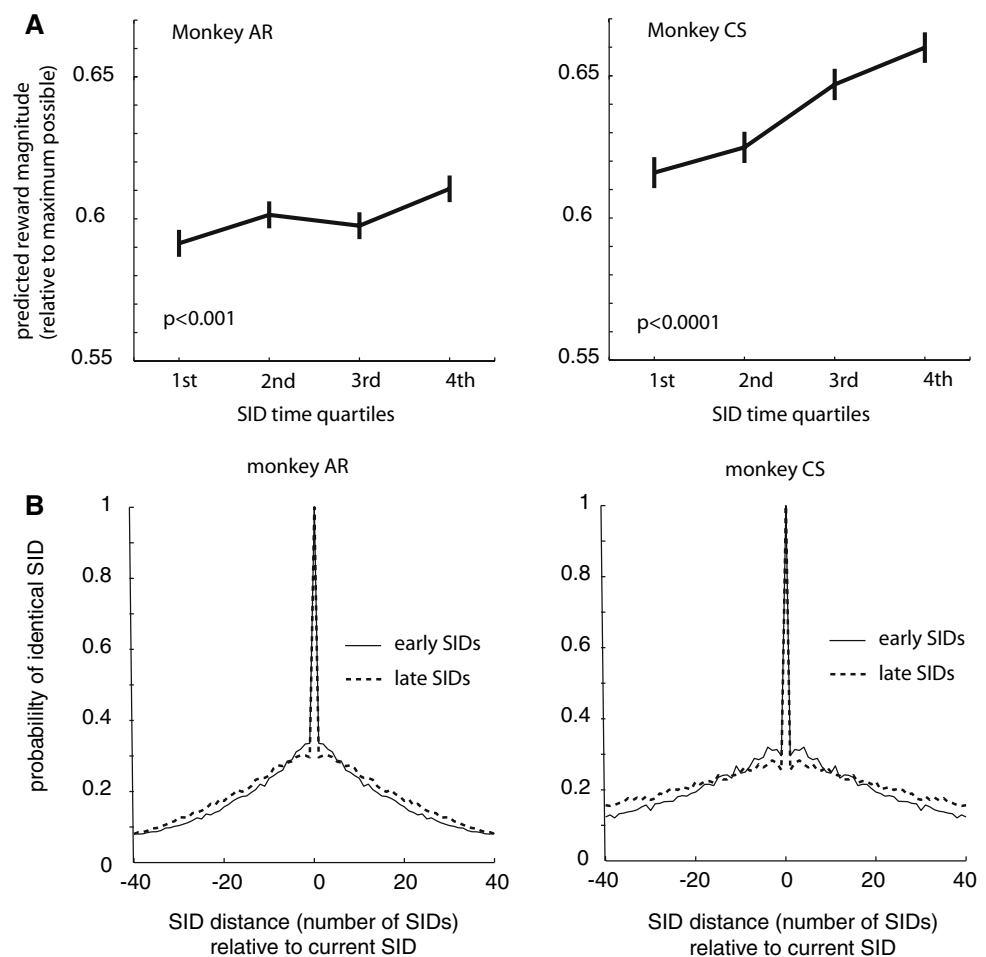
For late SIDs optimal decision making, i.e. reward maximization, should be based on choice history. Maximizing the reward yield on SID trials required the animals to distribute SIDs among the different targets in a manner that is adapted to the statistics to the task, i.e. their foraging should vary in line with changes in reward probabilities in the environment (here the experimental setting). Our task required the animal to dynamically alter choice distributions depending on the previous choices, as the reward probability and quantity were continuously adjusted by these previous choices. Thus, the animals should engage in a dynamic foraging task, where the dynamics of the changing reward structure were dependent on previous sampling.

If animals had learned at least parts of the reward contingencies of the task, but mostly exploited them during late SIDs, then late SIDs should be associated with higher reward magnitude probabilities than early SIDs. The

predicted reward yield corresponds to the normalized reward magnitude that would be given if the SID resulted in a reward. The conditional ‘would be given’ is due to the fact that on each trial the computer decided randomly whether or not the SID will be rewarded. Thus, although on a given trial the reward probability could be 0.8, and the associated predicted reward magnitude could be 0.9 (relative to the maximum possible), the SID might still not yield a reward. To investigate the relationship between SID timing and predicted reward magnitude we took a variety of different approaches. Firstly, we ordered the SID times during each experiment, and subdivided these trials into four equally sized groups. Thus, group 1 contained the quartile of the fastest SIDs during that session, group 2 the second fastest quartile, etc. We then pooled the respective SID groups from the different sessions and calculated the average predicted reward magnitude (and probability) that was associated with these SID groups. This approach ensured that equal numbers of SIDs were present in each group, and allowed to average across different hazard functions that were used in monkey CS. The results are shown in Fig. 2a. We found that SIDs occurring earlier in a

trial were associated with a significantly lower predicted reward magnitude than late SIDs ($P < 0.001$, ANOVA). Post-hoc testing revealed that the fourth quartile group from monkey AR yielded significantly larger predicted reward magnitudes than the first quartile group ($P < 0.01$). In monkey CS the third and fourth quartile group yielded significantly larger predicted reward incomes than the first and second quartile groups ($P < 0.01$), while there was no significant difference between the two members of these two subgroups. Secondly, we calculated the predicted reward yield for SIDs that were defined as early and late SIDs based on the criteria described in “Methods” for each monkey. In both animals the predicted reward yield was significantly higher for late SIDs (predicted yield monkey AR, early = 0.591 ± 0.29 , late = 0.607 ± 0.23 ; $P = 0.018$; monkey CS, early = 0.637 ± 0.217 , late = 0.655 ± 0.216 ; $P < 0.001$, rank sum test). Thirdly, we calculated the correlation between predicted reward magnitude and the SID time. SID times were converted to z scores for each experimental session, as they could vary somewhat between sessions. Moreover, the conversion was necessary because we used two different hazard functions

Fig. 2 a Predicted reward magnitude associated with SIDs that occurred during different times of the trial. SIDs from each session were ordered according to time of occurrence relative to trial onset within the session, and these ordered SIDs were subdivided into four equally sized groups, thus the first quartile contained the fastest SIDs (relative to trial onset) during that session, the second quartile the second fastest SIDs etc. The figure demonstrates that the predicted reward magnitude that would occur with the respective SID direction increased from the first to the fourth quartile. Although this increase was somewhat smaller in monkey AR than in monkey CS, it was nevertheless highly significant (ANOVA). **b** Autocorrelation of SIDs. The figure shows the probability that the same SID direction will be indicated before or after the current SID (data point zero). The x -axis plots the distance in terms of the number of intervening SIDs relative to the current SID



in monkey CS. For both monkeys we found a small but significant positive correlation between choice times and the predicted reward income (monkey AR, $r = 0.08$, $P = 0.011$; monkey CS, $r = 0.104$; $P = 0.0068$). Finally we calculated the autocorrelation of choices for early and late SIDs. The autocorrelation shows the probability of indicating the same choice as a function of distance from the current SID (i.e. what is the probability of choosing the same SID again on the next SID trial, on the SID trial after the next, etc.). Figure 2b shows the autocorrelations for early and late SIDs for both monkeys. Both monkeys were more likely to quickly re-sample the same SID during early SIDs than during late SIDs. This was slightly more pronounced in monkey CS, but was also present in monkey AR. In both monkeys the autocorrelation functions associated with early and late SIDs were significantly different from one another ($P < 0.001$, Kolmogorov–Smirnov test). Moreover, they were significantly different from an autocorrelation yielded by response randomization ($P < 0.001$, Kolmogorov–Smirnov test; response randomization would yield a completely flat autocorrelation at a level of 0.25). The finding that the mean of the autocorrelation functions for both monkeys is lower than 0.25 is due to the fact that a limited number of SIDs occurred in each trial, and thus the first SID will not have a predecessor, or the SID that occupies position number 10 from the end of the experimental session will only have 9 other SIDs following, thus no entries are available for the autocorrelation from position 10 to 40.

