

# Encoding of Reward and Decoding Movement from the Frontal Eye Field during Smooth Pursuit Eye Movements

Adi Lixenberg and  Mati Joshua

The Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Expectation of reward potentiates sensorimotor transformations to drive vigorous movements. One of the main challenges in studying reward is to determine how representations of reward interact with the computations that drive behavior. We recorded activity in smooth pursuit neurons in the frontal eye field (FEF) of two male rhesus monkeys while controlling the eye speed by manipulating either reward size or target speed. The neurons encoded the different reward conditions more strongly than the different target speed conditions. This pattern could not be explained by differences in the eye speed, since the eye speed sensitivity of the neurons was also larger for the reward conditions. Pooling the responses by the preferred direction of the neurons attenuated the reward modulation and led to a tighter association between neural activity and behavior. Therefore, a plausible decoder such as the population vector could explain how the FEF both drives behavior and encodes reward beyond behavior.

**Key words:** eye movements; frontal eye field; primates; reward; smooth pursuit

## Significance Statement

Motor areas combine sensory and reward information to drive movement. To disambiguate these sources, we manipulated the speed of smooth pursuit eye movements by controlling either the size of the reward or the speed of the visual motion signals. We found that the relationship between activity in frontal eye field and eye kinematics varied: the eye speed sensitivity was larger for the different reward conditions than for the different target speed conditions. Decoders that pooled signals by the preferred direction of the neurons attenuated the reward modulations. These decoders may indicate how reward can be both encoded beyond eye kinematics at the single neuron level and drive movement at the population level.

## Introduction

The drive for rewards controls almost every facet of behavior, from stereotypic reflexive movements to complex voluntary actions. Reward-related signals are ubiquitous to nearly every sub-area of the brain (Paton et al., 2006; Joshua et al., 2008; Vickery et al., 2011; Wagner et al., 2017; Retailliau and Morris, 2018). However, it remains unclear how these widespread reward representations are related to the computations and transformations that are implemented in different brain structures. Sensorimotor transformations, and in particular the eye movement system, constitute an excellent model for studying reward at the level of computations and transformations since much is already known about the basic signals that are transmitted at different levels of

the circuit (Robinson, 1981; Sparks, 2002; Krauzlis, 2004; Lisberger, 2010) and the way reward drives ocular behavior (Takikawa et al., 2002; Joshua and Lisberger, 2012; Reppert et al., 2015).

Where and how does reward interact with the sensorimotor transformation in the eye movement system? The frontal eye field (FEF) has been shown to be a key structure in both controlling eye movements (Bizzi, 1968; Bruce and Goldberg, 1985; Gottlieb et al., 1994; Tanaka and Fukushima, 1998) and encoding reward (Roesch and Olson, 2003, 2005; Ding and Hikosaka, 2006; Glaser et al., 2016). The challenge involved in studying reward in the motor system lies in dissociating modulations that are related to reward from modulations that reflect the motor command. When examining the ways in which reward potentiates movement, the null hypothesis that modulations reflect movement and not reward processing must be entertained. Recently, Glaser et al. (2016) exploited the natural variability of behavior to show that reward drives larger FEF responses even when behavior is equalized. However, it is still unclear whether the structure of FEF selectivity for reward can translate into shifts in motor parameters as a function of reward size.

Here, we build on these results by studying the structure of reward and movement representations of a population of pursuit

Received July 2, 2018; revised Oct. 2, 2018; accepted Oct. 5, 2018.

Author contributions: A.L. edited the paper; M.J. wrote the first draft of the paper; M.J. designed research; M.J. performed research; A.L. and M.J. analyzed data; M.J. wrote the paper.

This work was supported by Human Frontier Science Program Career Development Award (HFSP-CDA to M.J.) and the Israel Science Foundation (ISF grant 380/17). We thank Stephen Lisberger for his valuable input and support in the early stages of this project.

The authors declare no competing financial interests.

Correspondence should be addressed to Mati Joshua, The Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, E-mail: mati.joshua@mail.huji.ac.il.

<https://doi.org/10.1523/JNEUROSCI.1654-18.2018>

Copyright © 2018 the authors 0270-6474/18/3810515-10\$15.00/0

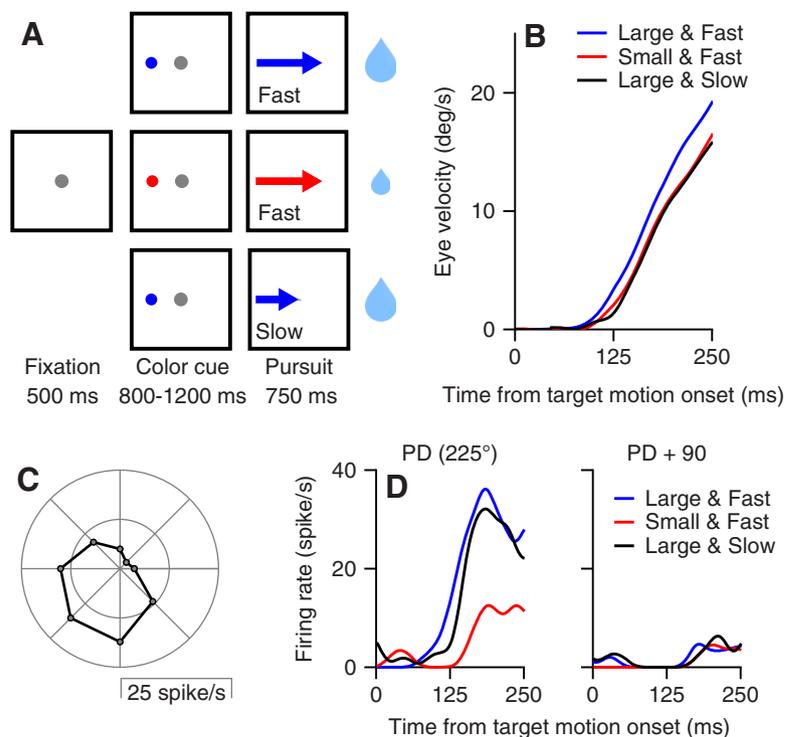
responsive neurons in the FEF. We used tasks where we independently manipulated pursuit eye movement either by controlling target speed or by controlling the expected reward size to better understand how mixed selectivity for reward and motor parameters are organized at the single neuron and population levels. To probe the associations between the encoding of reward and target motion and the decoding of movement, we tested the hypothesis that a population-vector decoding model of the FEF can account for the shifts in movement metrics with reward size. We found that the relationship between FEF activity and eye speed was not fixed. Specifically, eye speed sensitivity was larger when estimated in terms of the different reward conditions than in terms of the different target speed conditions. Averaging the responses of the cells by their preferred direction eliminated some of the reward modulations and resulted in a population response that better matched the behavioral modulation. Decoders such as the population vector that pool signals by the preferred direction of the neurons may thus explain how representations of reward drive behavior. Therefore, reward can be both encoded beyond eye kinematics at the single neuron level and drive movement at the population level.

## Materials and Methods

We collected neural and behavior data from two male rhesus macaque monkeys (*Macaca mulatta*). All procedures were approved in advance by the Institutional Animal Care and Use Committee at Duke University where the experiments were performed. Procedures were in strict compliance with the National Institutes of Health *Guide for the Care and Use of Laboratory Animals*. In each monkey, we implanted a head holder to restrain the monkey's head in the experiments, as well as a coil of wire on one eye to measure eye position using the magnetic search coil technique. After the monkeys had recovered from surgery, we trained them to track spots of light that moved across a video monitor placed in front of them. In a subsequent surgery in each monkey, we placed a recording cylinder stereotaxically over the frontal eye field.

Up to five quartz-insulated tungsten electrodes were lowered into the caudal parts of the FEF to record spikes using a Mini-Matrix System (Thomas Recording). Signals were high-pass filtered with a cutoff frequency of 150 Hz and digitized at a sampling rate of 40 kHz (Plexon Multichannel Acquisition Processor). For the detailed data analysis, we sorted spikes off-line (Plexon). For sorting, we used principal component analysis and corrected manually for errors. Sorted spikes were converted to time stamps with a time resolution of 1 ms and were inspected again visually to look for instability and obvious sorting errors. The neural activity in the preparatory activity of the large reward trials is described in a previous publication (Raghavan and Joshua, 2017).

