

ORIGINAL ARTICLE

Learning a New Selection Rule in Visual and Frontal Cortex

Chris van der Togt¹, Liviu Stănișor¹, Arezoo Pooresmaeili¹, Larissa Albantakis², Gustavo Deco³, and Pieter R. Roelfsema^{1,4,5}

¹Department of Vision and Cognition, Netherlands Institute for Neuroscience, An Institute of the Royal Netherlands Academy of Arts and Sciences, 1105 BA Amsterdam, The Netherlands, ²Madison School of Medicine, Department of Psychiatry, University of Wisconsin, 6001 Research Park Boulevard, Madison, WI 53719, USA, ³Dept. de Tecnologies de la Informació i les Comunicacions, Universitat Pompeu Fabra, C\ Tanger, 122-140, 08018 Barcelona, Spain, ⁴Department of Integrative Neurophysiology, Centre for Neurogenomics and Cognitive Research, VU University Amsterdam, Amsterdam, The Netherlands and ⁵Psychiatry Department, Academic Medical Center, 1105 AZ Amsterdam, The Netherlands

Address correspondence to email: p.roelfsema@nin.knaw.nl

Chris van der Togt and Liviu Stănișor contributed equally to this work.

Abstract

How do you make a decision if you do not know the rules of the game? Models of sensory decision-making suggest that choices are slow if evidence is weak, but they may only apply if the subject knows the task rules. Here, we asked how the learning of a new rule influences neuronal activity in the visual (area V1) and frontal cortex (area FEF) of monkeys. We devised a new icon-selection task. On each day, the monkeys saw 2 new icons (small pictures) and learned which one was relevant. We rewarded eye movements to a saccade target connected to the relevant icon with a curve. Neurons in visual and frontal cortex coded the monkey's choice, because the representation of the selected curve was enhanced. Learning delayed the neuronal selection signals and we uncovered the cause of this delay in V1, where learning to select the relevant icon caused an early suppression of surrounding image elements. These results demonstrate that the learning of a new rule causes a transition from fast and random decisions to a more considerate strategy that takes additional time and they reveal the contribution of visual and frontal cortex to the learning process.

Key words: frontal eye fields, learning, primary visual cortex, visual attention, visual routines

Introduction

Imagine that you want to learn a new game. There are multiple ways to learn the rules. For example, you can choose to study the manual. However, in many cases, the best way to learn is to try to optimize your strategy while playing, making many errors at the start. Virtually in all games, you will have to learn to attend

features that matter, for example, the shape of the chess pieces or the symbols on cards, and to ignore the rest. However, how do you distribute your attention and make decisions when you do not yet know the rules?

In previous work, researchers have gained insight into the neuronal mechanisms underlying sensory decisions. Many studies have focused on the parietal and motor cortex in tasks

where monkeys choose to move their eyes to one of a number of positions (Schall 1991; Schall et al. 1995; Hanes and Schall 1996; Kim and Shadlen 1999; Gold and Shadlen 2001, 2007; Murthy et al. 2001; Ding and Gold 2012). These studies revealed that eye movement decisions are implemented as a race or a competition between pools of neurons that code for different eye movement decisions. When the sensory evidence in favor of one of the eye movements is strong, neurons that code for this eye movement quickly ramp up their activity, but the buildup of activity is more gradual and reaction times are prolonged if the evidence is weak. In these previous tasks, the monkeys were familiar with the rule so that it was strategic to wait for more evidence and to postpone the response when the stimulus was weak. The optimal strategy may be different when the stimulus is clearly visible, but the rules are unknown. How is a decision made in this situation? Do monkeys also postpone their decision in the absence of evidence for one or the other decision?

In the present study, we were interested in the changes in the representations in the visual and frontal cortex when monkeys learn a new rule. The role of the frontal cortex in decision-making has been well established (Schall 1991; Schall et al. 1995; Hanes and Schall 1996; Kim and Shadlen 1999; Murthy et al. 2001; Ding and Gold 2012), but we here also investigated the influence of decision-making on neuronal activity in visual cortex because of the intimate relationship between decision-making and shifts of attention. The decision to make an eye movement is invariably associated with a shift of attention to the location of the eye movement target (Kowler et al. 1995; Deubel and Schneider 1996). Furthermore, learning an unknown rule implies that one learns to attend to the relevant features and to ignore irrelevant ones. Learning, therefore, has a pronounced influence on the distribution of attention. Specifically, visual objects that have been associated with a high reward in previous trials will usually attract attention in later trials (Della Libera and Chelazzi 2009; Raymond and O'Brien 2009; Hickey et al. 2010; Anderson et al. 2011; Chelazzi et al. 2013). Also at a neuronal level, the representations of stimuli that have been associated with high rewards are enhanced in visual, parietal, and frontal cortex. Their representations resemble the representation of attended stimuli (Kim and Shadlen 1999; Maunsell 2004; Serences 2008; Peck et al. 2009; Stănişor et al. 2013). Because the present study does not aim to dissociate the influence of attention from the influence of eye movement selection on neuronal activity, we will use the more neutral term “selection signal” to describe the effects of stimulus selection on activity in visual and frontal cortex.

Leveraging on earlier work that examined how monkeys learn arbitrary associations between stimuli and responses (Chen and Wise 1995a, 1995b, 1996; Asaad et al. 1998; Tremblay et al. 1998; Erickson and Desimone 1999; Jagadeesh et al. 2001; Wirth et al. 2003; Brasted and Wise 2004; Pasupathy and Miller 2005), we devised a new icon-selection task that allowed us to investigate selection signals related to motor planning in the frontal cortex and to shifts of attention in the visual cortex. Specifically, we presented a new pair of shapes (icons) to the monkeys in every session that were connected by a curve to an eye movement target (Fig. 1). The monkeys had to learn which of these icons was rewarded with juice and which one was not. This design provided a number of strengths. First, we could expose the monkeys to a new learning problem in every recording session. Second, we could place the curves in the receptive fields of V1 neurons to examine attentional selection signals while keeping the receptive field stimulus constant. Third, we could place the saccade targets in the receptive field of FEF neurons, again ensuring that the receptive field stimulation was identical across

conditions. Finally, we could vary the difficulty of the task to investigate how this factor influences the activity in V1 and FEF.

With this design, we aimed to address the following questions. 1) How are eye movements selected if the rule is unknown? 2) Are random guesses associated with selection signals in the visual and frontal cortex? 3) What is the influence of learning on the time course of selection signals in the visual and frontal cortex?

When the monkeys learned the new rule, we observed a transition from extremely fast but random decisions at the start of a session to slower and accurate decisions as the monkeys became proficient. The longer decision times were associated with a more gradual buildup of activity in the frontal cortex and with the later emergence of selection signals in the visual cortex. Our V1 recordings revealed a cause of the delay, because the learning induced a brief suppression of activity in the vicinity of the relevant icon allotting additional time for more considerate decisions in frontal cortex.

Materials and Methods

Three monkeys (A, J, and G) participated in this study. We recorded 2 datasets, one in FEF of monkeys A and J, and the other one in V1 of monkeys A and G. In a first operation, we implanted a head holder to stabilize the head and inserted a gold ring under the conjunctiva of one eye for the measurement of eye position. For the FEF recordings, we performed a separate surgery to make a trepanation over area FEF and to place a recording chamber. Before surgery, the FEF was localized with a magnetic resonance imaging scan and once the chamber was in place we confirmed its location by eliciting saccadic eye movements with microstimulation (generally $<50 \mu\text{A}$). For the V1 recordings, we chronically implanted arrays of 4×5 or 5×5 electrodes (Blackrock Microsystems). The surgical procedures were performed under aseptic conditions and general anesthesia. Details of the surgical procedures and the postoperative care have been described previously (Roelfsema et al. 1998; Khayat et al. 2009). All procedures complied with the US National Institutes of Health Guidelines for the Care and Use of Laboratory Animals were approved by the institutional animal care and use committee of the Royal Netherlands Academy of Arts and Sciences of the Netherlands.

The stimuli were presented on a monitor with a diagonal of 35.5 cm (14 inch), a resolution of 1024 by 768 pixels, and a refresh rate of 100 Hz at a distance of 75 cm from the monkey's eyes. The objects that appeared on the screen were colorful figures with a size of approximately 1.0° ; we will refer to these as “icons” (Fig. 1). The saccade targets were red disks with a diameter of 0.8° connected to the icons by a curve that was 2 pixels wide (0.05°). The eye position was measured using either the double magnetic induction technique built in house (Bour et al. 1984) (sampling rate 1 kHz) or an infrared camera system (Thomas Recording; ET-49B, sampling rate 250 Hz).

Behavioral Task

The animals performed a forced choice task where they had to select one of 2 icons and then made an eye movement to a circular disk that was connected to this icon by a curve (Fig. 1A). We presented a new (unfamiliar) pair of icons during every recording session and the monkey had to learn which of these 2 icons was associated with reward. A trial started as soon as the monkey's eye position was within a $3^\circ \times 3^\circ$ square window centered on a red fixation point (FP) for recordings in FEF or a $1.2^\circ \times 1.2^\circ$ window for recordings in V1. After an interval of 300 ms, the 2 icons as well as 2 curves and disks appeared on the screen (Fig. 1). We