How does the monkey's performance compare to well defined strategies? One possible strategy would be response randomization. Response randomization would yield an average predicted reward magnitude of 0.64 (based on 10^7 model SIDs). A better strategy is uniform choice allocation ('circling', e.g. up-right-down-left-up...). Circling would yield an average reward magnitude of 0.76 (based on 10^7 model SIDs). A yet better performance could be achieved by perfect memory, where the reward based on the last 1,000 SIDs could be predicted for each choice. An animal with perfect memory would then choose the direction associated with the highest reward magnitude/probability. This would yield an average reward magnitude of 0.77. The performance of both monkeys demonstrates that they used neither 'circling' nor did they have perfect memory. It shows that the strategies employed during early and late SIDs were suboptimal, as they did not come close to obtaining the maximum possible yield. The performance of monkey CS is close to what would be obtained by response randomization, but the autocorrelation clearly shows that the monkey did not use a randomization strategy. Although we do not know exactly what strategy was used by the animals, the important finding in relation to the behavioural data was that the monkeys allocated their

choices differently during early and late SIDs. They obtained significantly larger yields during late SIDs, and indicated fewer immediate choice repetitions during late SIDs. Such a change in strategy requires the animal to have at least some form of memory representation of the previous choice(s), and of the specific period within which the choice is made (early/late).

SIDs and associated neuronal activity in MT

We reasoned that early SIDs might have been triggered by an internal signal arising in sensory areas, possibly in MT, that the monkey could have mistaken for coding of an external stimulus. In that case activity of directionally selective neurons or synchrony among neurons in MT with similar preferred direction should have been enhanced before early SIDs, provided the direction of choice corresponded to the direction of motion these neurons preferred. This enhancement of activity should not occur if the direction of choice was opposite to the preferred direction of the neuron, i.e. a choice in anti-preferred direction of the neuron. We also predicted that the activity changes mentioned above should not occur prior to late SIDs in preferred direction and in anti-preferred direction, because these choices were more governed by attempts to maximize the reward income by taking choice history into account, where signals from sensory areas are unlikely to yield useful information. We tested these predictions by calculating the neuronal activity associated with early SIDs in preferred direction and in anti-preferred direction, and the activity associated with late SIDs in preferred direction and in anti-preferred direction of each neuron. Figure 3 shows two examples of neuronal activity associated with early and late SIDs, as well as SIDs in preferred direction and anti-preferred direction. For both cells the activity with early SIDs in preferred direction was larger than the activity with late SIDs in preferred direction, and it was larger than the activity with early or late SIDs in anti-preferred direction.

To assess the significance of the activity difference for the different types of SIDs (early vs. late, preferred direction vs. anti-preferred direction) at the population level we calculated the time resolved choice probability (Britten et al. 1996) for each cell for early and late SIDs, provided we had at least five decisions for each of the four different SID groups for a given cell (the median and range of SID numbers for our population of cells that fulfilled this criterion were: $PD_{\text{earlySIDs}}$ median $n = 17$, range $n = 101$; $APD_{\text{earlySIDs}}$, median $n = 18$, range $n = 86$; PD_{lateSIDs} , median $n = 21$, range $n = 74$; APD_{lateSIDs} , median $n = 22$, range $n = 77$). Figure 4 shows time resolved choice probabilities individually for the two monkeys and the pooled data sets. We calculated a two-factor repeated

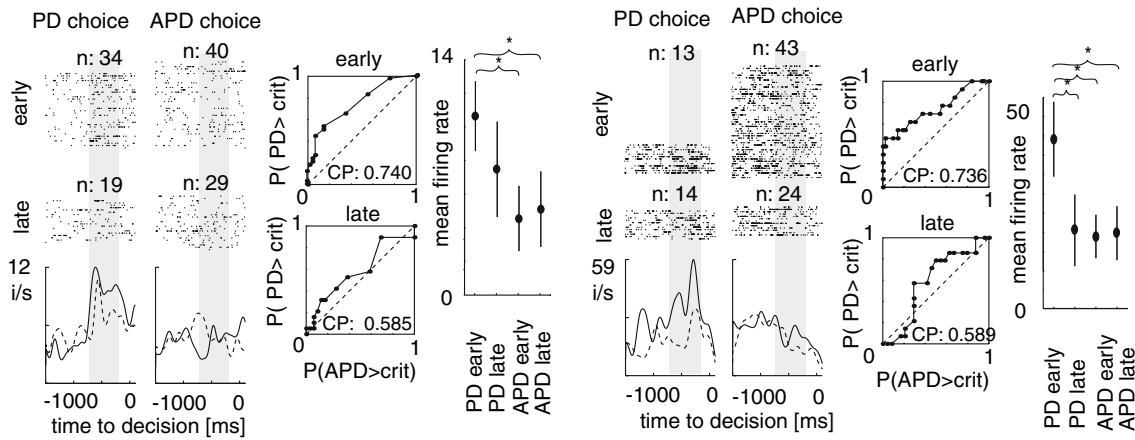


Fig. 3 Example cells of neuronal activity and choice probability for early and late SIDs as a function of choice. Raster plots show single trial activity associated with early SIDs (*upper rasters*) late SIDs (*lower rasters*) and choices in preferred (PD, *left*) and anti-preferred (APD, *right*) direction of the neuron. *Black solid* histograms show activity associated with early SIDs, *dashed histograms* activity with late SIDs. *Middle panel* of each subplot shows neuronal choice probability (CP) associated with early (*upper graph*) and late SIDs (*lower graph*) calculated from choices in preferred vs. anti-preferred

direction. It plots the probability that the activity with choices in preferred direction exceeds any given cut-off activity [$P(\text{PD} > \text{crit})$] versus the probability that the activity with choices in anti-preferred direction exceeds the given cut-off activity [$P(\text{APD} > \text{crit})$]. The time period (−700 to −200 ms) to calculate choice probability and mean activity is indicated by the grey area in the raster plots and histograms. *Panels on the right* in each subplot show the mean activities for the different types of SIDs during that time period. *Brackets* denote significant activity differences

measurement Kruskal–Wallis ANOVA on the population data to determine whether choice probability varied as a function of SID type (early vs. late SIDs, Factor 1), and the time prior to the SIDs (Factor 2). Choice probability was significantly affected by the SID type ($P = 2.813e-14$, Factor 1 main effect), and by the time prior to the SID ($P = 0.0149$, Factor 2 main effect). There was a significant

interaction, i.e. the size of the difference between early and late choice probabilities depended on the time bin relative to choice time ($P = 0.0016$, choice type \times time to choice interaction). The most consistent difference between early SIDs and late SIDs occurred from ~500 to ~100 ms before SID occurrence. We thus used this time window for all further analyses unless otherwise noted.