**Experimental design.** When lowering the electrodes, we looked for neurons that responded during pursuit eye movements. To test neurons for pursuit responses, we used a target (white circle, 0.5° diameter) that moved in one of eight directions. The targets stepped in one direction and moved by 20°/s in the other direction (Rashbass, 1961; step ramp). Typically, neurons in the parts of the FEF that are closer to the brain surface did not respond to our search stimuli, and we had to lower our



**Figure 1.** Behavioral task and example of reward enhancement of both pursuit eye movement and neural responses. **A**, The sequence of snapshots illustrates the structure of the behavioral task. The size of the drop of water represents the amount of reward given at the end of the trials. The length of the arrow represents target speed. **B**, Average eye speed in the large–fast (blue), small–fast (red), and large–slow (black) conditions in one recording session. **C**, Directional tuning of a sample neuron. The direction of each dot represents the target motion direction, and the eccentricity is the average activity in the first 250 ms after target motion onset. **D**, Average activity for the neuron presented with the tuning curve shown in **C**, which was recorded during the behavior presented in **B**. Left and right plots show motion in the preferred direction of the cells and in a direction rotated 90° from the preferred direction. Averages were smoothed with a Gaussian filter with a standard deviation of 10 ms.

electrodes deeper into the arcuate sulcus until we found neurons that responded to pursuit.

After we characterized pursuit tuning, the monkeys engaged in the main experiment protocol in which we manipulated the reward size (Fig. 1A). Each trial started with a bright white target that appeared in the center of the screen. After 500 ms of presentation, in which the monkey was required to acquire fixation, a colored target appeared 3° eccentric to the fixation target. The color of the target signaled the size of the reward the monkey would receive if it tracked the target. For Monkey Y, we used yellow to signal a large reward (0.1–0.2 ml) and green to signal a small reward (0.05 ml); in Monkey X, we reversed the associations. The eccentric color target appeared for 800–1200 ms, in which the monkeys continued to fixate on the white target in the center of the screen; gaze shifts resulted in abortion of the trial. Then the white target disappeared, and the eccentric target stepped from its 3° position and started to move continuously toward the center of the screen, prompting the initiation of smooth pursuit. For both monkeys, we used a fast target motion of 30°/s and step to a position 4–6° from the center of the screen. For the slow target motion, we used a target speed of 20°/s for Monkey Y and 15°/s for Monkey X, with a step to 2–4° eccentric to the center of the screen. The exact step size was optimized to minimize initial saccades. The small differences between monkeys in target speed were a result of our attempt to match the eye speed in the small reward condition to the slow target motion. The target moved for 750 ms and then stopped and stayed still for an additional 400–600 ms. When the eye was within a  $2 \times 2^\circ$  window around the target, the monkey received a juice reward.

We had three interleaved types of trials. Trials in which the monkey expected a large reward and the target moved fast, trials in which the monkey expected a small reward and the target moved fast, and trials in which the monkey expected a large reward and the target moved slower. The purpose of this design was to compare encoding of target speed to the

reward. We did not collect neural data for the conditions in which the target moved slowly and the monkey expected a small reward; therefore, we could not study the interactions between the encoding of reward and target speed.

In each recording session, we analyzed data from trials in which the targets moved in one of two orthogonal directions. In each session, one target direction was chosen to align with the preferred direction of at least one neuron recorded using our five-channel system. Tuning could differ between neurons on different recording channels; therefore, many of the neurons included in the analysis were recorded with a target that was not moving in their preferred direction. We opted for this experimental design rather than one involving extensively probing a single cell for many directions because often (especially in Monkey X) sessions with many trials with the small reward were difficult for the monkeys to complete. By probing only a subset of the directions, we ensured that the session would be short and enough data could be collected from a single cell per condition. Overall, we analyzed the activity of 176 neurons across animals ( $n = 111$  from Monkey Y,  $n = 65$  from Monkey X) that were recorded during at least 10 trials per condition. Some neurons were recorded in multiple blocks that had different task configurations with respect to movement directions. This added 26 recording sessions to our database (18 and 8 from Monkeys Y and X, respectively). For clarity of presentation in the text, we refer to the 202 analyzed sessions as “neurons.” Keeping to only one session per each neuron did not alter any of our findings. We used the 124 of 202 neurons that were directionally tuned ( $p < 0.05$  one-way ANOVA) for further analysis.

**Data analysis.** All analyses were performed using MATLAB (Mathworks). To study the time-varying properties of the responses, we calculated the peristimulus time histogram (PSTH) in 1 ms resolution. We then smoothed the PSTH with a 10 ms standard deviation Gaussian window. To calculate the modulation by target speed or reward, we averaged the responses in the first 250 ms after target motion. We then calculated the difference between the responses in the small–fast and large–fast conditions and the difference between the large–slow and large–fast target conditions. To calculate eye speed sensitivity, we divided these values by the difference in eye speed in the corresponding conditions.

To calculate the difference in time between either the reward condition with the same target speed or target speed conditions with the same reward, we calculated  $D'$  and the significance of the Wilcoxon rank sum test ( $p < 0.05$ ) between conditions in sliding 100 ms windows. We confirmed that increasing or decreasing the window size did not alter our conclusions.

The  $D'$  was calculated as  $D' = \frac{\mu_1 - \mu_2}{\sqrt{\frac{1}{2}(\sigma_1^2 + \sigma_2^2)}}$ , where  $\mu_i$  and  $\sigma_i$  are

the average and standard deviation of the measurement in the specific window of the analysis, respectively. We performed this analysis for the firing rate, eye speed, and eye acceleration. The analysis was performed for all neurons included in the analysis data set and corresponding behavior. In this analysis, we treated the time from 20 ms before the saccade to 20 ms after the saccade as missing data both for behavior and for neural activity. For the population analysis, we calculated the average of the absolute value of  $D'$ . We used bootstraps to test for the significance of the effects of  $D'$ . For each neuron, we randomly swapped the  $D'$  value for different reward conditions and the different target speed conditions. This procedure kept the temporal correlations of  $D'$  for a single cell but removed any main effect at the population level. We repeated this process 10,000 times and then derived the same measure we used to quantify the effects on the data and on the shuffled responses. A measure was considered significant if it was within 5% of the margins of the same measure of the shuffled population.

To calculate the population vector from the average PSTH, for each time bin, we calculated the weighted sum of the responses across the eight directions of movement. The population vector (PV) was defined as  $PV_x(t) = (1/n)\sum_{\theta} \text{PSTH}_{\theta} \cdot \cos \theta$  and  $PV_y(t) = (1/n)\sum_{\theta} \text{PSTH}_{\theta} \cdot \sin \theta$ , where  $\theta$  is the angle relative to the preferred direction,  $\text{PSTH}$  is the time-varying response, and  $n$  is the number of directions. We used the amplitude of the population vector to quantify the population response.

To calculate the tuning curves, we averaged the responses in the first 250 ms of the movement. The preferred direction of the neurons was characterized before the reward-size task. We recorded activity while the monkey tracked a target in one out of eight directions. In these trials, the monkey tracked a white target that did not explicitly signal the size of the upcoming reward. At the end of this trial, the monkey received an intermediate size reward. We calculated the preferred direction of the neuron as the direction that was closest to the vector average of the responses across directions (direction of the center of mass). Other characterizations such as fitting a von Mises function or choosing the direction with maximal response did not alter the findings. We used the preferred direction to calculate the population tuning curve by aligning all the responses to the preferred direction and averaging across cells. The preferred direction of the neurons was characterized before the reward-size task.