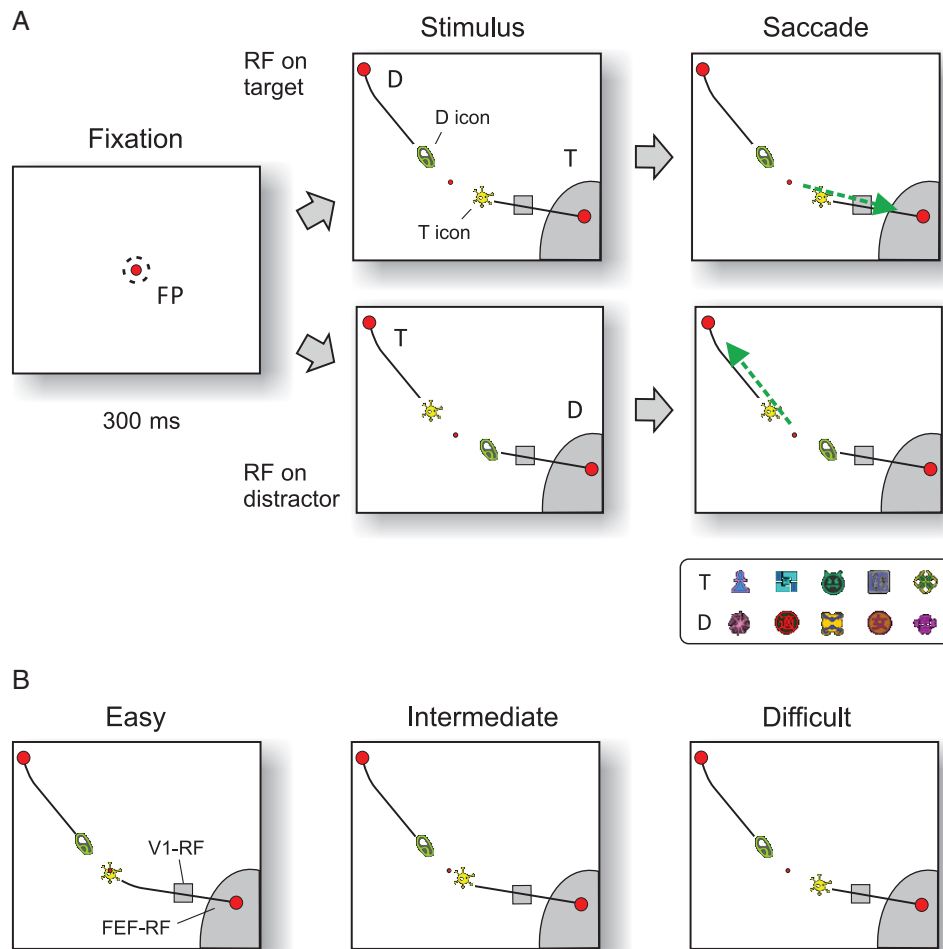


Figure 1. Icon-selection task. (A) After the monkey fixated a central fixation point for 300 ms, 2 icons appeared that were connected by curves (black lines) to 2 red disks. The monkey had to make a saccade (green arrow) to the target disk (T) that was connected to the target icon (yellow icon). The other icon, curve, and disk were distractors (D). The lower panel shows examples of target and distractor icons used in different sessions. (B) For easy and intermediate stimuli, the target icon was on top of or closer to the fixation point than the distractor icon. In the difficult condition, the distance between the icons and the fixation point was identical so that the monkeys could only rely on icon identity to make a correct choice. RF, receptive field; FP, fixation point.

presented 3 types of stimuli: easy, intermediate, and difficult. For the easy condition, the target icon was placed on top of the fixation point, so that the animal could trace this curve to identify the correct target for an eye movement (Fig. 1B). These trials were easy, because the monkeys had previously learned to trace a curve connected to the fixation point (without icons). In the intermediate condition, the target icon was closer to the fixation point than the distractor icon. In the difficult condition, the relevant and irrelevant icons were at the same distance from the FP (with a size of 0.3°) (Fig. 1B). In this condition, the monkey had to use the identity of the icons for correct performance. The easy, intermediate, and difficult trials were randomly interleaved. Note that the identity of the relevant and irrelevant icon was the same across difficulty levels. However, care was taken to use distinct icons (with distinct combinations of shapes and colors) in different sessions (examples are shown in the inset in Fig. 1A).

Before the 3 monkeys entered into the icon-selection task, they had ample experience with a curve-tracing task in which the target curve was connected to the fixation point. Prior to data collection, we specifically trained them to consider icon identity by using the same target and distractor icon across a number of successive sessions while varying the fraction of trials

of the easy, intermediate, and difficult condition. These early training sessions where icon identity was stable across days were not included in the present dataset.

During the recording sessions, the stimulus in the receptive field (RF) of FEF and V1 neurons was kept constant across conditions. We ensured that the V1 RFs (with a size of $\sim 1^\circ$ and mapped as described below) were activated only by the curves (small rectangles in Fig. 1) and the FEF RFs by curves and disks (larger gray regions in Fig. 1). Because FEF RFs are large (sometimes $>10^\circ$), the 2 disks were often positioned in opposite hemifields for the FEF recordings, whereas they usually were in the same hemifield for the V1 recordings. In other respects, the stimuli for the FEF and V1 recordings were comparable.

During the FEF recordings, the monkeys were allowed to initiate their response as soon as the stimulus appeared (reaction-time task). We imposed an additional fixation delay of 500 ms after stimulus onset for the V1 recordings before the monkey could make an eye movement to one of the disks, to have a longer time window to evaluate the influence of target selection in V1. In this task, the fixation point changed color from red to green as a cue to initiate a saccade. We rewarded correct choices with a drop of apple juice (~ 0.2 mL).

Analysis of Behavior

To investigate the progression of learning, we assigned a value of 0 to error trials and 1 to correct trials and then averaged accuracy across days while smoothing with a boxcar of 15 trials. The resulting accuracy values ranged from 0.5 (chance level) to 1 (100% correct). We also averaged the reaction time across sessions, only including the correct trials. We found that the animals' accuracy tended to decline toward the end of the sessions, and therefore, we excluded these trials (~20% of the trials at the end of the session). We considered that icon learning was successful when the accuracy reached a level higher than 75% in the difficult condition. For the FEF recordings, the animals performed on average 550 valid trials per session (~180 trials for each difficulty level); sessions with less than a total of 300 trials (e.g., due to loss of isolation of the unit) were not included in the analysis. For the V1 recordings, there were on average 700 trials per session.

Single Unit Recording and Data Analysis in FEF

Extracellular recordings from single neurons were obtained with tungsten microelectrodes, which were lowered through the dura mater with a hydraulic microdrive (Narishige). Action potentials were amplified, filtered, discriminated, and recorded on-line using spike-sorter software (Tucker-Davis Technologies). We confirmed that single units were in FEF by using the recording electrode for intracortical microstimulation (biphasic current pulses, 70 ms train duration, 400 Hz). We judged the penetration to be in FEF if we could trigger a saccade with a current <100 μ A (generally <50 μ A) (Bruce and Goldberg 1985).

The activity of 34 FEF single units was recorded from the 2 monkeys (10 from monkey J and 24 from monkey A). We first mapped the response field (RF) of the neuron by presenting a single saccade target at various eccentricities and directions. We then recorded separate blocks of trials with mapping tasks for eccentricity and radial tuning, to obtain a quantitative measure of the location and extent of the RF. Specifically, we fitted a Gaussian to the radial tuning profile and a spline to the eccentricity profile and assumed that the effects of eccentricity and direction on a neuron's response are separable, that is, the overall response can be described as a multiplication of the effect of these 2 factors (Khayat et al. 2009). We also employed a memory-guided saccade task to classify the cells according to the scheme of Bruce and Goldberg (1985). We classified the 34 cells included in this study as 7 visual neurons and 27 visuomotor neurons (we excluded motor neurons). In the learning task, one of the 2 target disks was placed within the neuron's RF and the other disk at an angle of 90° or 135° and at a similar eccentricity (ranging from 8° to 18°) (Fig. 1). Because some FEF neurons exhibit tuning to visual stimuli (Bichot et al. 1996; Peng et al. 2008), we placed the icons well outside the RFs, ensuring that differences in activity between conditions were caused by target selection.

Peristimulus and perisaccade time histograms of single-unit responses were computed with 10 ms bins and aligned to stimulus appearance and saccade onset, respectively. Only correct trials were included in our analyses, unless otherwise specified. Target selection activity of the FEF neurons was defined as the difference between activity elicited by the target and distracter curve/disk in spikes per second. This difference is expected to be zero at response onset followed by an increase or decrease reflecting the effect of target selection. On the basis of these assumptions, we fitted a piecewise (100 ms intervals) cubic spline using least square fitting (de Boor 1978) that was constrained to

be zero at onset. Fits were obtained using the shape prescriptive modeling toolbox implemented in MATLAB (Mathworks 2009 by John D'Ericco) and used to estimate the latency of the selection signal as the time point at which the fitted function reached 33% of its maximum. We used a bootstrap analysis to measure the significance of differences in the latency of the selection signal between conditions. We randomly selected the average responses in one or the other condition for each cell/MUA recording site and then averaged across all cells. We used the fitting procedure to compute latencies in 10 000 simulated samples and considered the real latency differences to be significant if they fell outside 95% of this simulated distribution.

Recording and Data Analysis of Multiunit Activity in V1

We recorded multiunit activity (MUA) from 28 recording sites simultaneously in 19 sessions in monkey A and from 7 recording sites simultaneously in 22 sessions in monkey G. MUA was recorded from chronically implanted multielectrode arrays (Blackrock Inc.) with TDT multichannel recording equipment. As in previous studies, the MUA signal was amplified, bandpass filtered (300–9000 Hz), full-wave rectified, low-pass filtered (<200 Hz), and sampled at a rate of 763 Hz. The MUA represents the pooled activity of a number of single units in the vicinity of the tip of the electrode (Logothetis et al. 2001; Supér and Roelfsema 2005; Cohen and Maunsell 2009; Xing et al. 2009). The population response obtained with this method is therefore expected to be identical to the population response obtained by pooling across single units and a previous study demonstrated that MUA indeed provides a reliable estimate of the average single-unit response (Supér and Roelfsema 2005).

We calculated the signal-to-noise ratio of the visual response in single trials at every recording site by dividing the average peak response (after smoothing with a Gaussian kernel with a σ of 10 ms) by the standard deviation of the spontaneous activity in a 200-ms window before stimulus onset. We only included recording sites with a signal-to-noise ratio larger than 1.5 (the mean was 2.5), so that the visual response was discernable in almost every trial. We estimated the RF dimensions of neurons at the recording site by measuring the onset and offset of the response to a slowly moving white bar on a black background, in each of 8 directions (Supér and Roelfsema 2005). The median area of the RFs was 0.52 deg² (ranging from 0.1 deg² to 2.4 deg²). RF eccentricity ranged from 1.2° to 5.0° with an average of 3.1°. V1 RFs were selected such that they all fell on one of the curves and never overlapped with the disks or with the icons.

We computed peristimulus time histograms in a time window ranging from 100 ms before stimulus onset and up to 800 ms afterward and smoothed the activity with a Gaussian kernel (SD 10 ms). To compute population responses, we normalized activity at individual recording sites to the peak response after subtraction of the spontaneous activity level (average activity in the last 100 ms before stimulus onset). The peak response was estimated as the highest activity in an interval between 0 and 150 ms after the stimulus onset. Target selection activity (attentional modulation) was defined as the difference between responses evoked by target and distracter curve, and the onset of the selection signal was estimated with piecewise cubic spline fitting as the time point when the fitted function reached 33% of its maximum, as described for the FEF data above. We measured neuronal activity at the same electrodes of the arrays across multiple days. For every recording site, we therefore first averaged the neuronal activity across sessions (only including sessions with a successful recording from that site) before

entering it into the dataset. Thus, every recording site contributed only a single data point to the statistics.

Results

We investigated the time course of learning of an icon-selection task and the expression of this learning process in FEF and V1. In every session, the monkeys learned to select one of 2 unfamiliar icons. These icons were connected to red disks by a curve and the monkeys indicated their choice by making an eye movement to one of the disks (Fig. 1). Eye movements to the disk connected to the target icon were rewarded, but eye movements to the other disk that was connected to the distractor icon were not. To facilitate learning, we used stimuli with 3 levels of difficulty. In easy trials, the target icon overlapped with the fixation point, and the monkey could simply trace the target curve to its other end to identify the target disk. In the intermediate condition, the target icon was closer to the fixation point than the distractor icon. These 2 conditions could be solved based on the distance between the icons and the fixation point. In the difficult condition, the target and distractor icons were at the same distance from the fixation point, and the monkey could only rely on icon identity for a correct response. There were 2 complementary stimuli in each condition, one for which the RF fell on the target curve and one where the RF fell on the distractor (Fig. 1) resulting in a total of 6 stimuli, which were randomly interleaved.