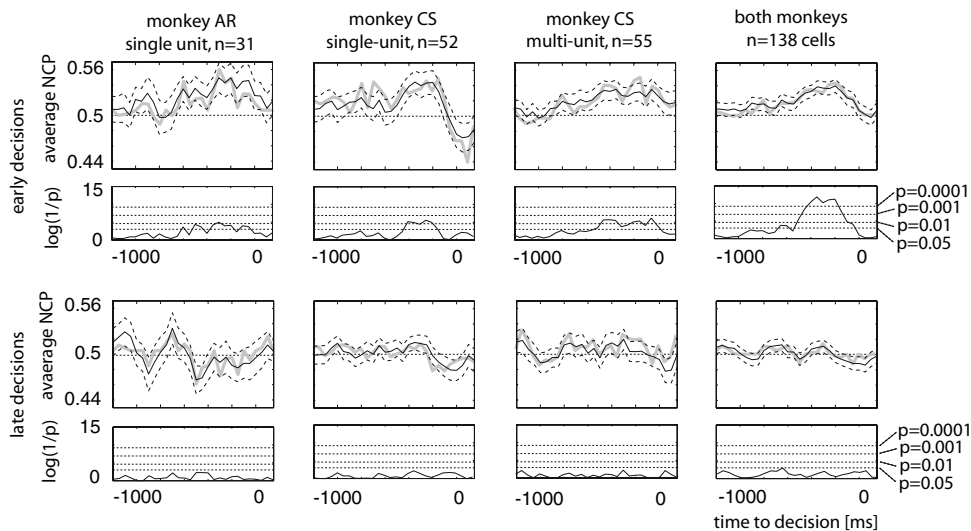


Fig. 4 Population choice probability as a function of time relative to SIDs for early and late SIDs. Choice probabilities were calculated within a sliding window of 200 ms width from −1,200 to +200 ms in 50 ms steps relative to the SID occurrence for each cell. These time resolved choice probabilities were averaged to yield population choice probability. The *black curve* shows the mean choice probability, *dashed curves* show SEM. The *grey curve* shows the median

choice probability. Population choice probability peaked at about 500–100 ms before SIDs during early decisions, while no significant deviation from 0.5 was present for late SIDs. The *boxes* below the average choice probabilities show the significance level for the respective time period (H_0 : median choice probability = 0.5, signed rank test)

We next calculated the choice probability for each neuron for early and late SIDs in order to determine the number of cells where choice probability was significantly affected by choice type (early vs. late) by using a permutation test (Britten et al. 1996). Distributions of choice probabilities are shown in Fig. 5. A total of 29 neurons showed significant choice probabilities during early SIDs, 22 of these were larger than 0.5. During late SIDs 19 neurons showed significant choice probabilities, 13 of

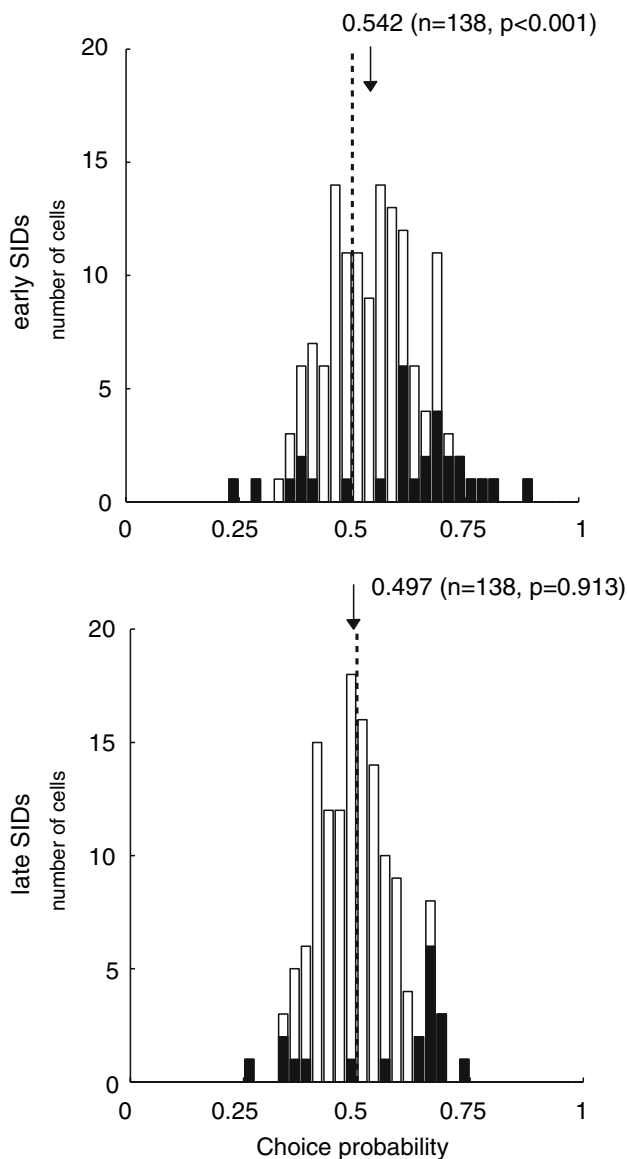


Fig. 5 Distribution of choice probabilities for our sample of units, calculated for the time period from -500 to -100 ms before the SID. The upper graph shows the distribution of choice probabilities for early SIDs, the lower graph the distribution for late SIDs. *Black bars* show the number of cells where the choice distribution was significant (permutation test). *Arrows* indicate the distribution mean based on 138 cells. *P* values indicate the probability that the mean was equal to 0.5

which were larger than 0.5. The median choice probability for early SIDs was significantly different from 0.5 ($P < 0.001$, signed rank test), while the median choice probability for late SIDs was not ($P = 0.913$, signed rank test).

Analyses of choice probability, median population activities, and activity ratios, were done on cells with average firing rates of >1 sp/s during SIDs ($n = 130/138$), because calculation of choice probabilities on cells that hardly fire is unreliable. We have also performed the analyses on the full data set (the above mentioned eight cells included). The outcome of the analysis on the full data set was similar to the outcome when only 130 cells were included; in effect the differences of the results between the various groups were even larger when these 8 cells were included.

The difference between choice probability values for early and late SIDs could be due to selectively increased activity associated with early SIDs in preferred direction (relative to the other three SID types. Alternatively, the difference in choice probability could be due to increased activity associated with late SIDs in anti-preferred direction (compared to late SIDs in preferred direction), while at the same time the activity level associated with SIDs in preferred direction was similar for early and late decisions. This is an important distinction, as the latter would not support our hypothesis that early SIDs were triggered by a signal arising in sensory areas. To investigate these scenarios we calculated the mean activity associated with early and late SIDs in preferred direction and anti-preferred direction during the time period from 500 to 100 ms before each SID. The activity before early SIDs in preferred direction was significantly higher than the activities before late SIDs in preferred direction and it was significantly higher than the activity before early or late SIDs in anti-preferred direction ($P < 0.001$, RM Kruskal–Wallis ANOVA, post hoc testing). None of the remaining 3 groups were significantly different from one another ($P > 0.05$, RM Kruskal–Wallis ANOVA, post hoc testing). This demonstrates that the activity for early SIDs in preferred direction was enhanced relative to all other SIDs, and could thus constitute a signal that the animals mistook for the presence of a stimulus (possibly akin to an illusory percept). Additional details regarding the different activity levels are given in Table 1, and the distributions of the log ratios of activities with different SID types are plotted in Fig. 6. The mean of the log ratio of activities with SIDs in preferred direction vs. anti-preferred was 0.201, which means that the activity with early SIDs in preferred direction was on average 22.3% higher than the activity associated with early SIDs in anti-preferred direction. The distribution's mean was significantly different from zero ($P < 0.001$, two tailed *t* test). For late SIDs the distributions mean was 0.064, i.e.

