ARTICLE IN PRESS

Nature Communications

https://doi.org/10.1038/s41467-026-71046-2

Article in Press

Enhanced reward coding and condition-independent dynamics in optogenetically identified corticostriatal neurons in monkeys

Received: 23 April 2025

Accepted: 9 March 2026

Adi Hovav-Lixenberg, Yirat Henshke, Tirzah Kreisel, Eran Lottem & Mati Joshua

Cite this article as: Hovav-Lixenberg, A., Henshke, Y., Kreisel, T. *et al.* Enhanced reward coding and condition-independent dynamics in optogenetically identified corticostriatal neurons in monkeys. *Nat Commun* (2026). https://doi.org/10.1038/s41467-026-71046-2

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

ARTICLE IN PRESS

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

Enhanced reward coding and condition-independent dynamics in optogenetically identified corticostriatal neurons in monkeys

Adi Hovav-Lixenberg[1], Yirat Henshke[1], Tirzah Kreisel[1], Eran Lottem[1] and Mati Joshua[1]

1. Edmond and Lily Safra Center for Brain Sciences (ELSC), The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Abstract

The basal ganglia integrate cortical inputs with dopaminergic signals to potentiate and select actions. The reward-related activity of dopamine neurons is well-studied, but the coding properties of cortical inputs to the basal ganglia remain largely unknown. We examined the activity of neurons in the frontal eye field of monkeys that were optogenetically identified as projecting to the basal ganglia. We found that the projecting neurons contained information about expected rewards and selected actions. The reward-related signal and modulations independent of task condition were stronger in optogenetically identified projecting neurons than in other neurons in the same area. These findings indicate that reward, choice, and sensorimotor information are already integrated into the cortical inputs to the basal ganglia, suggesting that the basal ganglia network integrates reward from both cortical and dopaminergic inputs rather than relying on a dopaminergic source alone.

Corresponding author: Mati Joshua, mati.joshua@mail.huji.ac.il

Introduction

One of the widespread assumptions about basal ganglia function is that the striatum integrates reward information from dopaminergic inputs with cortical information to potentiate movement and select actions[1–3]. While studies have extensively characterized dopaminergic inputs to the striatum[4–11], the content of the cortical input to the basal ganglia remains less well understood[12,13]. Cortical neurons have been shown to integrate reward, choice and sensorimotor information[14–19], but neurons projecting to the basal ganglia may not contain reward-related signals. Further, other sensory and movement signals transmitted from the cortex to the basal ganglia aside from reward-related signals remain poorly characterized.

To better understand the content of cortico-striatal signaling, we focused on the monkey eye movement system. In this system, the frontal eye field (FEF) projects to the caudate nucleus in the striatum[20,21], and both areas are causally linked to eye movements[22,23]. The FEF is a central hub for processing visual and movement signals that drive eye movements[18,24].

We aimed to functionally characterize the signal transmitted from the FEF to the caudate in eye movement tasks that included reward and choice manipulations. In earlier research protocols, antidromic electrical stimulation, where axons of neurons are stimulated at the target site and detected at the source has been used to identify input[12,25]. However, the yield of this method is often low, as it is difficult to stimulate and record from the same neuron at two sites simultaneously. Viruses that code channelrhodopsin and are transported retrogradely from the injection site (e.g., the caudate) to the source (e.g., FEF) provide an effective alternative[26,27]. Using this method, input neurons can be identified by their responses to light stimulation at the recording site[28–30], a method termed opto-tagging.

We opto-tagged neurons projecting from the FEF to the caudate and recorded their activity during eye movement tasks where reward size was manipulated. We found that the neurons that were opto-tagged encoded strong reward and choice information and exhibited more pronounced modulations in activity averaged across task conditions (termed condition-independent modulations) than other neurons, thus implying that the inputs to the basal ganglia have already integrated reward, choice and sensorimotor information. These findings thus suggest that the basal ganglia rely on a combination of reward signals from multiple sources rather than dopaminergic input alone.