## Results

### Expectation of reward potentiates smooth pursuit behavior and neural responses in the FEFs

We recorded the activity of neurons from the FEFs while monkeys were engaged in a smooth pursuit task in which we manipulated the reward size (Joshua and Lisberger, 2012) and target speed (Fig. 1A). Before target motion, a color target indicated the size of the reward the monkey would receive if it tracked the target. We had three interleaved experimental conditions. In the first condition (Fig. 1A, top), the monkey expected a large reward and the target moved fast (30°/s). In the second condition, (Fig. 1A, middle), the monkey expected a small reward and the target moved fast. In the third condition (Fig. 1A, bottom), the monkey expected a large reward but the target moved more slowly than in the other conditions (20°/s and 15°/s for Monkeys Y and X, respectively). We term these the large–fast, small–fast, and large–slow conditions.

Both reward and target speed modulated eye speed. The monkeys initiated smooth eye movements that were faster when they were expecting a large versus a small reward with the same target speed (large–fast vs small–fast; Fig. 1B; Joshua and Lisberger 2012). Their eyes moved faster for the faster target condition when the reward size was the same (large–fast vs large–slow; Fig. 1B). The initial 250 ms after target motion onset was mostly free of catch-up saccades. Therefore, by focusing our analysis on early pursuit, we could probe the effect of reward on the transformation from visual motion to pursuit while excluding the saccades. Reward and target speed modulated the firing rate of the cells. After we characterized the preferred direction of the cells (Fig. 1C), we recorded activity in at least two directions during the task (see Materials and Methods). Figure 1D shows a neuron that was recorded in the preferred direction and in a direction orthogonal to the preferred direction while reward and target speed were manipulated. In the preferred direction, the neuron responded maximally to the large–fast condition, and the response in the large–slow condition was slightly smaller. The response in the small–fast condition was substantially smaller than the responses in the other conditions. The initial eye speed during this session of neural recording was the fastest for the large–fast condition and was similarly slower in the small–fast and large–slow conditions (Fig. 1B). Thus, for this sample neuron, the response pattern could not be attributed to the differences in the concurrent eye speed.

### The eye speed sensitivity of neurons is larger for reward than for target speed

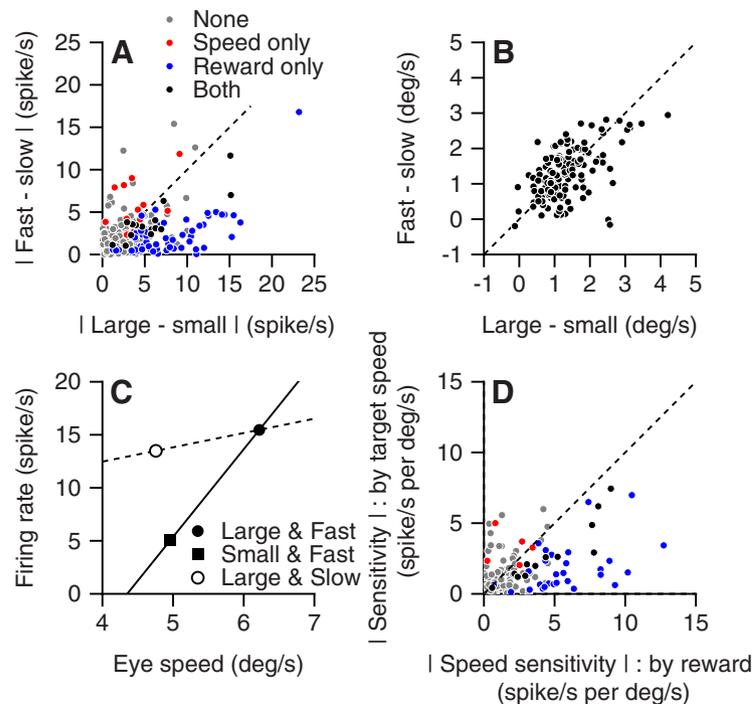
As found for the sample neuron (Fig. 1D), the rate modulation related to the reward tended to be larger than the modulation

related to the speed of the target. To assess the effect of the reward size and target speed on the perimovement neural response for each neuron, we averaged the responses in the first 250 ms after target motion onset. We then calculated the rate difference between reward conditions with the same target speed (large–fast vs small–fast) and between target speed conditions with the same reward (large–fast vs large–slow). Across the population, neurons modulated their firing rates more strongly between reward conditions than between target speed conditions ( $p < 0.01$ , signed rank test). This is depicted in Figure 2A, which shows the tendency of the dots to lie beneath the equality line. Furthermore, more neurons discriminated significantly between the large–fast and small–fast conditions than between the large–fast and large–slow conditions ( $p < 0.01$   $\chi^2$  test; Fig. 2A, fraction of blue and black vs red and black dots). Note that the number of trials was equalized, and different conditions were interleaved; thus, larger absolute modulations for the reward cannot be attributed to differences in noise in the estimates of the average responses.

The larger reward modulation did not reflect larger behavioral modulations. We chose the speed of the slow-motion target such that it led to a decrease in eye speed that was comparable to the decrease for small reward expectancy (Fig. 2B). The overall difference in the average eye speed during the first 250 ms of movement between the large–fast and small–fast conditions was not significantly different from the difference between the large–fast and large–slow target conditions (Fig. 2B;  $p = 0.3$ , signed rank test). Furthermore, when we selected only the sessions in which eye speed was more strongly modulated by target speed than by reward (Fig. 2B, only dots above the diagonal), the reward-related firing rate modulations were still larger than the target speed modulations ( $p < 0.01$ , signed rank test). Thus, the patterns of modulation in neural activity were different from the patterns of behavior modulations.

To compare target speed and reward modulations directly, we calculated the eye speed sensitivity from the target speed conditions and the reward conditions. Figure 2C illustrates this analysis for a sample neuron. We calculated the linear fit between the eye speed and the firing rate for the fast and slow target motion trials (Fig. 2C, dashed line) and the fit between firing rate and eye speed for the large and small rewards (Fig. 2C, solid line). If all the rate modulations could be attributed to the eye speed, we would expect the lines to have the same slope. In cases where eye speed sensitivity depended on the context, we would expect different slopes. The neuron in Figure 2C evidenced such a case in which the eye speed sensitivity was larger for the different reward conditions than for the different target speed conditions.

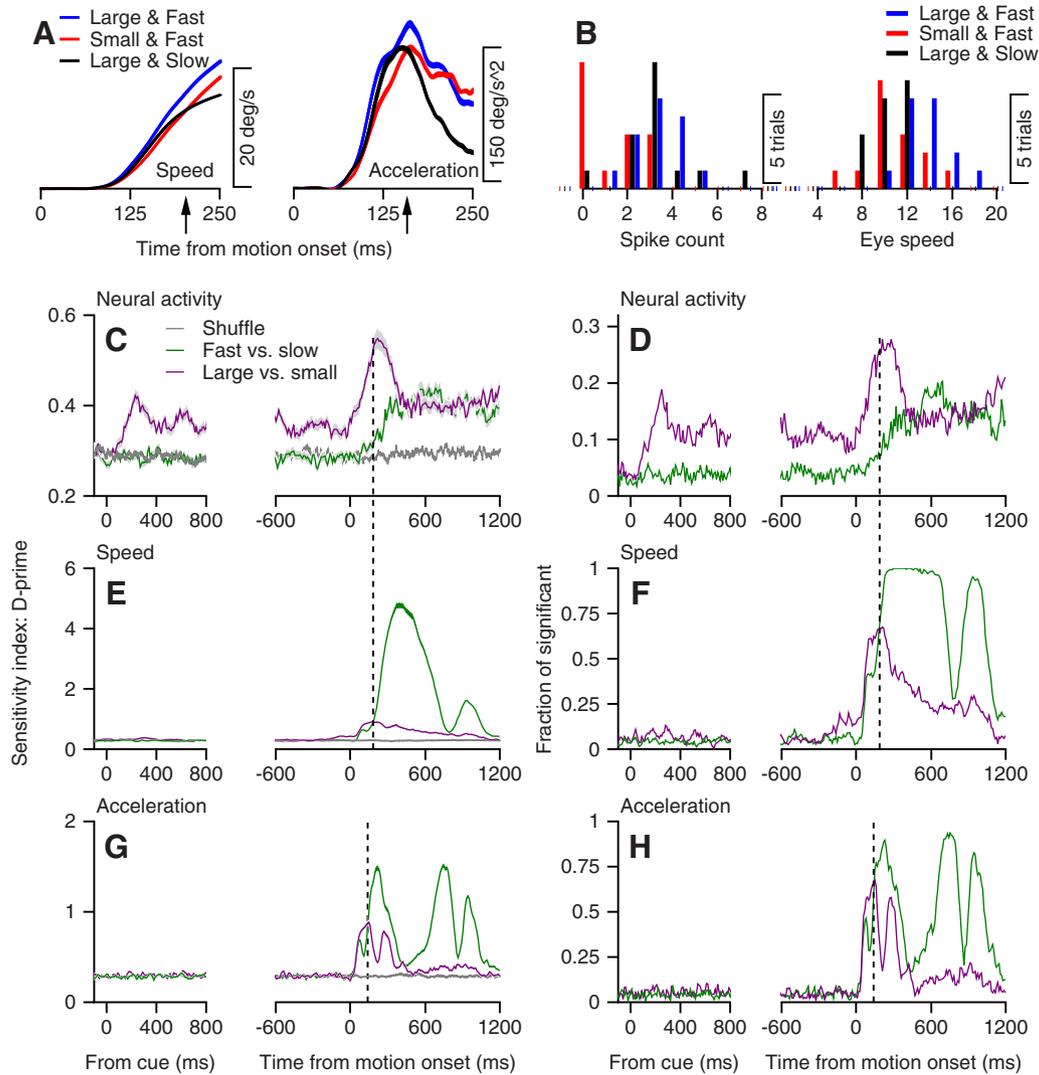
Overall, eye speed sensitivity was larger when calculated in terms of the reward conditions than the target speed conditions. Figure 2D compares the eye speed sensitivity when calculated with either the large and small rewards or the fast and slow target