Reaction Time Task

In the first experiment, we recorded single units in FEF in 2 monkeys (A and J). These animals performed a reaction time version of the icon-learning task and could make an eye movement immediately after stimulus presentation. We analyzed the behavioral performance data from a total of 104 sessions of monkey A and from 61 sessions of monkey J; FEF neurons were isolated in a subset of these sessions. The monkeys made hardly any errors with the easy stimuli (Fig. 2, blue traces). Accuracy in the intermediate condition also started above chance level (at 70–80%) and increased to almost 100% at the end of each session (Fig. 2, green traces). The performance improvement for the difficult stimuli was more gradual, in accordance with studies using related tasks (Asaad and Eskandar 2011; Heilbronner and Platt 2013). In the difficult condition, accuracy started at chance level and increased to values higher than 80% in both animals (Fig. 2, red traces). To analyze the significance of the improvements in performance, we calculated the accuracy in the first 30 (early) trials, and when learning had occurred, we chose trials 101 to 130 for comparison. Accuracy was significantly better in the later trials in the intermediate ($\chi^2 = 52$, $df = 1$, $P < 10^{-6}$) and difficult condition ($\chi^2 = 150$, $df = 1$, $P < 10^{-6}$). The improvement in performance was accompanied by an unexpected increase in reaction time in all 3 conditions. We used a 2-way ANOVA to test the effects of learning (first vs. later trials) and difficulty level (easy/intermediate/difficult) on reaction time. The reaction time increased from an average of 213 ms in the early trials to 238 ms in the later trials in monkey A ($F_{1,11921} = 1021$, $P < 10^{-6}$) and from 284 to 303 ms in monkey J ($F_{1,10614} = 200$, $P < 10^{-6}$). In addition, there was also a significant effect of task difficulty on reaction time (monkey A: $F_{2,11920} = 31.8$, $P < 10^{-6}$; Monkey J: $F_{2,10613} = 12.3$, $P < 10^{-6}$). In both monkeys, reaction times for the difficult task were longer than those in the other 2 conditions. In monkey J, the reaction time in the easy condition was shortest, but in monkey A the reaction time was shortest in the intermediate condition.

In many tasks, humans and monkeys trade off accuracy against speed to optimize reward income (Hanks et al. 2011; Dru-gowitsch et al. 2012; Heitz and Schall 2012). In the beginning of the task, the monkeys did not know the identity of the target icon and postponing their decision in the difficult condition could therefore not increase their accuracy. However, the trade-off between speed and accuracy might have changed by the end of the learning sessions when the monkeys knew which icon was the target. We, therefore, examined the reaction time distribution (Fig. 2C) and, in particular, the relation between reaction time and accuracy for the initial trials and for the late trials (Fig. 2D). For the initial trials, accuracy was around 60%, and as expected, it did not increase with reaction time. In contrast, accuracy in the late trials increased with reaction time from a value of approximately 60% at 200 ms to approximately 90% at 300 ms. Thus, in this phase of the session, it was advantageous to shift the reaction time distribution to higher values. We next computed the expected reward income (in mL juice per s) as function of reaction time, taking into account the relation between accuracy and reaction time as well as the waiting times (intertrial interval plus fixation time before the stimulus appeared).

$$\text{Reward income (RT)} = \frac{\text{Accuracy (RT)} \times 0.2 \text{ mL}}{(\text{waiting time} + \text{RT})}$$

In this function, juice per correct trial was 0.2 mL, and minimal waiting was 0.7 s for monkey A (a 300 ms fixation period plus an intertrial interval of 400 ms) and 0.6 s for monkey J (with an inter-trial time of 300 ms) (Fig. 2E). It can be seen that the reward income in later trials increased more by postponing the response (blue curves in Fig. 2E) than in early trials (purple curves). In early trials, the reward income was maximal for reaction times of 210 ms in monkey A and 290 ms in monkey J. The optimal reaction times increased to 260 and 350 ms, respectively, in the later trials, which is similar to the monkeys' increase in reaction time. These results suggest that the increase in reaction time during the learning sessions may have been a strategic adjustment that optimized reward income per unit time (Fig. 2E). It is of interest that we observed similar increases in reaction times in the easy and intermediate conditions (Fig. 2B). This result suggests that adjustments in the speed-accuracy trade-off for optimizing performance in the difficult condition also influenced reaction times for the easier levels.

Target Selection in FEF

We next investigated the neuronal correlates of target selection in FEF. Figure 3A shows the stimuli in one of the recording sessions. The RF of the example FEF neuron fell on the circular disk and on the more eccentric segment of one of the curves so that the stimulus in the RF was always identical. The neuron was activated more strongly when the target disk fell in the RF than if the distractor disk fell in the RF (Fig. 3B), a modulation caused by the selection of the target of an eye movement (Schall 1991; Schall et al. 1995; Murthy et al. 2001; Khayat et al. 2009; Pooresmaeili et al. 2014). The latency of this selection signal (the difference in activity elicited by target and distractor) increased with difficulty, as has been observed previously (Sato et al. 2001; Cohen et al. 2009; Pooresmaeili et al. 2014).

We observed the same effects when we pooled data across all FEF neurons, selecting all correct trials from a session (Fig. 4). Also at the population level, the activity evoked by the target disk was significantly stronger than the activity evoked by the distractor in all difficulty conditions (paired t-test 50–250 ms after stimulus onset, $n = 34$, all 3 P s $< 10^{-5}$) (Fig. 4). To examine

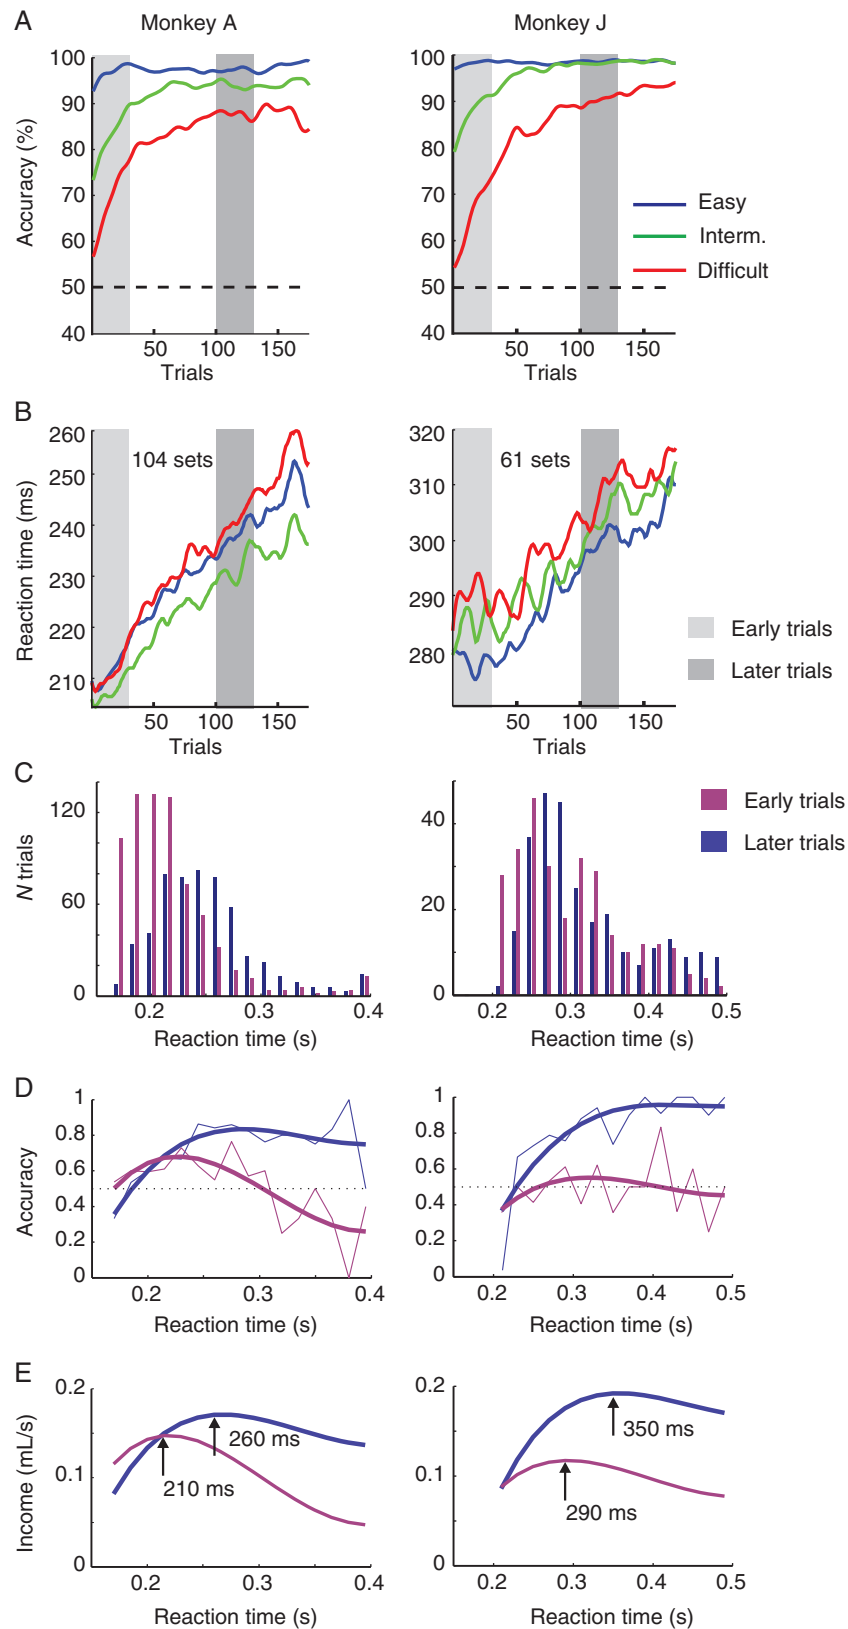


Figure 2. Behavioral performance and reward income in the reaction time task. (A) Accuracy during learning, averaged across sessions. Performance improved in the intermediate (green) and the difficult condition (red), and was always near 100% in the easy condition (blue). (B) Average reaction time, which increased during the sessions. (C) Reaction time distributions in early and late trials for the difficult condition. (D) Accuracy as function of reaction time for early (purple) and late (blue) trials, with a spline fit of the data (smooth thick line). After learning, accuracy increases with reaction time. (E) Reward income (mL juice per second) as a function of reaction time, computed with the following equation: Reward income (RT) = Accuracy (RT) × 0.2 mL/(waiting time + RT). Arrows indicate reaction times that maximize reward income in early and late trials.