Results

Reward impacts monkeys' behavior in eye movement tasks

Two monkeys were engaged in three eye movement tasks in which we manipulated the reward size (Fig. 1A). In these tasks, the monkeys were presented with color-coded targets that indicated whether they would receive a large or small reward at the end of the trial. The monkeys observed a colored target that appeared in a central (the pursuit task) or a peripheral location (the saccade task). After a variable delay, the monkeys were instructed to shift their gaze to the target (Go). On the saccade trials, the monkeys moved their eyes ballistically to the eccentric location (Fig. 1A and B, left), whereas on the pursuit trials the monkeys smoothly followed a continuously moving target (Fig. 1A and B, middle). The saccade and pursuit movements were directed to one of the four cardinal directions. On the choice task, two targets appeared and the monkeys had to select one (Fig. 1A and B right). At the end of each trial, the monkeys received either a small or large reward based on the color of the single or selected target. We refer to the epoch during which the monkeys observed the cues but were not allowed to move as the *cue epoch*, the epoch during which they moved their eyes to the colored target as the *movement epoch*, and the epoch during which they received the reward as the *outcome epoch*.

The monkeys' behavior indicated that they successfully associated the colored target with the upcoming reward. We analyzed behavior on days when we recorded neural activity, after the monkeys were well-trained on the task. On the saccade tasks, the monkeys moved earlier in the large reward target (Fig. 1C, p=8*10^-5 and 10^-6 for monkeys H and K, Wilcoxon signed-rank test). On the pursuit trials, in one monkey, the eye velocity at movement initiation was consistently faster for the large reward (Fig. 1D, eye velocity at 200-250 ms after target motion onset, p = 8*10^-7 and p =0.98 for monkeys H and K, Wilcoxon signed-rank test). The monkeys' preferences were confirmed on a choice task in which both monkeys almost always selected the target that was associated with the larger reward (Fig. 1E p = 5*10^-7 and p=8*10^-6 for monkeys H and K, Wilcoxon signed-rank test for differences from a chance level of 0.5), demonstrating that the monkeys associated color and reward size. Thus, the monkeys' behavior indicated that they were expecting the upcoming reward, which provided the foundation for studying reward-related signals.

Opto-tagging of neurons in the FEF projecting to the striatum

To identify neurons in the FEF projecting to the basal ganglia, we injected an AAVretro-ChR2 virus[26] into the eye movement areas of the caudate (Fig. 1F and Fig. S1). This virus is transported retrogradely from the axons in the caudate to the cell bodies in the cortex. As a result, in the cortex, only cortico-striatal projection neurons express ChR2. We confirmed the injection location over the course of several days during which we injected a solution with Mn+2 and imaged the injection site with an MRI scan (Fig. S1). We waited four weeks for viral expression and then on each recording day we lowered an optrode with 4 closely spaced recording contacts attached to a fiber for light stimulation (Thomas Recordings) into the FEF (Fig. S2). We searched for neurons that responded consistently with short latencies (see Methods) to brief light stimulations (durations of 5 ms at 10, 20 and 50 Hz). Figure 1G shows an example of the extracellular activity in one contact of the optrodes while we applied the light stimulation. The recorded neuron fired sparsely before the stimulation and very consistently with a very short latency during the stimulation (Fig. 1G, H). Based on the response to stimulation (see Methods) we categorized the neurons into (1) light responding neurons that were highly likely to be cortico-striatal projection neurons, (2) neurons that were recorded in the same location as the responding neurons but did not respond to the stimulation, or their response was inconsistent, and (3) neurons in which none of the neurons in the recording session responded consistently to the stimulation. A response to light was considered strong evidence for the cell to be a cortico-striatal neuron, but a lack of response was harder to interpret. For this reason, we classified the neurons into three groups: *tagged (n=108), untagged (n=266)* and *area-untagged (n=138)*. We performed an offline analysis to confirm the consistency of the extracellular waveforms of the connected neurons across the stimulation and behavioral sessions (Fig. S3 and Methods).