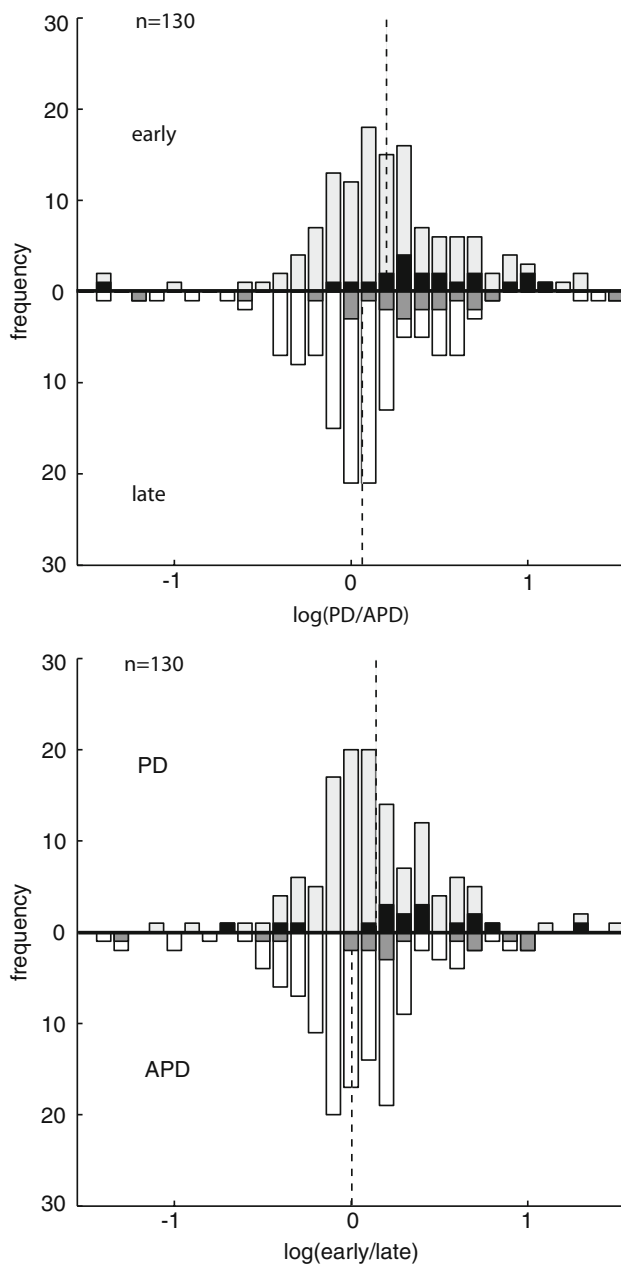


Fig. 6 Distribution of firing rate ratios for decisions in preferred (PD) and anti-preferred direction (APD) for early and late decisions (*upper panel*). The distribution of firing rate ratios between early SIDs in preferred direction and early SIDs in anti-preferred direction is plotted upwards (*light grey; black bars* show cells where the difference was significant), while the respective distribution for late SIDs is plotted downwards (*white bars, dark grey bars* show cells where the difference was significant). The *lower panel* shows the ratios between choices in preferred direction (early vs. late SIDs, *grey bars* for cells where the difference was not significant, *black bars* for cells where the difference was significant), and the respective distributions of activity ratios for choices in anti-preferred direction (*white bars* non-significant differences at the single cell level, *dark grey bars* significant differences at the single cell level). The *dashed lines* show the means of the log distributions

the activity with $\text{PD}_{\text{lateSID}}$ was on average 6.6% larger than the activity with anti-preferred direction $_{\text{earlySID}}$, which was not significantly different from zero ($P = 0.096$, two tailed t test). The two distributions ($\text{PD}_{\text{earlySIDs}}/\text{APD}_{\text{earlySIDs}}$ vs. $\text{PD}_{\text{lateSIDs}}/\text{APD}_{\text{lateSIDs}}$) were significantly different from one another ($P < 0.01$, two-tailed t test). Additionally we calculated the log ratios of early and late SIDs in preferred direction, and early and late SIDs in anti-preferred direction. The mean for $\text{PD}_{\text{earlySID}}/\text{PD}_{\text{lateSID}}$ log ratios was 0.141, i.e. the activity associated with $\text{PD}_{\text{earlySID}}$ was on average 15.1% larger than the activity with $\text{PD}_{\text{lateSID}}$ (Ho: distribution's mean = 0; $P < 0.01$, two-tailed t test). The mean of the log ratio $\text{APD}_{\text{earlySID}}/\text{APD}_{\text{lateSID}}$ distribution was 0.004, i.e. the activity with $\text{APD}_{\text{earlySID}}$ was on average 0.4% larger than the activity with $\text{APD}_{\text{lateSID}}$ (Ho: distribution's mean = 0, $P = 0.91$, two-tailed t test). The two distributions of the log ratios were significantly different from one another ($P < 0.01$, two-tailed t test). In summary we found that the activity with $\text{PD}_{\text{earlySID}}$ was significantly increased relative to all other types of SIDs.

We additionally determined how many cells showed a significant activity increase with choice in preferred direction during early SIDs. Significant differences occurred in a total 18/138 cells (13.4%, $P < 0.05$, two-tailed t test). To determine whether these few cells might account for the difference we see at the population level we eliminated them from our cell sample. Even without these cells the effect of increased activity associated with $\text{PD}_{\text{earlySID}}$ remained significant ($P < 0.001$, RM Kruskal–Wallis ANOVA, post hoc testing). We thus conclude that the effect was not due to a small minority of cells, but was present in the majority of cells, even if it did not reach significance levels for each of these individual units.

In principle these choice probabilities could have been due to small eye movements that induce selective activity in MT cells, which then trigger the animal's choice. We performed a variety of control analyses which demonstrate that small eye movements were not a contributing factor. These analyses are described in detail in supplementary materials.