**Figure 2.** Behavior and neural modulation by target speed and reward size. **A, B.** The difference between the average neural response (**A**) and eye speed (**B**) for the large and small reward conditions (horizontal) versus the fast and slow conditions (vertical). Each dot represents a single neuron (**A**) and the corresponding behavior (**B**). **C.** Example of the relationship between eye speed and firing rate for the large and small reward (solid line) and fast and slow target motion (dashed line). Symbols show the average eye speed (horizontal axis) and the average firing rate (vertical axis) in the 250 ms after target motion onset. **D.** Comparison of eye speed sensitivity calculated for the different reward conditions (horizontal) and different target speed conditions (vertical). Each dot represents the responses of a single neuron in trials with the same target direction condition. Colors in **A** and **D** indicate whether the neuron discriminated significantly ( $p < 0.05$ , Wilcoxon rank sum test) between reward conditions only (blue), target speed conditions only (red), both conditions (black), or neither condition (gray). For the speed sensitivity analysis (**D**), only behavioral sessions in which the behavior differed significantly between large and small reward conditions and between fast and slow target conditions were used.

speeds. Neurons tended to fall under the equality line ( $p < 0.01$ , signed rank test), indicating that the sensitivity was larger when calculated in the reward conditions. Thus, when we directly controlled for changes in behavior, we also found that reward modulated the activity beyond that which would be expected by target speed.

We designed the experiments, limited the analyses, and performed controls to ensure the validity of our conclusions. Very small behavioral differences (Fig. 2B, horizontal or vertical values) could lead to inflated sensitivity values. Therefore, to calculate the speed sensitivity, we used only cells that were recorded in behavior sessions in which the behavior was significantly different between the reward conditions and between the target speed conditions ( $p < 0.05$ , Wilcoxon rank sum test). We also confirmed that using a more stringent criterion ( $p < 0.01$ ) or setting the criterion on the magnitude ( $> 1.5^\circ/s$ ) did not alter any of our conclusions. Smaller differences between the behavior (e.g., smaller range in the horizontal vs vertical axis in Fig. 2A) might lead to larger slopes due to the noise in the estimation of the firing rate. The comparable effect of reward size and target speed on eye speed (Fig. 2B) indicated that this was not the case for our data set. Finally, to minimize the noise in the estimates of the firing rate, we averaged across time and did not use the time-varying neural and behavior signals (Ono and Mustari, 2009; Joshua et al., 2013). A linear relationship between the time-varying eye speed and firing rate implies a linear relationship between averages. Thus, the context-dependent relationship we found be-



**Figure 3.** Temporal course of encoding of reward and target speed in neural activity and behavior. **A**, Average eye speed and acceleration in the first 250 ms after target onset motion. The arrow under the horizontal axis indicates the time point at which the black and red lines intersected. **B**, The distribution across trials of spike counts (left) and eye speeds (right) from a sample neuron. The bars in the histogram were shifted slightly to enable comparison between conditions. **C, E, G**, The  $D'$  values of the responses of the neurons (**C**), eye speed (**E**), and acceleration (**G**) calculated for the different reward conditions (purple) and target speed conditions (green). The gray line shows the  $D'$  values of the responses that were shuffled across conditions. **D, F, H**, The fractions of neurons (**D**), sessions of eye speed (**F**), and acceleration (**H**) in which reward conditions (purple) and target speed conditions (green) were significantly different ( $p < 0.05$ , Wilcoxon rank sum test).

tween the average eye speed and neural activity implies that the relationship between the time-varying eye speed and firing rate was also context dependent. To sum up, the pattern of reward modulations in firing rate could not be attributed to a fixed relationship between activity in the FEF and eye speed.

**Temporal patterns of reward and target speed modulations in rate and behavior**

So far the analysis has been limited to the first 250 ms and has ignored the temporal patterns of the modulations. However, behavior and neural activity were modulated in time. Immediately after motion onset, the average eye speed and eye acceleration in the large–slow condition were larger than in the small–fast condition (Fig. 3A). As the movement progressed, this pattern inverted, and eye speed and acceleration in the large–slow condition became smaller than in the small–fast condition (Fig. 3A, time marked by arrows). We expected that if neural activity reflected only behavior, the cells should follow the same pattern as the behavior.

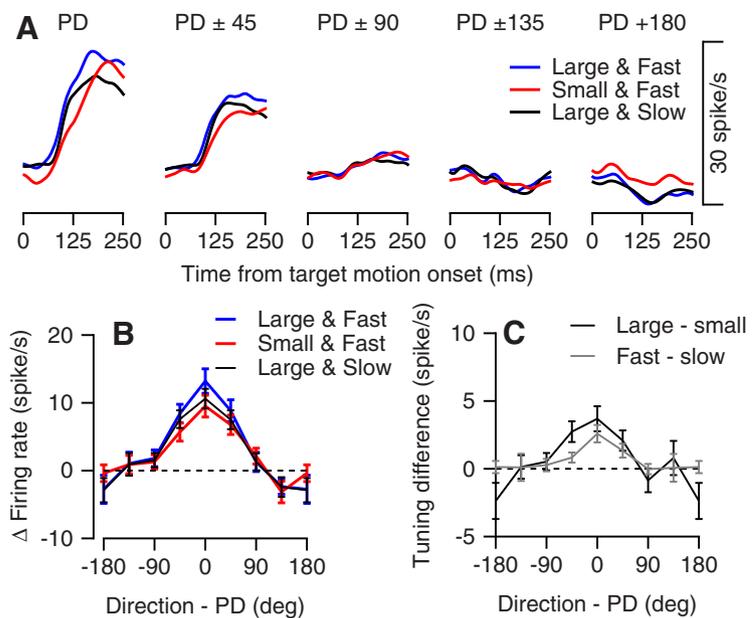
To study the temporal pattern of modulation with a finer time resolution and across all the trials, we used a sliding window of 100 ms to bin both behavior and the firing rate. For each cell and each time bin, we obtained a trial-by-trial distribution of neural and behavioral responses. We then quantified the difference between the target speeds or reward conditions by calculating  $D'$  (see Materials and Methods) or by testing for a significant difference using a nonparametric test ( $p < 0.05$ , Wilcoxon rank sum test). Figure 3B shows the distribution of spike counts and eye speeds for the sample neuron we present in Figure 1D (left) in the time window from 150 to 250 ms after motion onset. The  $D'$  value for the target speed quantifies the difference between the distributions depicted in blue and black, and the  $D'$  value for the reward size measures the difference between the distributions depicted in blue and red. In this specific neuron at this specific time window, the neural activity strongly encoded the reward ( $D' = 1.6$ ;  $p < 0.001$ , Wilcoxon rank sum test) but not the target speed ( $D' = 0$ ; i.e., averages are equal;  $p = 0.77$ ). The behavior, on the other hand, discriminated between speed conditions ( $D' =$

1.58;  $p < 0.01$ ) and between reward conditions ( $D' = 1.13$ ;  $p < 0.01$ ). We used the same measures for single cells and single behavioral sessions to obtain the same metrics for neural activity and behavior. Therefore, any systematic differences observed between neural activity and behavior indicate that behavior cannot simply reflect neural activity.