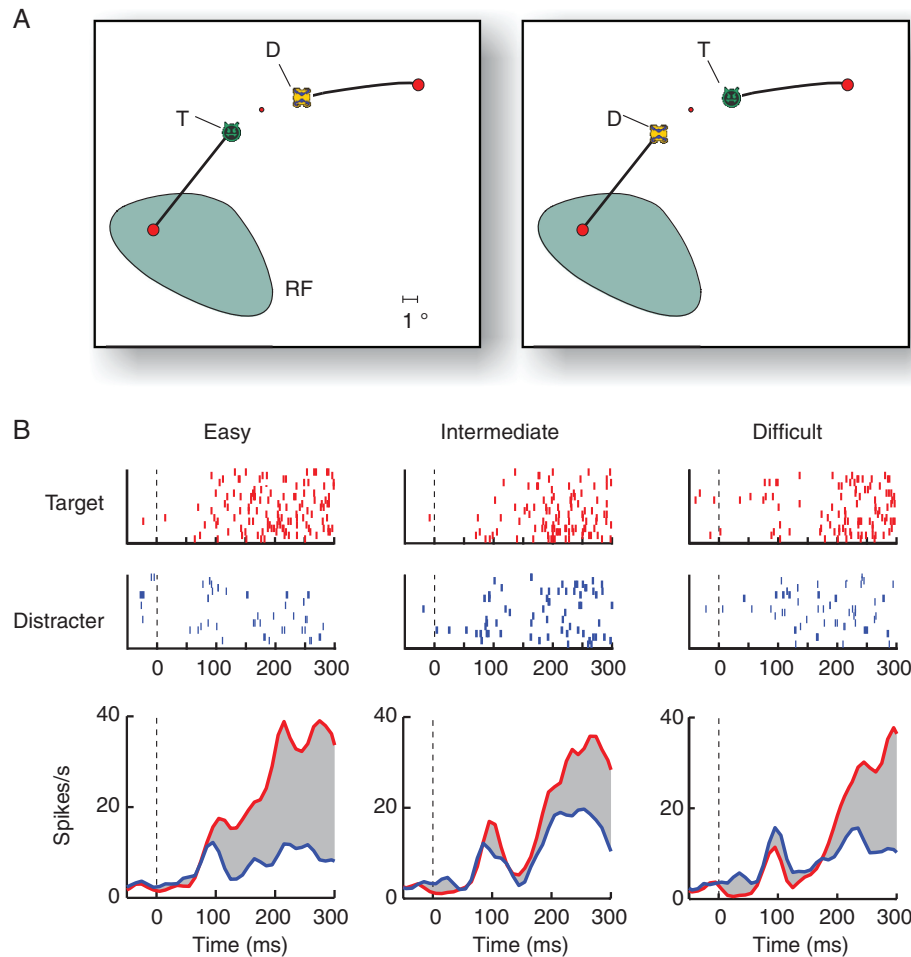


Figure 3. Example FEF neuron in the icon-learning task. (A) Stimuli of the difficult condition. The RF is shown in green. (B) The upper panels show raster plots (ticks are spikes in successive trials in the same condition) in the easy, intermediate, and difficult condition with the target (red) or distractor disk (blue) in the RF. The lower panels show the average responses of this neuron. The gray region illustrates the selection signal, defined as the difference in activity elicited by the target and distractor.

how the onset of the target-selection signal depended on task difficulty, we estimated the latency with a curve-fitting method (see Materials and Methods). Target selection occurred later when the task was more difficult. The latency was 145 ms in the easy condition; it increased to 155 ms in the intermediate condition and to 175 ms in the difficult condition (Fig. 4A) (bootstrap test, $N = 10\,000$, $P < 0.05$ for all 3 comparisons). Interestingly, the amplitude of the response elicited by the target curve reached a similar level across difficulty conditions in a 50-ms window just before saccade onset, but the distractor elicited more activity if the task was difficult (Fig. 4B). We found significant differences between distractor activity in the difficult condition and the 2 other conditions (paired t-test, $n = 34$, both P s < 0.01), but difference between the easy and intermediate conditions was not significant ($P > 0.05$). The present task thereby reproduces previous findings on saccade target selection. Neuronal activity accumulates toward a relatively fixed threshold, which is reached around the time of the saccade. The accumulation process takes more time if the task is difficult (Sato et al. 2001; Gold and Shadlen 2007).

The Effects of Learning in FEF

Our main aim was to investigate how learning influences the activity of FEF neurons. How does the time course of the neuronal

activity change when the monkeys learn which of the 2 icons is the target and which one is the distractor?

To investigate the effects of learning, we compared activity in the difficult condition in an early epoch (the first 15 correct trials with the RF on the target and the first 15 correct trials with the RF on the distractor in the difficult condition) when the monkeys' accuracy was low to that in a later epoch (trials 101–130) when performance had improved (Fig. 2). For both epochs, we then averaged activity elicited by the target or distractor curve across sessions (Fig. 5A). Besides a reduction in the initial visual peak response, which may have been caused by adaptation, we did not find a change in the total amount of activity evoked by target or distractor, nor an increase in the strength of modulation with learning or the peak activity around the time of the saccade ($P > 0.1$, bootstrap test, $N = 10\,000$). However, we did find a difference in the latency of the selection signal. The latency increased significantly from 145 ms at the start of the session to 185 ms in the later epoch (Fig. 5A) ($P < 0.05$, bootstrap test, $N = 10\,000$) for the difficult condition. When we analyzed the data of the 2 monkeys separately, we found that the increase in latency was 40 ms in both of them. The increase in latency with learning was significant in monkey J ($P < 0.05$) but not in monkey A ($P > 0.05$), presumably due to the small sample size in this animal ($n = 10$ single units). The latency of the selection signal did not change

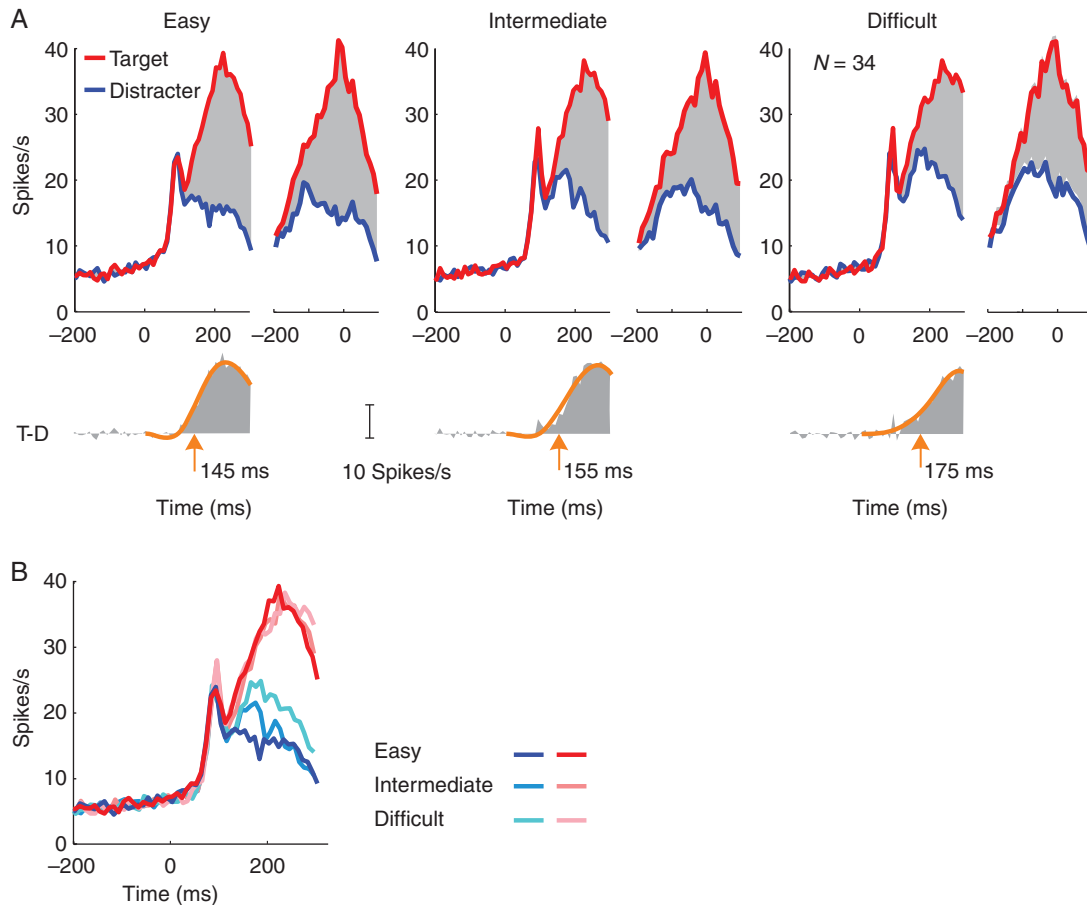


Figure 4. Population activity in FEF in correct trials of the easy, intermediate, and difficult conditions. (A) Average activity across all sessions and both monkeys, elicited by the target (T, red) and distractor (D, blue) and aligned on stimulus onset (left) or saccade (right). Lower insets depict the selection signal (T-D) aligned on stimulus onset. The latency of target discrimination was measured as the time point where a piecewise cubic spline fit (orange curve) reached 33% of its maximum. (B) Direct comparison of activity elicited by the target and distractor at different difficulty levels.

significantly with learning in the easy and intermediate condition (both P s > 0.1).

The early FEF selection signals in the difficult condition at the start of the task occurred during a period when the monkey had no preference for one or the other icon and performed at chance level (Fig. 2). Thus, our results suggest that FEF neurons initially implemented a fast but random selection of one of the disks. This interpretation was confirmed when we analyzed the error trials in the difficult condition. In error trials, the selection signal was inverted, that is, the FEF neurons selected the distractor curve (Fig. 5B), as has been observed previously (Thompson et al. 2005; Pooresmaeili et al. 2014). These results therefore demonstrate that learning was associated with a transition from a fast but random selection of one of the curves to a more accurate, yet delayed selection of the target curve (Fig. 5A). Thus, the monkeys allotted additional time for the accurate selection of the target icon (Fig. 2) and, accordingly, the selection FEF signal occurred at a later point in time.