SIDs and neuronal synchrony in MT

The selective increase of activity prior to early SIDs in preferred direction, albeit highly significant, was relatively small (Table 1). It has been proposed that integration of EPSPs in neurons benefits from synchronized inputs (Fries et al. 2001). According to this idea the impact on downstream stages would be larger when neurons with similar vector average preferred direction fire more synchronously (Singer 1999; Fries et al. 2001; Bruno and Sakmann 2006),

Table 1 median cell activity (25, 75%) associated with different types of decisions, and different eye-movement controls

SID type	<i>n</i>	Median (sp/s)	25%	75%
Cell sample with at least five trials per SID group				
PD early	138	17.12	5.83	34.29
APD early	138	13.92	5.38	31.19
PD late	138	14.07	5.27	29.52
APD late	138	12.78	4.58	30.00
Cell sample with at least five trials eye window ± 0.15				
PD early	122	18.00	6.00	33.89
APD early	122	14.87	5.44	30.56
PD late	122	13.77	5.28	30.28
APD late	122	13.30	4.58	31.43
Cell sample with at least five trials, saccades and drifts excluded				
PD early	112	17.42	6.00	28.05
APD early	112	13.74	5.28	25.03
PD late	112	13.64	5.44	23.61
APD late	112	13.78	5.23	24.14
Cell sample with at least nine trials per SID group				
PD early	75	24.18	8.40	39.34
APD early	75	18.15	7.72	34.34
PD late	75	17.58	6.27	32.32
APD late	75	16.34	6.57	35.03

and the strength of synchronization prior to a specific choice would then depend on the angular difference in preferred direction between the neuronal pairs. In both monkeys we recorded neuronal activity in area MT from up to seven different electrodes simultaneously. To test the above proposal we calculated the cross-correlation separately for early and late SIDs in preferred direction and anti-preferred direction. We did this for two groups of cell pairs, those whose vector average preferred direction differed by less than 60° , and those whose vector average preferred direction differed by more than 60° , and by less than 120° . We will refer to these two groups as the ‘ 60° group’ and the ‘ 120° group’. Neuronal synchrony for pairs of these groups was included in the analysis if at least ten trials were available for each SID type (see supplementary material for additional detail).

Figure 7a–d shows cross-correlograms from four different cell pairs for the different SID types. The examples in Fig. 7a–c are from the 60° group, the example in Fig. 7d is from the 120° group. The units shown in Fig. 7a–d exhibited strong synchrony prior to early SIDs in preferred direction, while neuronal synchrony prior to early SIDs in anti-preferred direction was lower. Synchrony prior to late SIDs in preferred direction and anti-preferred direction was also lower when compared to synchrony before early SIDs in preferred direction. From most of these examples it is additionally evident that the side flanks of the cross-

correlograms are higher during early SIDs in preferred direction compared to cross-correlograms associated with the other SIDs. This could be due to stronger firing rate correlation for these SIDs. To remove the possible confound of mixing effects of firing rate correlation and precise spike time synchrony we calculated the cross-correlation between neurons as described in detail in by Roelfsema et al. (2004), for which we also use the term R_{CCG}^* . We calculated R_{CCG}^* from spikes that occurred during the period from 500 ms until 100 ms before the SIDs for the four different SID types. We quantified the strength of precise spike timing correlation by calculating R_{CCG}^* for the time period of ± 5 , ± 10 , ± 20 , and ± 30 ms relative to the trigger centre bin. Qualitatively, the results were immune to the exact choice of the time period. Cell pairs of the 60° group exhibited significantly more spike time synchrony than cell pairs from the 120° group ($P < 0.001$, Wilcoxon rank sum test, compare Fig. 7e and f). Each group was further subdivided into subgroups according to their receptive field overlap. Overlap of receptive field was expressed as the total overlap in receptive field area (deg^2) divided by the sum of the individual receptive field areas. Each group was separated according to the median receptive field overlap, i.e. the first subgroup consisted of the cells with ‘less overlap’, while the second subgroup consisted of the group with ‘more overlap’. Receptive field overlap per se did not significantly affect the strength of spike time synchrony of our sample (60° group effect of RF overlap: $P = 0.645$, 120° group effect of RF overlap: $P = 0.890$, Wilcoxon rank sum test). For the 60° groups neither choice direction ($P = 0.2252$ more RF overlap subgroup; $P = 0.1091$ less RF overlap subgroup) nor choice type ($P = 0.3763$ more RF overlap subgroup; $P = 0.6253$ less RF overlap subgroup) alone had a significant influence on the strength of synchrony. However, a significant interaction was found between choice type and choice direction ($P = 0.041$ more RF overlap subgroup; $P = 0.0231$ less RF overlap subgroup, 2 factor RM-ANOVA). Post hoc comparison revealed that early SIDs in preferred direction yielded significantly more synchrony than early SIDs in anti-preferred direction or than for late SIDs in preferred direction for the 60° group with more RF overlap ($P < 0.05$, Fig. 7e grey bars), while late SIDs in anti-preferred direction yielded significantly more synchrony than late SIDs in preferred direction ($P < 0.05$, Fig. 7e black bars) for the 60° group with less RF overlap ($P < 0.05$, Fig. 7e grey bars). For the 120° group no significant differences were found for choice time, direction, or an interaction between these two factors for either subgroup ($P > 0.05$ 2 factor RM-ANOVA). These data demonstrate that precise spike time synchrony prior to SIDs were slightly, but significantly different before specific choices, provided cells