We applied  $D'$  and significant differences analyses for all the neurons in our data set both before and after the onset of target motion. Figure 3C–H shows the results of the population measures. As has been reported for saccadic movements (Roesch and Olson, 2003, 2005; Ding and Hikosaka, 2006) FEF neurons encoded the expectation of reward even before target motion onset. These modulations are also akin to FEF activity that precedes movement (Bruce and Goldberg, 1985) and have been linked to the allocation of spatial attention (Moore and Fallah, 2001; Gregoriou et al., 2012). After the onset of the color target and before motion onset, the FEF neurons encoded the reward condition, as indicated by the increase in  $D'$  and the number of significant neurons (Fig. 3C,D, purple lines). By contrast, the eye speed during this epoch mostly did not discriminate between the reward conditions before motion onset (Fig. 3E,F, with a minor exception just before target motion onset). Note that before motion onset, the target speed conditions were identical, thus precluding discrimination between target speed conditions by neurons or behavior. Hence, the different target speed conditions (Fig. 3, green lines up to motion onset) and data shuffling (Fig. 3, gray lines) provided baselines for  $D'$  when the tested conditions were identical.

After target motion onset and subsequent eye movement, neurons encoded both the reward size and target speed. The neurons encoded the reward significantly more strongly than the target speed for 359 ms after target motion onset (Fig. 3C, purple vs green;  $p < 0.05$ , signed rank test). This extends the results shown in Figure 2A to a finer temporal resolution (100 ms sliding window; aligned in Fig. 3 to the center bin). The analysis shows that even in time bins that were long after eye movement onset (100 ms after target motion; Fig. 3A,B), the encoding of reward size was stronger than the encoding of target speed. Therefore, the stronger encoding of reward we found in the coarser time resolution (Fig. 2A) was not only the result of reward modulations in the preparatory activity.

After motion onset, eye speed also discriminated between reward conditions and between target speed conditions (Fig. 3E,F). Just as we found for the average in behavior (Fig. 3A), immediately after motion onset, the eye speed differentiated more strongly between reward conditions than between target speed conditions (Fig. 3E,F, values to the left of the dotted lines). Very rapidly, this pattern switched, and the eye speed discriminated better between target speed conditions than between reward conditions. The dashed vertical lines in Figure 3 mark this time point (184 ms for  $D'$ , 188 for the number of significant neurons). After this point, the  $D'$  value and the number of signifi-



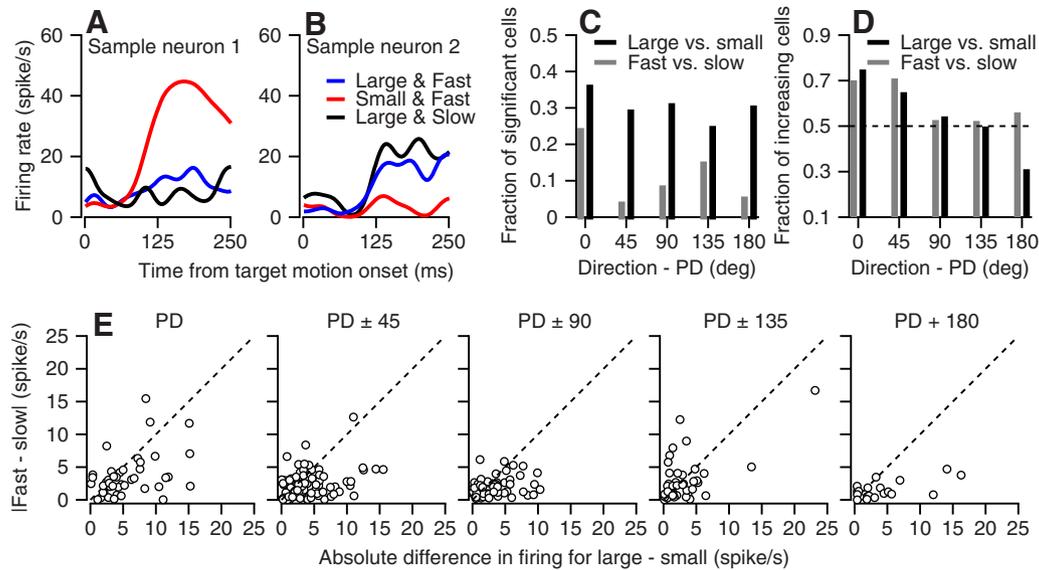
**Figure 4.** The effects of reward and target speed on population directional tuning. **A**, Population average of the neural activity for the large–fast, (blue) small–fast (red), and large–slow conditions. Different plots correspond to the differences in the direction of motion from the preferred direction of the neurons. The baseline firing rate was defined as the average rate at motion onset across all recorded conditions and was subtracted individually for each neuron. **B**, The population tuning curve for large–fast (blue), small–fast (red), and large–slow (black) conditions. **C**, The difference in tuning curves between the large and small rewards (black) and fast and slow target motion (gray). Error bars show the SEM.

cant sessions became substantially larger for the target speed conditions than for the reward conditions.

The pattern of neural activity did not coincide with the pattern of behavior. The larger encoding of reward continued for slightly less than 175 ms after the period in which the eye speed discriminated better between reward conditions than between target speed conditions (Fig. 3C,D, values to the right of the dashed line). This pattern of larger encodings of reward for an additional 175 ms after behavioral change was highly significant ( $p < 0.001$ , bootstrap; see Materials and Methods). These values also contrasted sharply with the 20–40 ms latency between electrical stimulations in the FEF and smooth pursuit movements (Gottlieb et al., 1993; Tanaka and Lisberger, 2001). Later in the trial, although eye speed discriminated very strongly between target speed but not between reward conditions, the encoding of reward and target speed was similar. Using the eye acceleration instead of eye speed for the behavioral measurement did not alter the findings (Fig. 3G,H). Overall, we thus found that neurons in the FEF encoded the reward condition before and during movement. The temporal patterns of the rate modulations did not match the temporal patterns in the behavioral modulations.

#### Interaction of encoding of reward and target speed with directional tuning

To probe how population rate modulation interacted with directional tuning, we examined activity in different directions relative to the preferred direction. In each recording session, we recorded data from trials in which the targets moved in one of two orthogonal directions (Fig. 1D). In each session, we recorded with up to five channels and tuning could differ between neurons on different recording channels. Therefore, many of the neurons included in the analysis were recorded with targets that were not moving in their preferred direction. This allowed us to study the population direction tuning for the different task conditions.



**Figure 5.** Reward modulations are attenuated when responses are averaged according to the tuning direction. **A, B**, Examples of the activity of two neurons when the target moved in a direction perpendicular their preferred direction. **C, D**, The fraction of neurons that significantly encoded reward or target speed (**C**) and the fraction of neurons in which reward or target speed was encoded by increases in the firing rate (**D**). Horizontal values show the difference in direction from the preferred direction. Black and gray bars show the data for large–fast versus small–fast and large–fast versus large–slow conditions. **E**, The absolute difference between the average neural responses for the large and small reward conditions (horizontal) versus the fast and slow conditions (vertical). Different plots correspond to different directions of motion from the preferred direction of the neurons. Each dot represents a single neuron.