Delayed Response Task

We next investigated how learning in the icon-learning task changes activity in the visual cortex. We targeted area V1, where the representation of relevant curves and icons is enhanced relative to the representation of distractors (Roelfsema

et al. 2003; Moro et al. 2010). For the V1 recordings, we used a delayed version of the icon-selection task that gave us the opportunity to analyze neural responses in a longer time window. The monkeys had to maintain fixation for 500 ms after stimulus onset before they were cued to make a saccade by a change in the color of the fixation point (Fig. 6). We obtained behavioral data from 22 recording sessions in each monkey (A and G). Figure 7 shows the influence of learning on accuracy and reaction time, averaged across sessions. The increase in accuracy was pronounced in the difficult condition, whereas the monkeys hardly made any errors for the easy and intermediate stimuli. In the difficult condition, monkey A reached an accuracy higher than 75% after approximately 100 trials, comparable to his learning rate in the reaction time task used with the FEF recordings (Fig. 2) and the learning curve of monkey G was somewhat steeper. A comparison of early trials (1–30) and later trials in the difficult condition (trials 121–150 in monkey A; 61–90 in monkey G; gray bars in Fig. 7) revealed that the increase in accuracy was highly significant (monkey A: $\chi^2 = 95.3$, $P < 10^{-6}$; monkey G: $\chi^2 = 60.6$, $P < 10^{-6}$, Chi-Square test).

The reaction time profile differed from the reaction time task (Fig. 2B), because the monkeys were now required to fixate an extra 500 ms. Reaction times in monkey G were often relatively long and decreased slightly during learning, whereas monkey A tended to respond quickly after the cue to make a saccade.

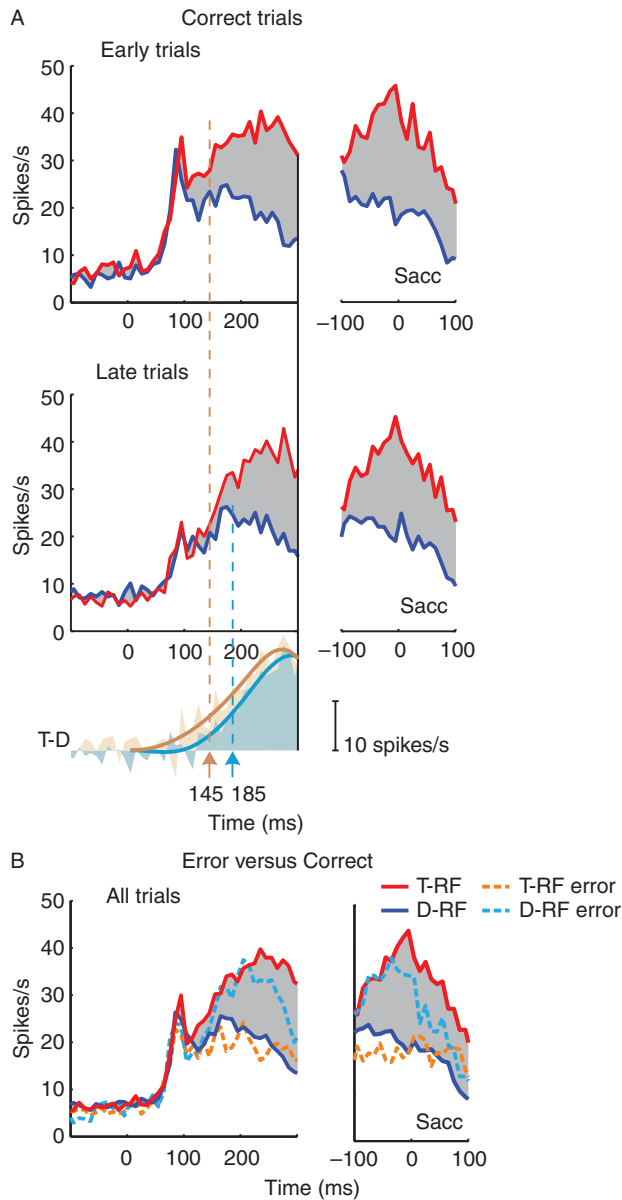


Figure 5. Influence of learning and errors on neuronal activity in FEF. (A) Left panels show stimulus-aligned activity, right panels saccade-aligned activity in early and late trials. The panel below shows the time course of the selection signal (T-D) in the early (orange) and later trials (light blue) aligned on stimulus onset and a curve that was fitted to determine the latency. (B) Comparison of activity in correct (continuous curves) and error trials (dashed) elicited by the target (red) and distractor (blue). The number of error trials decreased with learning so that we could not compare activity between early and late error trials.

A 2-way ANOVA with factors “epoch” (early vs. late) and difficulty (easy, intermediate, difficult) revealed that only the reaction times of monkey G decreased significantly during learning ($F_{1,3262} = 37$, $P < 10^{-6}$). However, in both monkeys, reaction times were longer in the difficult condition than in the other conditions (monkey A, $F_{2,3357} = 30$, $P < 10^{-6}$; monkey G, $F_{2,3261} = 40$, $P < 10^{-6}$).

Target Selection in V1

We always configured the stimuli in such a way that the RFs of the MUA recording sites either fell on the target curve or on the

distracter curve, and the RF stimulus was identical across conditions (Fig. 6). Figure 8A shows the average normalized MUA responses of all recording sites (28 recording sites in monkey A and 7 in monkey G combined), an average that includes all correct trials from the sessions. The initial visual response triggered by the appearance of the curve in the RF was identical for the 2 conditions, but after a delay of 100–200 ms, the activity elicited by the target curve became stronger than that elicited by the distractor in all 3 conditions. Thus, V1 neurons also select the relevant curve.

The temporal profile of the selection signal differed between the difficulty levels. In the easy and intermediate conditions, response modulation started early with a brief period of stronger modulation. The selection signal built up more gradually in the difficult condition and, accordingly, the strength of the response modulation differed between conditions in both monkeys (ANOVA; monkey A, $F_{2,81} = 6.91$, $P < 0.01$; monkey G, $F_{2,18} = 26.6$, $P < 0.01$), whereas the strength of the selection signal became comparable around the time of the saccade. The main difference between difficulty levels was the latency of the selection process. It was 104 ms in the easy condition, 113 ms in the intermediate condition, and 202 ms in the difficult condition. In particular, the latency in the difficult condition was strikingly longer compared with that in other 2 conditions. The differences in latency between conditions were significant for both monkeys (bootstrap test, $N = 10\,000$, $P_s < 0.05$) with exception of the difference between easy and intermediate for monkey G.

The Effects of Learning in V1

To investigate the effect of learning on the V1 selection signal, we compared activity in the early epoch to that in the later epoch (for the definition of these epochs, see Fig. 7). Learning delayed the onset of the selection signal in the difficult condition from 205 to 269 ms ($N = 35$, $P < 0.01$, bootstrap test, $N = 10\,000$). Learning did not delay the onset of the selection signal in the intermediate and easy conditions ($P > 0.05$) (Fig. 8B and 9A). The increased latency in the difficult condition was associated with an early inversion of the selection signal (“+” symbols in Fig. 8B), implying that the activity elicited by the target curve first decreased below the activity elicited by the distractor. This suppression became significantly stronger when the animals became proficient in the difficult condition (paired t-test, $P < 10^{-5}$) (Fig. 9B). Thus, learning caused a brief episode of suppression of the V1 representation of contour elements of the target curve adjacent to the target icon, which the animals learned to select. This suppression was associated with an additional delay in the onset of the selection signal for the target curve.

Error Trials in V1

We hypothesized that the early suppression of the target curve is caused by an initial selection of the target icon, which might be associated with a suppression of the activity of the adjacent elements of the target curve. To investigate this possibility, we analyzed the error trials where the monkeys selected the wrong icon. Error trials were associated with an early suppression of activity elicited by the distractor curve, followed by an increase in the activity evoked by this curve, that is, an inversion of the activity profile in correct trials (Fig. 9C). These results suggest that the selection of one of the icons caused a brief suppression of the adjacent contour elements of the connected curve, which changed into a later enhancement, as has been illustrated in Figure 10. Interestingly, this suppression did not occur in the easy and

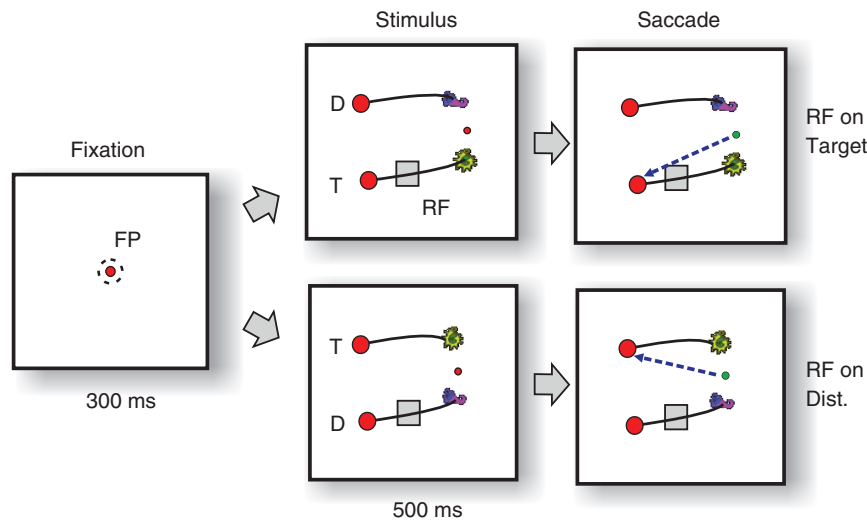


Figure 6. Delayed version of the icon-selection task. The monkeys directed their gaze to a central fixation point (FP) for 300 ms and then the stimulus appeared. Either the target (T) or the distractor curve (D) fell in the RF of a group of neurons in area V1. After an additional delay of 500 ms, the color of the fixation changed, cueing the monkeys to indicate their choice by making a saccade to the disk connected to the target icon (green arrow).

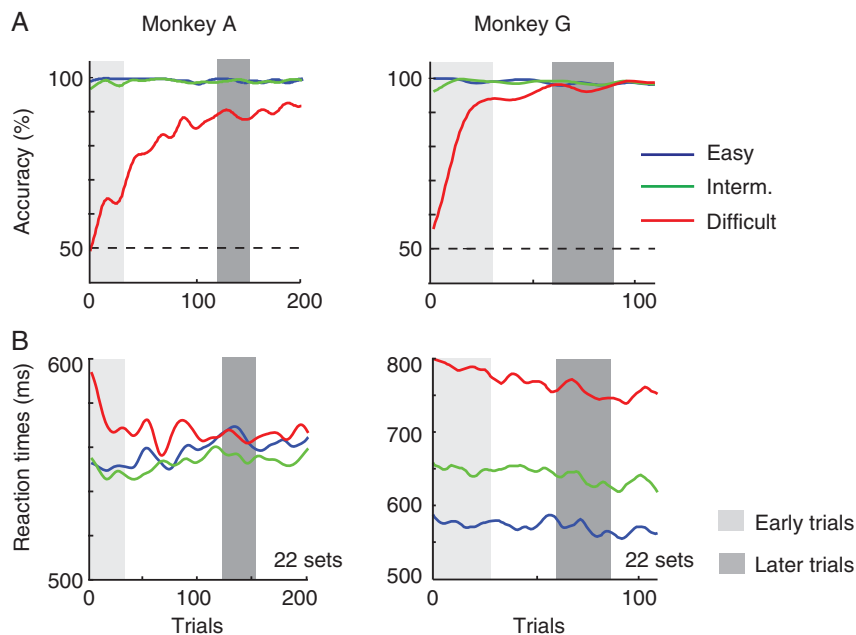


Figure 7. Behavioral performance in the delayed response task. (A) Accuracy as a function of trial number in the easy (blue), intermediate (green), and difficult condition (red). Gray bars illustrate bins of 30 trials (early: light gray and late: dark gray) that were selected for further analysis. (B) Reaction times in trials with the 3 difficulty levels.

intermediate conditions where the animals could select the icon closest to fixation so that they did not have to rely on its identity.