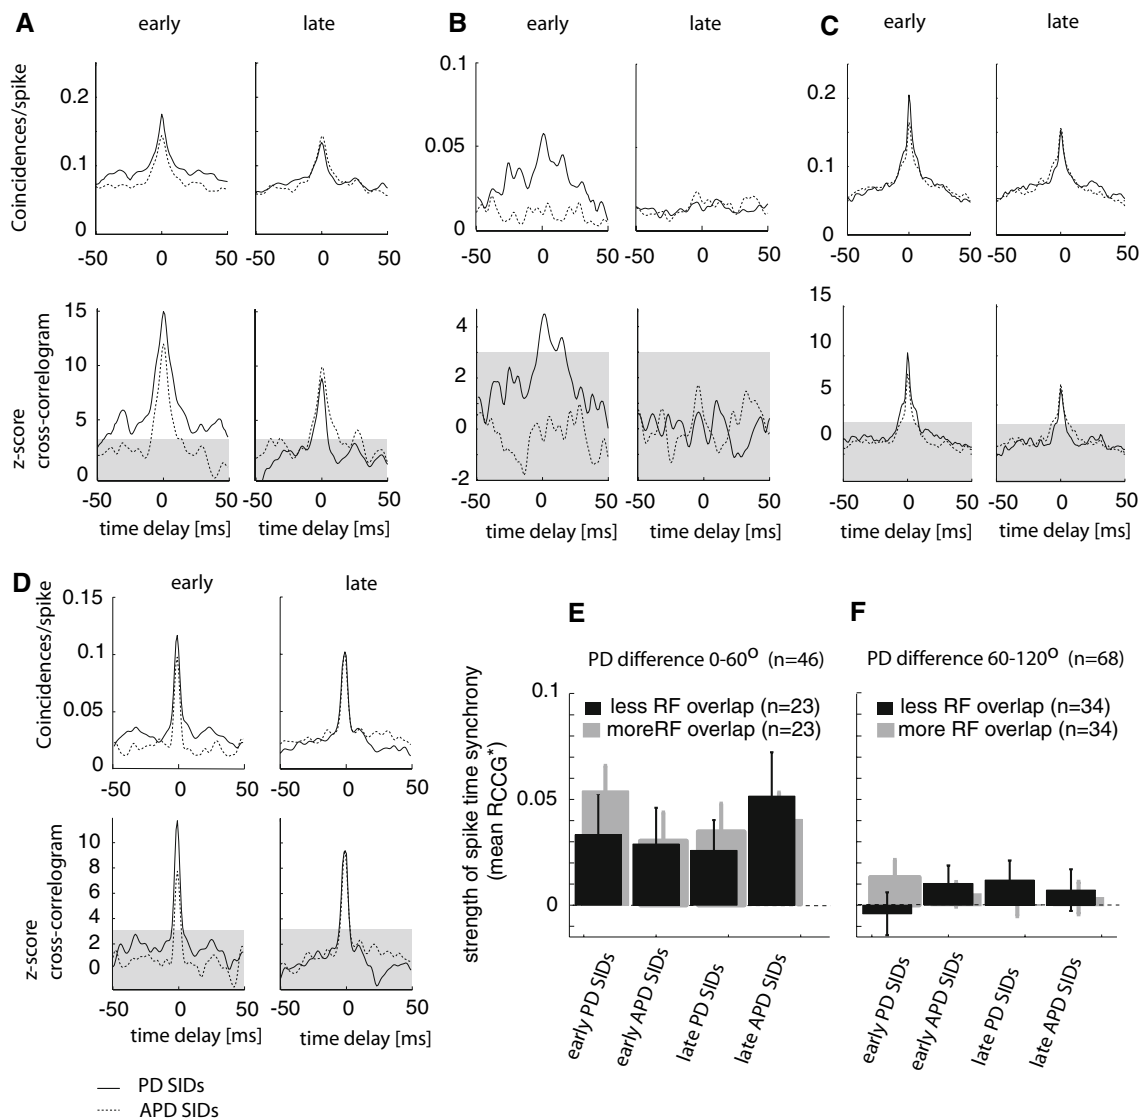


Fig. 7 a–d Cross-correlograms for pairs of units prior to SIDs (500–100 ms window) depending on SID type. *Solid curves* show cross correlograms for choices in preferred direction, *dashed lines* for choices in anti-preferred direction. *Left columns* in each *sub-panel (a–d)* show the cross correlograms associated with early SIDs, *right columns* the respective cross correlograms associated with late SIDs. *Upper rows* show cross correlograms normalized to coincidences per spike, the *lower row* shows the cross correlograms normalized to the z score. Peaks extending from the *grey area* show significant correlations. Cross-correlograms were smoothed with a Gaussian kernel of 3 ms standard deviation. All pairs were more synchronous before

shared a similar preferred direction. However, the pattern of results defies a straightforward explanation within the context of our experiment.

Firing rate correlation

Synchrony assesses neuronal interaction on a very short time scale. Another measure of increased coupling between

early SIDs in preferred direction than in anti-preferred direction and than before any late SID. **e** Average strength of synchrony for the population of unit pairs with similar vector average preferred direction, separated according to the degree of receptive field overlap. Strength of synchrony was quantified by the correlation coefficient calculated from ± 20 ms relative to the correlogram centre. **f** Strength of synchrony for cells that had more dissimilar vector average preferred directions (within 60° and 120°), separated according to the degree of receptive field overlap. The *four bars on the left* show the strengths of synchrony as a function of SID type, calculated during the period from -500 to -100 ms before the SID

units is firing rate correlation. We calculated firing rate correlation between simultaneously recorded units, and subdivided these into units with similar vector average preferred direction (0°–60°, i.e. the 60° group) and those with less similar vector average preferred directions (60°–120°, i.e. the 120° group). We used two different approaches to calculate firing rate correlation. (a) We measured firing rate correlation that can occur during short

time periods in the two units by applying a spike time cross-correlation based method (Bair et al. 2001; Roelfsema et al. 2004). We will refer to this measure as ‘fine firing rate correlation’. (b) We measured firing rate correlation using single trial based mean firing rates of the two units. We will refer to this measure as ‘coarse firing rate correlation’. We will first report the results for the fine firing rate correlation. Figure 8a shows the average fine firing rate correlation calculated for the period from 500 ms until 100 ms before the SIDs for the four different SID types. Cell pairs of the 60° group exhibited significantly more fine firing rate correlation than cell pairs from the 120° group ($P < 0.0001$, Wilcoxon rank sum test). Neither choice direction ($P = 0.4137$) nor choice type ($P = 0.4545$) alone had a significant influence on the strength of the fine firing rate correlation. However, a significant interaction was found between choice type and choice direction ($P = 0.0078$, 2 factor Kruskal–Wallis RM ANOVA). Post hoc comparison revealed that early SIDs in preferred direction yielded significantly higher fine firing rate correlation than early SIDs in anti-preferred direction (Fig. 8a grey bars). Notably, the increase in fine firing rate correlation was not present throughout the entire trial (Fig. 8a, compare rightmost bar to the first bar). For the

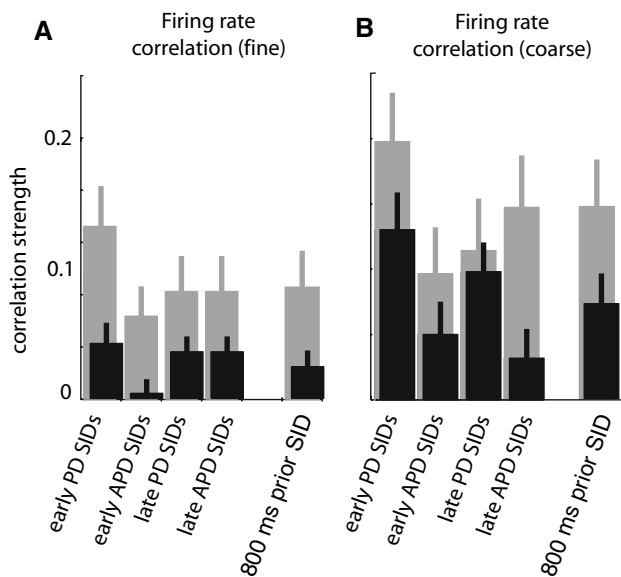


Fig. 8 Firing rate correlation between simultaneously recorded units. **a** Fine firing rate correlation calculated from the area under the cross-correlogram (see “Methods”). Units were separated according to similarity of vector average preferred direction in line with the approach for the synchrony analysis. Firing rate correlation was calculated within a 400 ms window from -500 to -100 ms before the SID (left bars) and from $-1,000$ to -600 ms before early SIDs in preferred direction (right bar). **b** Coarse firing rate correlation calculated from the trial by trial mean firing rates (see “Methods”). All conventions as in **a**. Grey bars show cell pairs with similar vector average preferred direction (difference $<60^\circ$), black bars cell pairs with more dissimilar vector average preferred direction

120° group no significant differences were found for choice time ($P = 0.3922$), direction ($P = 0.0743$), or an interaction between these two factors ($P = 0.3112$). Additional information regarding the distribution of fine firing rate correlation strengths, its time course, dependence on receptive field overlap and single vs. multi unit recording sites is provided in supplementary material.

Coarse firing rate correlation yielded very similar results, although overall correlation strength was somewhat higher for this measure (compare Fig. 8a, b). For the 60° group neither SID time ($P = 0.4439$, 2-factor RM ANOVA) nor direction ($P = 0.2670$, 2-factor Kruskal–Wallis RM ANOVA) alone yielded significant differences in coarse firing rate correlation, but a significant interaction was found ($P = 0.0104$, time \times direction interaction, 2-factor Kruskal–Wallis RM ANOVA), whereby the coarse firing rate correlation was significantly increased before early SIDs in preferred direction. The 120° group yielded significantly smaller coarse firing rate correlation overall ($P < 0.0001$, rank sum test) when compared to the 60° group. In this group we found a significant main effect of choice direction ($P = 0.0023$), but not of choice type ($P = 0.2113$). We did not find an interaction between the two main factors for the coarse firing rate correlation in the 120° cell group. While for this latter group coarse firing rate correlation was significantly larger for choices in preferred direction irrespective of the time (early vs. late SIDs), it is important to emphasize that the overall coarse firing rate correlation was significantly lower ($\sim 50\%$ less) when compared to the 60° group (Fig. 8b). In line with the findings for fine firing rate correlation, the increase was restricted to a short time period of ~ 500 – 100 ms before the occurrence of the SID (Fig. 8b, rightmost bar in comparison to the first bar). For additional information regarding the distribution of coarse firing rate correlation and its dependence on receptive field overlap see supplementary material.