The average peristimulus time histogram of the response of the neurons recorded in the preferred direction and rotated by 45° was the largest for the large–fast condition (blue traces in Fig. 4A). In these directions, the pattern of the average response in the large–slow and small–fast conditions was comparable to the pattern we found for behavior (Fig. 3A). Initially, the responses in the large–fast and large–slow conditions were similar, but as of 125–150 ms the rate was smaller for the large–slow target condition (Fig. 4A, blue vs black lines in PD and PD ± 45). As we found for behavior (Fig. 3A) in the small–fast condition, the neural response was initially smaller than the large–slow target condition but exhibited a reversal toward the end of the analysis period, as indicated by the intersection of the red and black lines in Figure 4A (PD and PD ± 45). In directions rotated 90° and 135° from the preferred direction, the population PSTHs in all conditions were very similar. In the antipreferred direction (180°), there was a slight trend in the neurons to respond less to conditions with the large versus small reward (Fig. 4A, blue and black vs red in PD + 180). Note that the resemblances of the patterns of the average PSTH to behavior were apparently inconsistent with the results of the  $D'$  analysis. This apparent inconsistency is a result of pooling according to the preferred direction; we return to this point below.

To quantify tuning at the population level for the large–fast, small–fast, and large–slow conditions, we calculated the population direction tuning curves. For each neuron, we averaged the firing rate across the 250 ms after target motion onset for the different conditions and subtracted the average across all conditions of the firing rate before target motion. We then averaged the responses of the neurons that were recorded at that specific difference from the preferred direction for each direction separately. This resulted in a population-directional tuning curve (Fig. 4B) for the different conditions. To highlight the effects of the reward and target speed on the tuning curve, we calculated the differences between the large–fast condition and either the small–fast reward or the large–slow target condition (Fig. 4C). As found for the time-varying response (Fig. 4A) in the direction close to the

preferred direction, reward increased the firing rate of the cell. In the antipreferred direction, modulation was in the opposite direction. The shape of this modulation thus supported the existence of a multiplicative effect of the reward and target speed on the population direction tuning. This is similar to how orientation tuning curves are potentiated by top-down information in sensory areas such as V4 (McAdams and Maunsell, 1999), but differs from the anterior cingulate cortex in which reward was reported to add a constant modulation to the directional tuning of the populations (Hayden and Platt, 2010).

### The population vector averages out reward modulations and fits neural activity to behavior

A more fine-grained examination of the direction-dependent average PSTHs suggested similarities between the patterns of responses and behavior (Figs. 3A, 4A). For example, after 250 ms, the encoding of reward at the population average was weaker than the encoding of target speed in that the average response in the large–fast condition was closer to that in the small–fast than to that in the large–slow condition (Fig. 4A, PD and PD ± 45). This appears to be inconsistent with the stronger encoding of reward versus target speed that we found at the single cell level that continued beyond the first 250 ms (Fig. 3C,D). This can be explained by the fact that grouping the cells by the preferred direction results in averaging out some of the rate modulations.

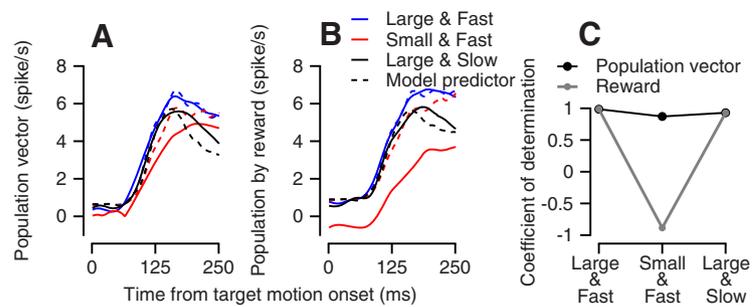
The responses of the cells shown in Figure 5, A and B, exemplify how opposite modulations are averaged out. Both neurons were recorded in directions perpendicular to their preferred direction. The similarity between the blue and black lines indicates that in both cells, the target speed was encoded weakly. The difference between the blue and red lines indicates that in both cells, the encoding of the reward was stronger. However, the sign of the modulations was reversed; in cell 1, the response in the small reward was larger than the response to the large reward, whereas in cell 2, we found the opposite relationship. Therefore, averaging the activity of these two cells would reduce the extent of the reward modulations.

In all the different directions, more neurons significantly encoded the reward than the target speed (Fig. 5C). In all but one condition (PD  $\pm$  135), reward modulations were significantly larger than the speed modulations (Fig. 5E;  $p < 0.05$ , signed rank test). However, this pattern of larger and more frequent encoding of reward was averaged out in some cases due to opposite rate modulations (Fig. 5D). In the PD and PD  $\pm$  45 directions, most of the neurons increased their firing rate in the large–fast condition compared to the other conditions; therefore, in these directions, reward and target speed were also encoded at the population level (Fig. 4). In directions PD  $\pm$  90 and PD  $\pm$  135, the numbers of increasing and decreasing cells were almost identical (Fig. 5D, values close to the dashed line); therefore, at the population level, we did not find a modulation related to target speed or reward. In the direction 180° from the PD, many neurons decreased their firing rate in the large–fast compared to the small–fast condition, but close to 50% of the cells decreased their rate in the large–fast versus large–slow condition. Therefore, at the population level, reward but not target speed was encoded in this direction (Fig. 4).

To demonstrate that pooling the population response by the preferred direction resulted in modulations that were closer to behavior, we calculated the temporal pattern of the amplitude of the population vector from the average PSTH (see Materials and Methods). This analysis yielded time-varying population responses for each experimental condition (Fig. 6A, solid lines). We then fit the population vector to a linear weighted sum of the average eye acceleration and speed (Shidara et al., 1993; Ono and Mustari, 2009; Joshua and Lisberger, 2014) as follows:  $Population\ response(t) = a_0 + a_1\dot{E}(t - \Delta t) + a_2\ddot{E}(t - \Delta t)$ , where  $\dot{E}$  is speed,  $\ddot{E}$  is acceleration,  $a_i$  are the behavior sensitivity parameters, and  $\Delta t$  indicates how much time the eye movement averages needed to be shifted to optimize the fit to the average firing rate. We did not use the eye position to prevent overfitting. The position had very little effect on the firing rate at movement initiation since eye position changes were small during the first 250 ms after target motion ( $\sim 1.5^\circ$ ) and the position sensitivity of neurons in the FEF is small (Ono and Mustari, 2009).

To test whether behavior could account for the population modulations, we fit the model to the large–fast condition and tested how well the model predicted the population response in the large–slow and small–fast conditions. We found that the model provided a very good fit for these conditions (Fig. 6A). For the small–fast and large–slow conditions, 87 and 93% of the variability in the population response could be explained by the fit (Fig. 6C). The difference between predictors and the actual traces was a result of a small shift in baseline activity. Note that we used the fit in the large–fast condition and tested in the other conditions; thus, overfitting would decrease values rather than increase them. This suggests that behavior can explain a large proportion of the variability in the population response.

Next, we tested whether averaging the responses based on the preferred direction was indeed critical for improving the fit between behavior and neural activity. Instead of constructing the population responses from the PSTH sorted by the preferred direction, we averaged the responses of the cells by the reward modulation (Fig. 6B). We did so by flipping the responses of the



**Figure 6.** Match to behavior of the neural activity is improved when responses are averaged according to the direction tuning. **A, B**, The amplitude of the population vector (**A**) and the average response calculated by the sign of the reward modulations (**B**). The blue, red, and black lines correspond to the large–fast, small–fast, and large–slow conditions. Each dashed line shows the predictor of the linear model that was calculated based on the large–fast condition. **C**, The coefficient of determination of the model predictor (dashed lines in **A, B**) and neural response (solid lines in **A** and **B**). Black and gray lines correspond to the population vector (**A**) and the average population response (**B**). Values could be below zero, because the fit could be worse than the average of the trace. (In linear regression analysis, the values are bounded between 0 and 1.)

cells in which the response decreased with reward size (Fig. 5A, cell 1). Using this procedure could bias the population averages due to noise in the estimates of the sign of the modulations. Therefore, we used half of the trials from each cell to calculate the sign of the reward modulation and the other half to calculate the PSTH. We then fit the linear model to the response in the large–fast condition and compared the fit to that in the other conditions. The model failed to predict the response for the small–fast condition. The response in the small–fast condition was substantially smaller than the behavioral prediction (Fig. 6B, red solid vs dashed line). The coefficient of determination was below zero, indicating that taking a constant value for the fit would be better than the fit of the linear model (Fig. 6C; see Materials and Methods). In this case, the failure stemmed mostly but not solely from the large baseline shift between trials with small and large rewards. Subtracting the baseline shift resulted in a better fit between the model predictor and the response to the small–fast condition (coefficient of determination, 0.7).