Discussion

We investigated the time course of response selection in a task where monkeys had to learn a new rule and obtained a number of new findings. First, we found that the learning process starts with a phase of fast and random decisions and ends with a phase with more accurate but slower decisions. Learning increased the reaction time by approximately 40 ms (Fig. 2) and it induced a comparable delay in the selection signals in frontal cortex (Fig. 5), implying that the monkeys allotted additional

time for an influence of newly learned icon identity on the decision. The results differ from previous studies that investigated decision-making in tasks where the rule was known, but the strength of sensory evidence varied (Kim and Shadlen 1999; Roitman and Shadlen 2002). If the rule is familiar, the decision process takes longer if the sensory evidence is weak (Pooresmaeili et al. 2014).

Thus, these 2 sources of uncertainty, unfamiliar rules, and weak evidence, have fundamentally different influences on response selection. This difference can be explained by a trade-off between speed and accuracy that optimizes reward income (Bogacz et al. 2010; Heitz and Schall 2012). If the rule is known but evidence is weak, it is strategic to wait for more evidence to

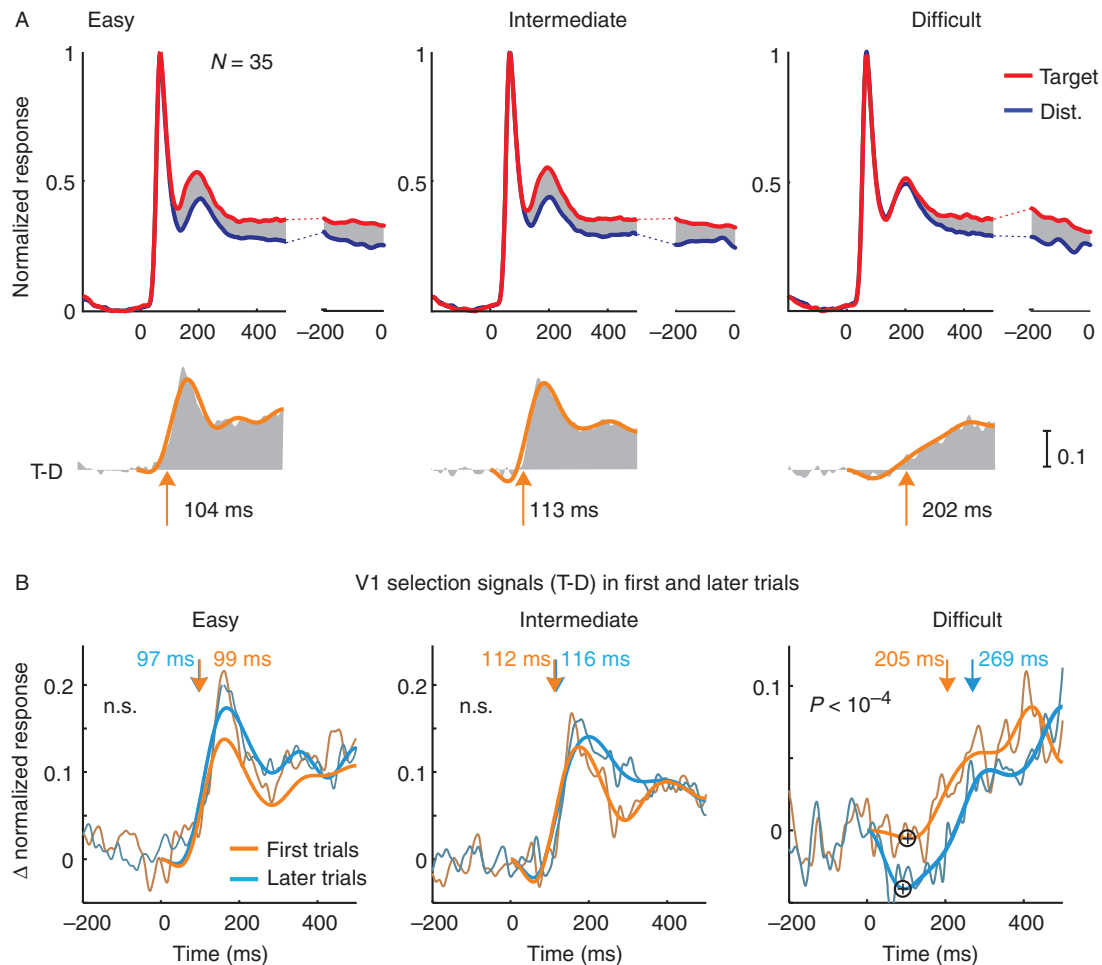


Figure 8. Neuronal activity in V1 in the icon-selection task. (A) Activity elicited by the target (red trace) and distractor curve (blue trace), averaged across all V1 recoding sites in the easy, intermediate, and difficult condition and over both monkeys, aligned on stimulus onset (left panels) and saccade (right). Lower insets show the selection signal and a curve that was fitted to determine latency. Arrows depict latency measured as the time point where the fitted function reached 33% of its maximum. (B) Selection signal (target minus distractor response) in the first (orange) trials and later (blue) trials of the easy, intermediate, and difficult condition. The influence of learning on the onset of the selection signal in the difficult condition was significant, but it was not significant in the easy and intermediate conditions (n.s.). Plus signs in the panel of the difficult condition mark the inversion of the selection signal, that is, when the target curve elicits weaker activity than the distractor.

improve accuracy before committing to a choice. If the rule is unknown, however, postponing the decision neither increases accuracy nor reward income, and decisions can therefore be made fast. However, as soon as the monkeys learned the icon identities in our task, the trade-off favored longer reaction times and they postponed their decisions. The additional delay was associated with a brief suppression of V1 activity elicited by the target curve near the target icon, which counteracted the later enhancement of the neuronal response elicited by this curve (Fig. 10). This result suggests that the monkeys started to focus on the target icon, causing a delay in the selection of the connected curve and eye movement target.

The second new finding is that selection signals are fed back to early visual cortex even at the start of the learning task when the monkeys are still guessing. V1 neurons selected the same curve as did FEF neurons, the right one in correct trials and the wrong one in erroneous trials. Theoretical studies suggested that feedback signals about the chosen action to visual cortex may play a role in the guidance of neuronal plasticity, by focusing synaptic changes on the visual representations that are related to the choice (Roelfsema et al. 2010; Rombouts et al. 2012, 2015). The feedback signals are delayed relative to the onset of the visual

stimulus, presumably because they depend on recurrent interactions within and between areas of the frontal, parietal, and visual cortex where neurons jointly select one of the 2 curves (Bruce and Goldberg 1985; Schall et al. 1995; Khayat et al. 2009; Pooresmaeili et al. 2014) and in accordance with theories that stress the importance of reentrant processing (Di Lollo et al. 2000; Lamme and Roelfsema 2000; Roelfsema 2005; Cichy et al. 2014; Tsotsos and Kruijne 2014).

Fast but Random Decisions

In the easy condition of the icon-selection task, FEF neurons started to discriminate between the target and distractor curve after approximately 145 ms, which is comparable with the timing of FEF selection signals in pop-out search (Thompson et al. 1996) and a simple curve-tracing task (Pooresmaeili et al. 2014). At the same time, the selection signal occurred earlier than the value of 235 ms in a motion discrimination task (Kim and Shadlen 1999; Ding and Gold 2012). If this motion discrimination task is made difficult because motion is weak, the accumulation of evidence for an eye movement decision can last up to several hundreds of milliseconds (Kim and Shadlen 1999; Roitman and Shadlen

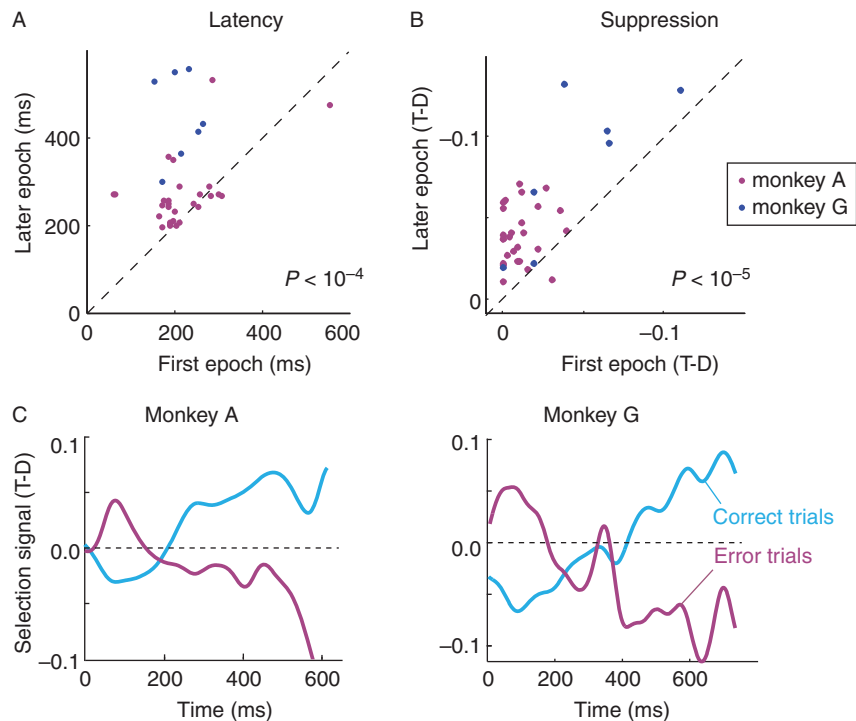


Figure 9. Influence of learning and errors on the V1 selection signal in the difficult condition. (A) Latency of the selection signal (onset of positive modulation) at individual V1 recording sites at the start of the sessions (x-axis) and at a later epoch when the monkeys' accuracy had increased (y-axis). Magenta, monkey A, Blue, monkey G. (B) Strength of suppression, quantified as the difference in activity elicited by the target and distractor curve (T-D, note the negative values on the axes, these values were normalized to the peak response; time window from 100 to 200 ms after stimulus onset) at the start of the session (x-axis) and in the later epoch (y-axis). (C) Time course of the selection signal in correct (light blue) and erroneous trials (magenta).