Discussion

During the course of a trial animals were required to shift from reliance on sensory information to reliance on internally stored memory and reward prediction signals. This necessity was reflected in the animals’ behaviour. They exploited the reward structure more efficiently during late SIDs (their choices yielded higher reward probabilities, although these differences were small). In conjunction with these differences regarding choice for early and late SIDs we found systematic activity and firing rate correlation changes in area MT prior to directional decisions while animals actively waited for a sensory stimulus (which never occurred).

The difference between the firing rate on SIDs in preferred direction and those in anti-preferred direction is reminiscent of the activity differences found with hits and misses during trials when stimuli were presented (Thiele et al. 1999). There it was found that neurons were more active on correct trials than incorrect trials, whereby the direction of the error additionally influenced the activity levels (errors 90° off the correct direction resulted in less reduction than errors opposite a stimulus that moved in preferred direction).

It is possible that the increased activity and firing rate correlation was erroneously interpreted by downstream neurons (and ultimately by the animal) to be elicited by a motion stimulus, and thus triggered the decision. An increase in firing rate with the direction of choice for MT neurons has been reported previously (Britten et al. 1996), where animals were confronted with motion displays of various level of motion coherence and fixed viewing duration. The increased neuronal activity seen by Britten et al. could have been caused by selective attention to a particular direction of motion. Even at 0% coherence individual dots have motion energy in the attended direction, and attention can selectively increase motion related activity in area MT (Treue and Maunsell 1996). In our experiments animals indicated their choice in complete absence of a visual motion stimulus. Thus choice dependent activity in MT in our study cannot be assigned to feature based attention (Treue and Trujillo 1999). We also employed a task where monkeys were free to indicate choices at any time throughout the trial, while Britten et al. (1996) employed a fixed viewing time of 2 s. They reported that choice dependent firing rate differences emerged from early on during the response, it was unaffected as the trial progressed, but was absent during fixation prior to stimulus presentation (Britten et al. 1996). Conversely, we found an increase during the fixation period (the entire SID trial period was ‘fixation only’), but the increase was restricted to a short period from ~500 to 100 ms before the animal’s choice. Thus, the presence or absence and the duration of these increases appear strongly coupled to task timing and structure. Whether the increased neuronal activity is generated within MT or inherited from other areas is currently unknown. It could occur in direction selective neurons in V1, V2, or V3, whereby selective pooling in area MT possibly amplified the signal. Alternatively, attention signals operating in the absence of visual stimuli could produce the increases. Animals probably knew that early in a trial there was a greater chance of stimulus appearance and as time progressed the residual stimulus probability decreased. In response, monkeys could have been more focused on activating MT through attentional feedback mechanisms early in the trials and less so later in the trials. This would require highly selective feedback, whereby a

specific directional neuronal subpopulation was activated more strongly than subpopulations coding for other directions of motion. A mechanism, similar to the biased competition model of attention (Desimone 1998) could further increase the differences in activity between different populations of direction selective neurons. These specific directionally selective feedback signals could arise, e.g. in area LIP (Kusunoki et al. 2000; Roitman and Shadlen 2002; Bisley and Goldberg 2003; Bisley and Goldberg 2006; Saalmann et al. 2007) or area VIP (Cook and Maunsell 2002). In addition to its high selectivity such a feedback signal would also have to activate MT neurons within a short time frame, since the activity increase was restricted to ~500–100 ms before the choice. Additional experiments are necessary to determine whether the selectively increased activity in area MT is generated in MT itself, whether it is generated at even earlier stages (e.g. V1, V2, V3), or whether it is due to feedback from higher areas. We currently favour the hypothesis that the signal is generated in MT or earlier areas, rather than by selective feedback, due to the temporal relation (and precision) of its occurrence and the ensuing choice.

We failed to find a significant increase in MT neuronal activity before choices in preferred direction associated with late SIDs. Negative results are difficult to interpret, as they do not allow for the conclusion that there was no difference in activity between late SIDs in preferred and anti-preferred direction. However, our results demonstrate conclusively that the increase with choices in preferred direction during early SID was significantly larger than the activity level with choices in preferred direction during late SIDs. It could be argued that some of the late SIDs were delayed early SIDs, i.e. the animal made a decision at some point early in the trial due to a briefly increased activity in the population that codes for the direction of choice but delayed its indication. Related activity increases would then get blurred for these particular late SIDs and this could contribute to the fact that we do not see increased choice probabilities for late SIDs. However, the increased reward exploitation and different choice allocation strategy during late SIDs argues against the idea that late SIDs are simply delayed early SIDs.

In addition to the selective increase of firing rates we also found a dependence of precise spike time synchrony with directional choices, but these were only systematically related to the direction and timing of choice for neurons with similar preferred direction provided they had some receptive field overlap. A pattern opposite to that expected was found for cells with similar preferred direction but relatively little or no RF overlap. We currently have no intuitive explanation for the latter finding. A previous study failed to find synchrony differences as a function of directional choices in area MT (Bair et al. 2001). However,

they used a measure that could confound precise spike time correlation and rate correlation between neurons (Roelfsema et al. 2004). We will therefore discuss their finding in the section below where we discuss our data on rate correlation. In line with previous results we found that strength of precise spike timing synchrony depended on similarity of preferred direction (de Oliveira et al. 1997). Interestingly, we found no relation to the degree of RF overlap, which relates to the distance of recording sites.

A measure of neuronal cooperation that correlated systematically with the time and direction of choice was firing rate correlation. We found selectively increased firing rate correlation in neuronal pairs with similar vector average preferred direction (within 60°), but not (or less) in neuronal pairs with more dissimilar vector average preferred directions (60°–120°). A previous study failed to find fine firing rate correlation differences as a function of directional choices in area MT (Bair et al. 2001). However, important differences exist between our and that report. We found that firing rate correlation was significantly increased during a limited time window before choices in preferred direction of the neurons. Bair et al. (2001) used a fixed viewing time (2 s), and the animals had to indicate their choice thereafter. The animal's decision may have occurred during a short time period within these 2 s, and it is impossible to know when this was. Since we show that the increased firing rate correlation only occurs within a short time period prior to the choice, averaging over 2 s in Bair et al.'s study may have concealed such selective differences. Whether the increase in firing rate correlation is generated locally, by recurrent connections between neurons that share the same preferred direction, or whether it is due to common input from lower or higher areas is impossible to decide from our data. It has been argued that an increase in choice probability in MT neurons is due to common input (Shadlen et al. 1996; Dodd et al. 2001). If true, then increased common input could result in choice dependent firing rate correlation between MT neurons.