Factoring the response of single neurons to the experimental variables

To study the coding properties of the neurons, we factored the single-neuron activity into components related to direction, reward, and those independent of the task condition. We depict this factorization in Figure 2A-E which shows the response of an example neuron to the appearance of an eccentric color cue in the saccade task. Averaging across either the reward size or direction of the visual cue (and upcoming movement) (Fig. 2C and D) indicated that the neuron was modulated by the direction and the reward size. Averaging across both reward size and direction showed that the neuron exhibited overall temporal modulations Fig. 2E) which we term condition-independent modulations.

To quantify the contribution of these different experimental variables to the overall activity, we calculated the partial ω^2 effect size (ω^2, see Methods)[31–33]. ω^2 is an unbiased estimator of the variability explained by an experimental variable, normalized by the same variance plus a noise term (Fig. 2F). The noise term is an estimator of the variance that is unaccounted for by any variable in the experiment; i.e., the trial-by-trial variability within task conditions. Note that we term this ‘noise’, but it may still contain meaningful trial-by-trial activity. ω^2 is distributed around zero (unbiased) when a neuron does not respond to the experimental variable and has a value of 1 when all the trial-by-trial variability is explained by an experimental variable. Intermediate values indicate how much of the neuron’s activity can be accounted for by a specific variable in comparison to the trial-by-trial variability that cannot be explained by any of the variables.

To calculate the effect size, we used an ANOVA model for each neuron, with reward size, direction of visual cue, and the time bin within the trial as variables (see Methods, e.g., values in Fig. 2C–E). The direction effect size estimates the magnitude of the neuron's directional tuning. In the choice task, we define the direction effect size as the side (e.g., left or right) on which the large-reward target appeared. The reward effect size estimates how different the responses were in the large- versus small-reward condition. The time effect size estimates the overall average modulation across all conditions (e.g., Fig. 2E), i.e., the condition-independent modulations.

The results of the effect size analysis confirmed findings from previous studies: (1) The direction of the visual cue and reward effect sizes were significantly different from zero (Fig. 2H–I, p < 10^-6 Wilcoxon signed-rank test) as expected from neurons in the FEF[15,17], (2) The effect size was widely distributed, with most of the neurons having small values close to zero but some exhibiting large effect sizes (e.g., Fig. 2G) demonstrating the typical long-tailed coding of task parameters[34], (3) The time effect size tended to be the largest (Fig. 2J, and see for example the comparison in Fig. 2G between reward and time effect sizes), consistent with previous studies reporting that condition-independent modulations can account for the largest share of the population activity[16,35].

Tagged neurons had the largest time and reward effect sizes during the cue epoch

We compared the effect sizes of the three populations of neurons. On all three tasks, the average time effect size of the tagged neurons was larger than in the untagged and the area-untagged populations (Fig. 2J). In the tasks where we manipulated the upcoming reward size (saccade and pursuit), the reward effect size was larger in the tagged neurons (Fig. 2I). On the tasks where the visual eccentric target provided information about the direction of the upcoming movement (saccade and choice), the direction effect size was not significantly different across populations (Fig. 2H). Thus, the effect size analysis of the cue period indicated that reward and condition-independent modulations were the most strongly observed in inputs to the basal ganglia.

Larger time and reward effect sizes in the tagged neurons persisted into the movement epoch

The FEF demonstrated coding of multiple experimental variables during the movement epoch as well. Figure 3A shows an example neuron that coded multiple experimental variables during the saccade tasks around the time of the saccade. This neuron responded with a burst of activity around the time of the saccade. Interestingly, its response was stronger in the small reward condition (red versus blue lines in Fig. 3B,D). Later we return to an analysis that quantified whether and when reward size was coded by an increase or decrease in activity. The same neuron also responded differently for different directions of movement (Fig. 3C), and the overall modulations were not averaged out across conditions (Fig. 3E). This coding of multiple components was reflected in the positive values of the effect size.