Thus, overall, reward appeared to be encoded more strongly than target speed across the population. However, grouping the activity with respect to the preferred direction of the cells averaged out some of the reward modulations, leading to population activity that matched the behavior quite accurately.

## Discussion

The relationship between the firing rate of FEF neurons and eye speed depends on the context. The eye speed sensitivity of the neurons was larger for the reward conditions than for the target speed conditions (Figs. 2, 3). What best explains the fact that FEF activity encodes reward more strongly than target speed even when taking behavioral modulations into account? Below we discuss the implications of the results with regard to the ways in which reward is encoded and movement is decoded from neural activity.

### Encoding reward and decoding movement from the FEF

The key to understanding how reward drives behavior and is modulated beyond behavior lies in the effects of reward on directional tuning. We found that reward potentiated the directional tuning of the population. This result indicates that a plausible decoder that utilizes directional selectivity, such as the vector average of the population (Georgopoulos et al., 1986; Lisberger and Ferrera, 1997), would read out the FEF activity to drive movement. The outcome of a population vector of the FEF activity resulted in modulations that were closely related to behavior

(Fig. 6A). The population vector averaged out some of the reward modulations, and the remaining modulations drove the behavioral differences. This contrasts with decoders that averaged activity based on reward modulations that did not follow the pattern of eye kinematics (Fig. 6B,C).

The anatomy and physiology of the pursuit system suggest a way in which the downstream structures could implement the transformation that averages out reward. Neurons are known to be directionally tuned across the pursuit system. Neurons with the same directional tuning are often anatomically localized (Lisberger and Fuchs, 1978; Albright, 1984; Gottlieb et al., 1994). Electrical stimulation in different locations drives movement in different directions (Gottlieb et al., 1994; Lisberger et al., 1994; Groh et al., 1997; Krauzlis and Miles, 1998). This interarea principle of organization suggests that transformations between areas are organized to maintain the preferred direction. In fact, models of the pursuit system often assume either implicitly or explicitly that cells with the same directional selectivity are connected (Robinson et al., 1986; Yang et al., 2012).

By simply complying with direction-selective connectivity, the downstream structures can average out the reward modulations. For example, in the results depicted in Figure 5, the downstream population of neurons that is randomly connected to neurons that are perpendicular to the preferred direction of motion would not be modulated by reward. When the target moves in the preferred direction (Fig. 4A, left), the modulations of reward are not averaged out, and the reward signal propagates to downstream neurons but in ways that correspond to behavioral effects. We thus posit that downstream, such as in the cerebellum and brainstem, reward activity is fully transformed into motor activity.

Other decoders could use other dimensions of neural population activity and better decode reward size. Akin to other prefrontal regions (Miller and Cohen, 2001; Machens, 2010; Mante et al., 2013; Rigotti et al., 2013), the FEF encodes a mixture of task-relevant parameters such as target speed and reward size. A critical question when considering mixed representations is how they are exploited by the decoder. We found that the movement decoder that is organized by the preferred direction of the cells averages out some of the reward modulation.

Other decoders that are not organized by the preferred direction might pool data according to the motivationally potent dimensions to extract the motivational value of the stimulus. In studies that have compared the size of reward modulations in the motor structures across different frontal cortex populations, the motor structures emerged as the areas with the strongest reward modulations (Roesch and Olson, 2003). The presence of reward modulation beyond behavior suggests that the FEF itself might encode the abstract motivational facets of movement and not only the motor properties of the movement.

Decoders that use the reward information could also potentiate sensory processing. It is possible that the reward related activity is also the neural source for attention since reward and attention are strongly, and perhaps necessarily, linked (Maunsell, 2004). The modulations before target motion (Fig. 3C,D) were found to be similar to the activity of FEF neurons during attention tasks (Thompson et al., 2005; Gregoriou et al., 2012). It is possible, for example, that decoders that are organized by the spatial selectivity of neurons might use reward modulation to drive spatial attention (Moore and Armstrong, 2003; Noudoost and Moore, 2011). We did not attempt to control for attention in our task; therefore, further work is needed to test whether and how activity in the FEF is related to spatial and feature attention

during pursuit (van Donkelaar, 1999; Garbutt and Lisberger, 2006; Lovejoy et al., 2009; Spering and Carrasco, 2012).

### Other possible explanations for context-dependent eye speed sensitivity

Our explanation assumes that activity is decoded by the preferred direction of the neurons. Beyond this, we do not assume specialized connectivity or functions for the neurons in the FEF. Other mechanisms that assume more specific connectivity could also explain mismatches between behavior and neural activity. One possibility is that not all neurons in the population drive behavior. For instance, neurons projecting downstream could be more closely related to behavior such that their firing rate is independent of the context. Neurons that project from the FEF downstream to the reticular formation tend to respond more strongly at movement initiation and thus are more sensitive to eye acceleration (Ono and Mustari, 2009). We, however, did not find any relationship between the pattern of the response (initiation vs steady state) and whether the cells were context dependent (differences in sensitivity for reward and target speed).

Another possible explanation for modulations beyond behavior is that the FEF controls the strength, or gain, of visuomotor transmission (Tanaka and Lisberger, 2001). The gain hypothesis does not require the neurons to strictly encode the details of movement since they do not provide the direct drive for movement. In this case, the neurons that encode reward beyond behavior could represent the larger gain of the movement in the trial with larger reward. However, these hypotheses are not mutually exclusive in that gain could be implemented with modulations that are proportional to behavior.

### Other experiments that have studied how reward interacts with FEF activity

The interaction between reward information and eye movements has been studied in the saccade movement system. Rewarded saccades have shorter latency and larger peak velocities and are less variable than nonrewarded saccades (Takikawa et al., 2002; Reppert et al., 2015). The size or presence of the reward modulates the activity of FEF neurons in preparation for and during saccades (Roesch and Olson, 2003, 2005; Ding and Hikosaka, 2006; Glaser et al., 2016). In one study, the saccade-by-saccade variability was used to show that in natural scene searches, FEF activity is enhanced beyond the eye kinematics (Glaser et al., 2016).

Our study expands these works to the pursuit system in two important ways. First, we controlled the kinematics of the movement in the pursuit system experimentally by controlling target motion. Thus, we compared the target speed and reward drive for pursuit directly to show that the relationship between FEF activity and behavior is context dependent. Second, we characterized the directional tuning curve of the cells and showed how it interacts with reward modulation. The effect on the tuning curve suggests a plausible decoding framework that may explain how the reward both is represented beyond behavior and can drive behavior.