Importantly, we did not find an increase in MT activity or firing rate correlation during decisions after the hazard function had maxed out. These late decisions were associated with a significantly increased predicted reward yield and a significantly altered SID allocation strategy, which suggests that they were more governed by memories of previous choices and reward prediction or at least by the knowledge that stimulus appearance is no longer possible. Thus, decisions triggered by an attempt to maximize reward income in this direction discrimination task were not calculated in MT nor did they influence MT activity and firing rate correlation systematically via feedback connections. Previous studies have demonstrated that signals injected into area MT have little effect on an animal's

behaviour when MT activity is behaviourally irrelevant (Seidemann et al. 1998). This finding has been interpreted as evidence for active gating of the flow of sensory information, a mechanism by which irrelevant information is filtered out such that it does not interfere with behaviour. Our data show that these signals do not even arise in MT during such time periods.

Choices based on memory and reward prediction are almost certainly calculated in higher areas of the parietal and prefrontal cortex, in conjunction with the basal ganglia and midbrain dopaminergic neurons (Platt and Glimcher 1999; Schultz et al. 2000; Barraclough et al. 2004; Dorris and Glimcher 2004; Tsujimoto and Sawaguchi 2005; Ding and Hikosaka 2006; Ichihara-Takeda and Funahashi 2006; Morris et al. 2006). If the increased activity prior to early SIDs in preferred direction was due to feedback signals from higher areas, rather than generated within sensory areas themselves, then these feedback signals must originate in a neuronal pool, which is different from the pool that calculates the choice direction for the late choices. Otherwise we would expect to see the same activity differences for early and late SIDs. Alternatively a mechanism must exist by which feedback can be enabled and disabled at the level of the area which receives the feedback, and such enabling would depend on the behavioural state and the task constraints.

In summary, our results demonstrate that the relative weight given to a particular source of information varies continuously depending on task conditions and is properly gated within the cerebral cortex on short time scales that directly match task constraints.

Acknowledgments We would like to thank J. Read, P. Dayan, and N. Daw for valuable help, feedback, and discussions. Supported by Deutsche-Forschungs-Gemeinschaft, Neurovision, SFB 509. AT was supported by the Wellcome Trust.

References

- Aertsen AM, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J Neurophysiol* 61:900–917
- Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci* 21:1676–1697
- Barraclough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 7:404–410
- Bisley JW, Goldberg ME (2003) Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299:81–86
- Bisley JW, Goldberg ME (2006) Neural correlates of attention and distractibility in the lateral intraparietal area. *J Neurophysiol* 95:1696–1717
- Bradley DC, Chang GC, Andersen RA (1998) Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature* 392:714–717

- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765
- Britten KH, Newsome WT, Shadlen MN, Celebriani S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100
- Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* 312:1622–1627
- Cook EP, Maunsell JH (2002) Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci* 5:985–994
- de Oliveira SC, Thiele A, Hoffmann KP (1997) Synchronization of neuronal activity during stimulus expectation in a direction discrimination task. *J Neurosci* 17:9248–9260
- Desimone R (1998) Visual attention mediated by biased competition in extrastriate visual cortex. *Philos Trans R Soc Lond B Biol Sci* 353:1245–1255
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222
- Ding L, Hikosaka O (2006) Comparison of reward modulation in the frontal eye field and caudate of the macaque. *J Neurosci* 26:6695–6703
- Dodd JV, Krug K, Cumming BG, Parker AJ (2001) Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J Neurosci* 21:4809–4821
- Dorris MC, Glimcher PW (2004) Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44:365–378
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563
- Hoffmann KP, von Seelen W (1980) Performance in the cat's visual system: a behavioral and neurophysiological analysis. *Behav Brain Res* 12:101–120
- Ichihara-Takeda S, Funahashi S (2006) Reward-period activity in primate dorsolateral prefrontal and orbitofrontal neurons is affected by reward schedules. *J Cogn Neurosci* 18:212–226
- Konishi S, Nakajima K, Uchida I, Kameyama M, Nakahara K, Sekihara K, Miyashita Y (1998) Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nat Neurosci* 1:80–84
- Kreiter AK, Singer W (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J Neurosci* 16:2381–2396
- Krug K (2004) A common neuronal code for perceptual processes in visual cortex? Comparing choice and attentional correlates in V5/MT. *Philos Trans R Soc Lond B Biol Sci* 359:929–941
- Kusunoki M, Gottlieb J, Goldberg ME (2000) The lateral intraparietal area as a salience map: the representation of abrupt onset, stimulus motion, and task relevance. *Vision Res* 40:1459–1468
- Mansouri FA, Matsumoto K, Tanaka K (2006) Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin Card Sorting Test analog. *J Neurosci* 26:2745–2756
- Miller EK, Li L, Desimone R (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13:1460–1478
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167
- Morris G, Nevet A, Arkadir D, Vaadia E, Bergman H (2006) Midbrain dopamine neurons encode decisions for future action. *Nat Neurosci* 9:1057–1063
- Nakahara K, Hayashi T, Konishi S, Miyashita Y (2002) Functional MRI of macaque monkeys performing a cognitive set-shifting task. *Science* 295:1532–1536
- Newsome WT, Britten KH, Salzman CD, Movshon JA (1990) Neuronal mechanisms of motion perception. *Cold Spring Harb Symp Quant Biol* 55:697–705
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238
- Rainer G, Asaad WF, Miller EK (1998) Memory fields of neurons in the primate prefrontal cortex. *Proc Natl Acad Sci USA* 95:15008–15013
- Roberts M, Delicato LS, Herrero J, Gieselmann MA, Thiele A (2007) Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner. *Nat Neurosci* 10:1483–1491
- Roelfsema PR, Lamme VA, Spekreijse H (1998) Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395:376–381
- Roelfsema PR, Lamme VA, Spekreijse H (2004) Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nat Neurosci* 7:982–991
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489
- Saalman YB, Pigarev IN, Vidyasagar TR (2007) Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316:1612–1615
- Salzman CD, Britten KH, Newsome WT (1990) Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346:174–177
- Schultz W, Tremblay L, Hollerman JR (2000) Reward processing in primate orbitofrontal cortex and basal ganglia. *Cereb Cortex* 10:272–284
- Seidemann E, Zohary E, Newsome WT (1998) Temporal gating of neural signals during performance of a visual discrimination task. *Nature* 394:72–75
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510
- Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations. *Neuron* 24:49–65
- Smith P (1995) Psychophysically principled models of visual simple reaction time. *Psychol Rev* 102:567–593
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304:1782–1787
- Thiele A, Hoffmann KP (1996) Neuronal activity in MST and STPp, but not MT changes systematically with stimulus-independent decisions. *Neuroreport* 7:971–976
- Thiele A, Distler C, Hoffmann KP (1999) Decision-related activity in the macaque dorsal visual pathway. *Eur J Neurosci* 11:2044–2058
- Thiele A, Dobkins KR, Albright TD (2001) Neural correlates of chromatic motion perception. *Neuron* 32:351–358
- Treue S, Maunsell JHR (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–541
- Treue S, Trujillo JCM (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579
- Tsujimoto S, Sawaguchi T (2005) Neuronal activity representing temporal prediction of reward in the primate prefrontal cortex. *J Neurophysiol* 93:3687–3692
- von Seelen W, Hoffmann KP (1976) Analysis of neuronal networks in the visual system of the cat using statistical signals. Part I Biol Cybern 22:7–20

- Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. *J Neurophysiol* 90:1790–1806
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080