Figure 10. Hypothetical 2-step selection process that emerged during learning in the difficult condition. When the monkeys became proficient, the selection of the relevant icon initially suppressed the activity elicited by the adjacent target curve, because the animal first selected the relevant icon (yellow), which may have caused a suppressive zone (blue; Mexican hat profile) of attention. In the subsequent phase, V1 neurons select the target curve and FEF cells the correct target for an eye movement.

2002). In contrast, target selection in the difficult condition of the icon-selection task during the phase of random eye movement selection was rapid (see also Uchida et al. 2006). We do not know what drove these fast yet random decisions, but the rapid ramping of activity in the icon-selection task may be related to so-called urgency signals (Churchland et al. 2008; Drugowitsch et al. 2012), which have been proposed to enforce fast decisions if they are strategic. If the odds are unknown, these fast responses optimize reward income and can be used to quickly probe target identity.

A Transition to Slower and More Accurate Decisions

The monkeys' reaction times increased by approximately 40 ms when the animals became proficient with the difficult version

of the icon-selection task. In the early phase, trials of the difficult conditions could be correct only by chance, whereas in the later trials the accuracy in the difficult condition increased to >80%. This additional delay of 40 ms was apparently sufficient for the newly learned meaning of the icons to influence the decision. Interestingly, this delay is in accordance with Stanford et al. (2010) who systematically varied the amount of time where monkeys could sample a visual stimulus before they made their decision in a color discrimination task. Also in their task, a delay of 30 ms was sufficient for the transition from random decisions to near-perfect performance. The increase in reaction time was accompanied by a delay in the onset of the FEF selection signal of approximately 40 ms. Comparable variations in the onset time of the activity increase in FEF have been observed previously

(Pouget et al. 2011). In the easy and intermediate conditions, the monkeys' reaction times also increased with trial number. In these conditions, the monkeys could rely on the distance between the fixation point and the icons, which is a rule that they knew well and did not require new learning. Accordingly, the timing of the selection signal in FEF and in V1 in the easy and intermediate conditions hardly changed. At first sight, it may seem surprising that the reaction time increased, while the timing of the selection signal was stable. However, previous studies in FEF demonstrated that there is a variable delay between attentional selection and movement initiation so that delays in attentional selection contribute to reaction time, but are not the only determinant (Thompson et al. 1996; Murthy et al. 2001; Sato et al. 2001). A recent psychophysical study with human participants that carried out a variant of the curve-tracing task also obtained evidence for a variable and time-consuming process that transforms the perceptual decision into a motor response (Zylberberg et al. 2012). Our finding that the reaction time also increases in the easy and intermediate conditions even if it is only beneficial for the accuracy in the difficult condition suggests that the process responsible for the longer reaction time does not fully discriminate between the difficulty levels.

In contrast to the easy and intermediate conditions, we did observe an increase in latency of the FEF selection signal in the difficult condition, where the monkeys could only rely on the newly learned identity of the target icon. This increase in the latency of selection differs from results in prefrontal cortex and the basal ganglia in a task where monkeys had to learn the association between a symbol and a specific eye movement and their reversals (Asaad et al. 1998; Pasupathy and Miller 2005). In this task, learning decreased the onset of a selection signal from 700 to 250 ms. However, in these studies the monkeys fixated for 1.5 s, preventing them from making fast but random decisions when the odds are unknown.

It is likely that the increase in the latency of the selection signal contributed to the increase in reaction times in the difficult condition. We also considered the possibility that the slowing was caused by fatigue. However, the reaction times increased from the very beginning of the learning sessions, when the animals were highly motivated to do the task (Fig. 2B). Furthermore, we have also reanalyzed the pattern of reaction times in a search-then-trace task where we cued the color of the target icon at the fixation point so that new learning was unnecessary (Roelfsema et al. 2003). In this task, the reaction times did not increase with trial number, which suggests that the increase of reaction time observed in the present study is related to the learning of the icons.

Selection Signals in Visual Cortex

By examining the development of selection signals within a single recording session, the present results complement studies that examined perceptual learning in the visual cortex across many days (Schoups et al. 2001; Lee et al. 2002; Schwartz et al. 2002; Furmanski et al. 2004; Kourtzi et al. 2005; Sigman et al. 2005; Law and Gold 2008; Li et al. 2008). Learning across days increased the neuronal sensitivity to relevant feature variations, with some studies reporting strongest effects in early visual cortex (Sigman et al. 2005) and others in areas outside visual cortex (Law and Gold 2008). Unlike these previous studies, we did not measure the neuronal tuning for the to-be-learned visual stimuli, but signals related to the selection of one of these stimuli for a behavioral response.

The selection of one of the 2 icons boosted the V1 representation of the curve connecting it to the eye movement target. The

easy condition in the icon-selection task was comparable to previous curve-tracing studies, because monkeys could simply trace the curve that started at the fixation point (Roelfsema et al. 1998; Roelfsema 2006). During curve tracing, the V1 cells initially code the contours, but after a delay the target curve elicits stronger activity than a distractor. The latency of this selection signal depends on task difficulty (Roelfsema and Spekreijse 2001; Pooresmaeili et al. 2014). In the easy and intermediate conditions, the latency appeared to be even shorter in V1 than in FEF (104 and 113 ms vs. 145 and 155 ms). However, caution is warranted, because the V1 and FEF data were obtained in different monkeys and in different versions of the task (reaction time vs. delayed response). Although many theories propose that selection signals in visual cortex are caused by feedback from areas of frontal and parietal cortex (Hochstein and Ahissar 2002; Moore and Armstrong 2003; Ekstrom et al. 2008; Buffalo et al. 2010; Noudoost and Moore 2011; Zhou and Desimone 2011), selection signals in V1 and FEF in the curve-tracing task have a similar latency, which is compatible with a more active participation of visual cortex (Pooresmaeili et al. 2014). In accordance with this view, V1 neurons selected the wrong curve in error trials, which were abundant at the start of the learning process (see also Roelfsema and Spekreijse 2001).

The Emergence of a 2-Step Selection Process in Visual Cortex

The V1 selection signal occurred later in the difficult condition than in the easier conditions, because the animals had to rely on icon identity. When the animals learned to consider icon identity, this delay increased even further, just as in FEF. Learning to select the correct icon caused a brief early phase of inhibition of the V1 representation of the connected curve. This suppression is reminiscent of studies demonstrating that attentional selection sometimes takes the form of a Mexican hat with an enhanced representation of the relevant shape and a suppression of nearby items (Tsotsos 1990; Schall et al. 1995; Müller and Kleinschmidt 2004; Hopf et al. 2006; Tsotsos et al. 2008) (Fig. 10). The inferotemporal cortex is a putative source for top-down influences that boost the representation of the relevant icon and suppress the adjacent curve (middle panel of Fig. 10). Jagadeesh et al. (2001) trained monkeys to select one of 2 pictures, which was consistently paired with reward, with a design that resembles the present icon-selection task. They found that the representation of the rewarded picture was enhanced in the inferotemporal cortex during learning, whereas the representation of unrewarded pictures was suppressed. Neurons in the inferotemporal cortex (and perhaps area V4) that represent the rewarded icon with enhanced activity might feed back to early visual cortex to increase the representation of this icon and to suppress the adjacent curve. Interestingly, suppression of the target curve did not occur for the easy and intermediate difficulty levels (Fig. 8B), and it has also not been observed in previous curve-tracing studies where all contour elements of the target curve were always highlighted with extra neuronal activity (Roelfsema 2006; Pooresmaeili and Roelfsema 2014). Thus, the suppression in the difficult condition may be specifically related to the learning of the identity of the target icon. When the animals had learned the target icon, it took additional time before the inhibition was overcome, and the target curve was highlighted with additional activity.

These results imply that the monkeys' final strategy transformed into a 2-step selection process (Thompson et al. 1996; Carpenter 2004), where they first selected the relevant icon and

subsequently the connected curve and eye movement target. Previous work examined multistep selection processes during visual routines that were composed of successive curve-tracing and visual search operations. In these tasks, neurons in visual cortex selected visual objects consecutively, at the time that they become relevant so that the progression of the routine could be monitored with high temporal precision (Ullman 1984; Roelfsema et al. 2003; Moro et al. 2010). For example, if the monkey first has to search for an icon with a particular color that cues the start of the target curve, V1 activity elicited by this target icon increases before the enhanced activity spreads across the target curve (Roelfsema et al. 2003; Moro et al. 2010). In the present icon-selection task, we therefore obtained evidence for the insertion of an additional icon-selection process before curve tracing that delayed the selection of the eye movement target in FEF. The present results thereby also contribute to our understanding of the emergence multistep selection processes in the brain during learning. The neuronal correlates of such a multistep visual selection task in low-level visual areas as well as in the frontal cortex lend strong support for reentrant models of visual processing (Di Lollo et al. 2000; Lamme and Roelfsema 2000; Roelfsema 2005; Cichy et al. 2014; Tsotsos and Kruijne 2014).

Funding

This work was supported by NWO (Brain and Cognition grant 433-09-208 and ALW grant 823-02-010), the European Union Seventh Framework Program (project 269921 “BrainScaleS” and Marie-Curie Action PITN-GA-2011-290011), and an ERC advanced grant (Cortic_al_gorithms; no 339490) awarded to P.R.R.

Notes

The authors thank Kor Brandsma for technical support. *Conflict of Interest:* The authors declare no competing financial interests.