### References

- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130. [CrossRef Medline](#)
- Bizzi E (1968) Discharge of frontal eye field neurons during saccadic and following eye movements in unanesthetized monkeys. *Exp Brain Res* 6:69–80. [Medline](#)
- Bruce CJ, Goldberg ME (1985) Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53:603–635. [CrossRef Medline](#)
- Ding L, Hikosaka O (2006) Comparison of reward modulation in the fron-

- tal eye field and caudate of the macaque. *J Neurosci* 26:6695–6703. [CrossRef Medline](#)
- Garbutt S, Lisberger SG (2006) Directional cuing of target choice in human smooth pursuit eye movements. *J Neurosci* 26:12479–12486. [CrossRef Medline](#)
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419. [CrossRef Medline](#)
- Glaser JI, Wood DK, Lawlor PN, Ramkumar P, Kording KP, Seiringer MA (2016) Role of expected reward in frontal eye field during natural scene search. *J Neurophysiol* 116:645–657. [CrossRef Medline](#)
- Gottlieb JP, Bruce CJ, MacAvoy MG (1993) Smooth eye movements elicited by microstimulation in the primate frontal eye field. *J Neurophysiol* 69:786–799. [CrossRef Medline](#)
- Gottlieb JP, MacAvoy MG, Bruce CJ (1994) Neural responses related to smooth-pursuit eye movements and their correspondence with electrically elicited smooth eye movements in the primate frontal eye field. *J Neurophysiol* 72:1634–1653. [CrossRef Medline](#)
- Gregoriou GG, Gotts SJ, Desimone R (2012) Cell-type-specific synchronization of neural activity in FEF with V4 during attention. *Neuron* 73:581–594. [CrossRef Medline](#)
- Groh JM, Born RT, Newsome WT (1997) How is a sensory map read out? effects of microstimulation in visual area MT on saccades and smooth pursuit eye movements. *J Neurosci* 17:4312–4330. [CrossRef Medline](#)
- Hayden BY, Platt ML (2010) Neurons in anterior cingulate cortex multiplex information about reward and action. *J Neurosci* 30:3339–3346. [CrossRef Medline](#)
- Joshua M, Lisberger SG (2012) Reward action in the initiation of smooth pursuit eye movements. *J Neurosci* 32:2856–2867. [CrossRef Medline](#)
- Joshua M, Lisberger SG (2014) A framework for using signal, noise, and variation to determine whether the brain controls movement synergies or single muscles. *J Neurophysiol* 111:733–745. [CrossRef Medline](#)
- Joshua M, Adler A, Mitelman R, Vaadia E, Bergman H (2008) Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials. *J Neurosci* 28:11673–11684. [CrossRef Medline](#)
- Joshua M, Medina JF, Lisberger SG (2013) Diversity of neural responses in the brainstem during smooth pursuit eye movements constrains the circuit mechanisms of neural integration. *J Neurosci* 33:6633–6647. [CrossRef Medline](#)
- Krauzlis RJ (2004) Recasting the smooth pursuit eye movement system. *J Neurophysiol* 91:591–603. [CrossRef Medline](#)
- Krauzlis RJ, Miles FA (1998) Role of the oculomotor vermis in generating pursuit and saccades: effects of microstimulation. *J Neurophysiol* 80:2046–2062. [CrossRef Medline](#)
- Lisberger SG (2010) Visual guidance of smooth-pursuit eye movements: sensation, action, and what happens in between. *Neuron* 66:477–491. [CrossRef Medline](#)
- Lisberger SG, Ferrera VP (1997) Vector averaging for smooth pursuit eye movements initiated by two moving targets in monkeys. *J Neurosci* 17:7490–7502. [CrossRef Medline](#)
- Lisberger SG, Fuchs AF (1978) Role of primate flocculus during rapid behavioral modification of vestibuloocular reflex. I. Purkinje cell activity during visually guided horizontal smooth-pursuit eye movements and passive head rotation. *J Neurophysiol* 41:733–763. [CrossRef Medline](#)
- Lisberger SG, Pavelko TA, Broussard DM (1994) Responses during eye movements of brain stem neurons that receive monosynaptic inhibition from the flocculus and ventral paraflocculus in monkeys. *J Neurophysiol* 72:909–927. [CrossRef Medline](#)
- Lovejoy LP, Fowley GA, Krauzlis RJ (2009) Spatial allocation of attention during smooth pursuit eye movements. *Vision Res* 49:1275–1285. [CrossRef Medline](#)
- Machens CK (2010) Demixing population activity in higher cortical areas. *Front Comput Neurosci* 4:126. [Medline](#)
- Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503:78–84. [CrossRef Medline](#)
- Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 8:261–265. [CrossRef Medline](#)
- McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441. [CrossRef Medline](#)
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202. [CrossRef Medline](#)
- Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421:370–373. [CrossRef Medline](#)
- Moore T, Fallah M (2001) Control of eye movements and spatial attention. *Proc Natl Acad Sci U S A* 98:1273–1276. [CrossRef Medline](#)
- Noudoost B, Moore T (2011) Control of visual cortical signals by prefrontal dopamine. *Nature* 474:372–375. [CrossRef Medline](#)
- Ono S, Mustari MJ (2009) Smooth pursuit-related information processing in frontal eye field neurons that project to the NRTP. *Cereb Cortex* 19:1186–1197. [CrossRef Medline](#)
- Paton JJ, Belova MA, Morrison SE, Salzman CD (2006) The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* 439:865–870. [CrossRef Medline](#)
- Raghavan RT, Joshua M (2017) Dissecting patterns of preparatory activity in the frontal eye fields during pursuit target selection. *J Neurophysiol* 18:2216–2231. [Medline](#)
- Rashbass C (1961) The relationship between saccadic and smooth tracking eye movements. *J Physiol* 159:326–338. [CrossRef Medline](#)
- Reppert TR, Lempert KM, Glimcher PW, Shadmehr R (2015) Modulation of saccade vigor during value-based decision making. *J Neurosci* 35:15369–15378. [CrossRef Medline](#)
- Retailleau A, Morris G (2018) Spatial rule learning and corresponding CA1 place cell reorientation depend on local dopamine release. *Curr Biol* 28:836–846.e4. [CrossRef Medline](#)
- Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. *Nature* 497:585–590. [CrossRef Medline](#)
- Robinson DA (1981) The use of control systems analysis in the neurophysiology of eye movements. *Annu Rev Neurosci* 4:463–503. [CrossRef Medline](#)
- Robinson DA, Gordon JL, Gordon SE (1986) A model of the smooth pursuit eye movement system. *Biol Cybern* 55:43–57. [CrossRef Medline](#)
- Roesch MR, Olson CR (2003) Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *J Neurophysiol* 90:1766–1789. [CrossRef Medline](#)
- Roesch MR, Olson CR (2005) Neuronal activity dependent on anticipated and elapsed delay in macaque prefrontal cortex, frontal and supplementary eye fields, and premotor cortex. *J Neurophysiol* 94:1469–1497. [CrossRef Medline](#)
- Shidara M, Kawano K, Gomi H, Kawato M (1993) Inverse-dynamics model eye movement control by Purkinje cells in the cerebellum. *Nature* 365:50–52. [CrossRef Medline](#)
- Sparks DL (2002) The brainstem control of saccadic eye movements. *Nat Rev Neurosci* 3:952–964. [CrossRef Medline](#)
- Spering M, Carrasco M (2012) Similar effects of feature-based attention on motion perception and pursuit eye movements at different levels of awareness. *J Neurosci* 32:7594–7601. [CrossRef Medline](#)
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291. [CrossRef Medline](#)
- Tanaka M, Fukushima K (1998) Neuronal responses related to smooth pursuit eye movements in the periarculate cortical area of monkeys. *J Neurophysiol* 80:28–47. [CrossRef Medline](#)
- Tanaka M, Lisberger SG (2001) Regulation of the gain of visually guided smooth-pursuit eye movements by frontal cortex. *Nature* 409:191–194. [CrossRef Medline](#)
- Thompson KG, Bischof KL, Sato TR (2005) Neuronal basis of covert spatial attention in the frontal eye field. *J Neurosci* 25:9479–9487. [CrossRef Medline](#)
- van Donkelaar P (1999) Spatiotemporal modulation of attention during smooth pursuit eye movements. *Neuroreport* 10:2523–2526. [CrossRef Medline](#)
- Vickery TJ, Chun MM, Lee D (2011) Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron* 72:166–177. [CrossRef Medline](#)
- Wagner MJ, Kim TH, Savall J, Schnitzer MJ, Luo L (2017) Cerebellar granule cells encode the expectation of reward. *Nature* 544:96–100. [CrossRef Medline](#)
- Yang J, Lee J, Lisberger SG (2012) The interaction of Bayesian priors and sensory data and its neural circuit implementation in visually guided movement. *J Neurosci* 32:17632–17645. [CrossRef Medline](#)