References

- Anderson BA, Laurent PA, Yantis S. 2011. Value-driven attentional capture. *Proc Natl Acad Sci USA*. 108:10367–10371.
- Asaad WF, Eskandar EN. 2011. Encoding of both positive and negative reward prediction errors by neurons of the primate lateral prefrontal cortex and caudate nucleus. *J Neurosci*. 31:17772–17787.
- Asaad WF, Rainer G, Miller EK. 1998. Neural activity in the primate prefrontal cortex during associative learning. *Neuron*. 21:1399–1407.
- Bichot NP, Schall JD, Thompson KG. 1996. Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature*. 381:697–699.
- Bogacz R, Wagenmakers EJ, Forstmann BU, Nieuwenhuis S. 2010. The neural basis of the speed-accuracy tradeoff. *Trends Neurosci*. 33:10–16.
- Bour LJ, van Gisbergen JA, Bruijns J, Ottes FP. 1984. The double magnetic induction method for measuring eye movement—results in monkey and man. *IEEE Trans Biomed Eng*. 31:419–427.
- Brasted PJ, Wise SP. 2004. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *Eur J Neurosci*. 19:721–740.
- Bruce CJ, Goldberg ME. 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol*. 53:603–635.
- Buffalo EA, Fries P, Landman R, Liang H, Desimone R. 2010. A backward progression of attentional effects in the ventral stream. *Proc Natl Acad Sci USA*. 107:361–365.
- Carpenter RHS. 2004. Contrast, probability, and saccadic latency; evidence for independence of detection and decision. *Curr Biol*. 14:1576–1580.
- Chelazzi L, Perlato A, Santandrea E, Della Libera C. 2013. Rewards teach visual selective attention. *Vision Res*. 85:58–72.
- Chen LL, Wise SP. 1996. Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurosci*. 16:3067–3081.
- Chen LL, Wise SP. 1995a. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurophysiol*. 73:1101–1121.
- Chen LL, Wise SP. 1995b. Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. *J Neurophysiol*. 73:1122–1134.
- Churchland AK, Kiani R, Shadlen MN. 2008. Decision-making with multiple alternatives. *Nat Neurosci*. 11:693–702.
- Cichy RM, Pantazis D, Oliva A. 2014. Resolving human object recognition in space and time. *Nat Neurosci*. 17:455–462.
- Cohen JY, Heitz RP, Woodman GF, Schall JD. 2009. Neural basis of the set-size effect in frontal eye field: timing of attention during visual search. *J Neurophysiol*. 101:1699–1704.
- Cohen MR, Maunsell JHR. 2009. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci*. 12:1594–1600.
- de Boor C. 1978. A practical guide to splines. Springer-Verlag.
- Della Libera C, Chelazzi L. 2009. Learning to attend and to ignore is a matter of gains and losses. *Psychol Sci*. 20:778–784.
- Deubel H, Schneider WX. 1996. Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res*. 36:1827–1837.
- Di Lollo V, Enns JT, Rensink R a. 2000. Competition for consciousness among visual events: the psychophysics of reentrant visual processes. *J Exp Psychol Gen*. 129:481–507.
- Ding L, Gold JI. 2012. Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb Cortex*. 22:1052–1067.
- Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A. 2012. The cost of accumulating evidence in perceptual decision making. *J Neurosci*. 32:3612–3628.
- Ekstrom LB, Roelfsema PR, Arsenault JT, Bonmassar G, Vanduffel W. 2008. Bottom-up dependent gating of frontal signals in early visual cortex. *Science*. 321:414–417.
- Erickson C a, Desimone R. 1999. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J Neurosci*. 19:10404–10416.
- Furmanski CS, Schluppeck D, Engel SA. 2004. Learning strengthens the response of primary visual cortex to simple patterns. *Curr Biol*. 14:573–578.
- Gold JI, Shadlen MN. 2007. The neural basis of decision making. *Annu Rev Neurosci*. 30:535–574.
- Gold JI, Shadlen MN. 2001. Neural computations that underlie decisions about sensory stimuli. *Trends CognSci*. 5:10–16.
- Hanes DP, Schall JD. 1996. Neural control of voluntary movement initiation. *Science*. 274:427–430.
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN. 2011. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci*. 31:6339–6352.
- Heilbronner SR, Platt ML. 2013. Causal evidence of performance monitoring by neurons in posterior cingulate cortex during learning. *Neuron*. 80:1384–1391.

- Heitz RP, Schall JD. 2012. Neural mechanisms of speed-accuracy tradeoff. *Neuron*. 76:616–628.
- Hickey C, Chelazzi L, Theeuwes J. 2010. Reward changes salience in human vision via the anterior cingulate. *J Neurosci*. 30:11096–11103.
- Hochstein S, Ahissar M. 2002. View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron*. 36:791–804.
- Hopf J-M, Boehler CN, Luck SJ, Tsotsos JK, Heinze H-J, Schoenfeld MA. 2006. Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proc Natl Acad Sci USA*. 103:1053–1058.
- Jagadeesh B, Chelazzi L, Mishkin M, Desimone R. 2001. Learning increases stimulus salience in anterior inferior temporal cortex of the macaque. *J Neurophysiol*. 86:290–303.
- Khayat PS, Pooresmaeili A, Roelfsema PR. 2009. Time course of attentional modulation in the frontal eye field during curve tracing. *J Neurophysiol*. 101:1813–1822.
- Kim JN, Shadlen MN. 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci*. 2:176–185.
- Kourtzi Z, Betts LR, Sarkheil P, Welchman AE. 2005. Distributed neural plasticity for shape learning in the human visual cortex. *PLoS Biol*. 3:e204.
- Kowler E, Anderson E, Doshier B, Blaser E. 1995. The role of attention in the programming of saccades. *Vision Res*. 35:1897–1916.
- Lamme V a F, Roelfsema PR. 2000. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci*. 23:571–579.
- Law C-T, Gold JI. 2008. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci*. 11:505–513.
- Lee TS, Yang CF, Romero RD, Mumford D. 2002. Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nat Neurosci*. 5:589–597.
- Li W, Pièch V, Gilbert CD. 2008. Learning to link visual contours. *Neuron*. 57:442–451.
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature*. 412:150–157.
- Maunsell JHR. 2004. Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci*. 8:261–265.
- Moore T, Armstrong KM. 2003. Selective gating of visual signals by microstimulation of frontal cortex. *Nature*. 421:370–373.
- Moro SI, Tolboom M, Khayat PS, Roelfsema PR. 2010. Neuronal activity in the visual cortex reveals the temporal order of cognitive operations. *J Neurosci*. 30:16293–16303.
- Müller NG, Kleinschmidt A. 2004. The attentional “spotlight’s” penumbra: center-surround modulation in striate cortex. *Neuroreport*. 15:977–980.
- Murthy A, Thompson KG, Schall JD. 2001. Dynamic dissociation of visual selection from saccade programming in frontal eye field. *J Neurophysiol*. 86:2634–2637.
- Noudoost B, Moore T. 2011. Control of visual cortical signals by prefrontal dopamine. *Nature*. 474:372–375.
- Pasupathy A, Miller EK. 2005. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*. 433:873–876.
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J. 2009. Reward modulates attention independently of action value in posterior parietal cortex. *J Neurosci*. 29:11182–11191.
- Peng X, Sereno ME, Silva AK, Lehky SR, Sereno AB. 2008. Shape selectivity in primate frontal eye field. *J Neurophysiol*. 100:796–814.
- Pooresmaeili A, Poort J, Roelfsema PR. 2014. Simultaneous selection by object-based attention in visual and frontal cortex. *Proc Natl Acad Sci USA*. 111:6467–6472.
- Pooresmaeili A, Roelfsema PR. 2014. A growth-cone model for the spread of object-based attention during contour grouping. *Curr Biol*. 24:2869–2877.
- Pouget P, Logan GD, Palmeri TJ, Boucher L, Paré M, Schall JD. 2011. Neural basis of adaptive response time adjustment during saccade countermanding. *J Neurosci*. 31:12604–12612.
- Raymond JE, O’Brien JL. 2009. Selective visual attention and motivation: the consequences of value learning in an attentional blink task. *Psychol Sci*. 20:981–988.
- Roelfsema PR. 2006. Cortical algorithms for perceptual grouping. *Annu Rev Neurosci*. 29:203–227.
- Roelfsema PR. 2005. Elemental operations in vision. *Trends Cogn Sci*. 9:226–233.
- Roelfsema PR, Khayat PS, Spekreijse H. 2003. Subtask sequencing in the primary visual cortex. *Proc Natl Acad Sci USA*. 100:5467–5472.
- 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, Spekreijse H. 2001. The representation of erroneously perceived stimuli in the primary visual cortex. *Neuron*. 31:853–863.
- Roelfsema PR, van Ooyen A, Watanabe T. 2010. Perceptual learning rules based on reinforcers and attention. *Trends Cogn Sci*. 14:64–71.
- 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.
- Rombouts J, Roelfsema P, Bohte S. 2012. Neurally plausible reinforcement learning of working memory tasks. *NIPS*. 1880–1888.
- Rombouts JO, Bohte SM, Roelfsema PR. 2015. How attention can create synaptic tags for the learning of working memories in sequential tasks. *PLoS Comput Biol*. 7:1–34.
- Sato T, Murthy A, Thompson KG, Schall JD. 2001. Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron*. 30:583–591.
- Schall JD. 1991. Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *J Neurophysiol*. 66:559–579.
- Schall JD, Hanes DP, Thompson KG, King DJ. 1995. Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. *J Neurosci*. 15:6905–6918.
- Schoups A, Vogels R, Qian N, Orban G. 2001. Practising orientation identification improves orientation coding in V1 neurons. *Nature*. 412:549–553.
- Schwartz S, Maquet P, Frith C. 2002. Neural correlates of perceptual learning: a functional MRI study of visual texture discrimination. *Proc Natl Acad Sci USA*. 99:17137–17142.
- Serences JT. 2008. Value-based modulations in human visual cortex. *Neuron*. 60:1169–1181.
- Sigman M, Pan H, Yang Y, Stern E, Silbersweig D, Gilbert CD. 2005. Top-down reorganization of activity in the visual pathway after learning a shape identification task. *Neuron*. 46:823–835.
- Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E. 2010. Perceptual decision making in less than 30 milliseconds. *Nat Neurosci*. 13:379–385.
- Stănişor L, van der Togt C, Pennartz CMA, Roelfsema PR. 2013. A unified selection signal for attention and reward in primary visual cortex. *Proc Natl Acad Sci USA*. 110:9136–9141.

- Supèr H, Roelfsema PR. 2005. Chronic multiunit recordings in behaving animals: advantages and limitations. *Prog Brain Res.* 147:263–282.
- Thompson KG, Bichot NP, Sato TR. 2005. Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J Neurophysiol.* 93:337–351.
- Thompson KG, Hanes DP, Bichot NP, Schall JD. 1996. Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J Neurophysiol.* 76:4040–4055.
- Tremblay L, Hollerman JR, Schultz W. 1998. Modifications of reward expectation-related neuronal activity during learning in primate striatum. *J Neurophysiol.* 80:964–977.
- Tsotsos JK. 1990. Analyzing vision at the complexity level. *Behav Brain Sci.* 13:423–445.
- Tsotsos JK, Kruijne W. 2014. Cognitive programs: software for attention's executive. *Front Psychol.* 5:1260.
- Tsotsos JK, Rodríguez-Sánchez AJ, Rothenstein AL, Simine E. 2008. The different stages of visual recognition need different attentional binding strategies. *Brain Res.* 1225:119–132.
- Uchida N, Kepecs A, Mainen ZF. 2006. Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat Rev Neurosci.* 7:485–491.
- Ullman S. 1984. Visual routines. *Cognition.* 18:97–159.
- Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. 2003. Single neurons in the monkey hippocampus and learning of new associations. *Science.* 300:1578–1581.
- Xing D, Yeh C-I, Shapley RM. 2009. Spatial spread of the local field potential and its laminar variation in visual cortex. *J Neurosci.* 29:11540–11549.
- Zhou H, Desimone R. 2011. Feature-based attention in the frontal eye field and area V4 during visual search. *Neuron.* 70:1205–1217.
- Zylberberg A, Ouellette B, Sigman M, Roelfsema PR. 2012. Decision making during the psychological refractory period. *Curr Biol.* 22:1795–1799.