We then compared the effect sizes across populations. As was the case for the cue, the tagged neurons had the largest reward and time effect sizes (Fig. 3G and H). The direction effect size was not significantly different across populations. Thus, analysis of the movement epoch further demonstrated the coding of reward-related information by inputs to the basal ganglia and the larger condition-independent modulations of these inputs. We also calculated the effect sizes during the outcome (Fig. 3I-K) which further supported the larger time and reward effect size by the tagged neurons (Fig. 3K and J).

In addition to the effect size differences between the populations, we also found that the baseline firing rate defined by the pre-cue activity (Fig. 1A pre-cue) was larger for the tagged neurons. Even when taking this firing rate baseline into account, the time and reward effect sizes were larger for the projecting neurons (Fig. S4).

Beyond effect size: taking into account the sign of the modulation and dissociating reward from coding of eye movement

The effect size analysis has the advantage of partitioning activity according to the experimental variables, thereby providing a convenient tool for comparing populations. However, it has some drawbacks that require further analysis to better understand coding in the FEF. The effect size does not indicate whether the modulations are a result of an increase or decrease in rate. For example, the neurons shown in Figure 2A and 3A responded differently to the large and small reward conditions and therefore had a positive reward effect size; however, their responses were qualitatively different since one neuron responded more strongly to the large reward whereas the other responded more strongly to the small reward. Another shortcoming of effect size is that the nature of the interaction between variables can be difficult to interpret. For example, the interaction between reward and direction could result from either reward potentiation or attenuation of directional tuning, or other patterns. Finally, the effect size analysis did not dissociate the effect of reward itself from the effect of reward on eye movements. In particular, we found that saccade latency was shorter for the large reward (Fig. 1C). This raises the question of whether the reward-related modulations observed in the cue and movement epochs reflect differences in latency.

To further study if and how reward was processed in the different populations, we conducted the next analyses by taking the population average of the different populations and studying the distribution of responses across neurons. We aimed to test whether (1) the sign of the reward modulation was consistent across neurons, (2) how reward modulations interacted with the directional tuning curves, and (3) whether the reward modulations could be explained by the effect of reward on eye movements.

Reward potentiated directional tuning in the cue epoch the most strongly for tagged neurons

To study how reward modulated the coding of direction, we first tested whether each neuron was directionally tuned (ANOVA for direction, $p < 0.05$) and then calculated the preferred direction (PD, see Methods). We then calculated the average activity in the PD and the opposite direction (Null) for the large and small reward conditions. Directional tuning was enhanced by the reward, as indicated by the larger responses in the PD for the large versus small reward condition (Fig. 4A-C). Quantification of reward modulation in the PD at the single-neuron level showed that neurons across all the populations tended to respond more strongly in the large reward condition (Fig. 4E-G). The effect of the reward in the PD was the largest in the tagged neurons (Fig. 4D, solid black line): the difference in rate between the large and small rewards was significantly greater for the tagged neurons than for the other groups ($p=0.02$ and $p=0.002$ for tagged versus untagged and tagged versus area-untagged, Wilcoxon rank-sum test). Thus, the directional tuning of the visual responses of inputs to the basal ganglia was already strongly potentiated by reward. This potentiation of visual responses by reward is akin to reports of modulations of spatial attention[36,37], suggesting that inputs to the basal ganglia already carry spatial attention information.

Reward did not consistently modulate the sign of activity in neurons that were not directionally tuned. Although reward had a significant effect in many of these neurons (70/365 open dots in Fig. 4H-J), the neurons that increased their activity in the large reward or small reward conditions were overall evenly distributed. As a result, the overall population difference between the large and small rewards did not reach significance (Fig. 4D, dashed line). Similarly, during the cue epoch of the pursuit task when directional information was not available, many neurons coded the reward (174/512), but the sign of the modulation was inconsistent across neurons (Fig. 4D gray line).

The modulations described so far indicated that, in the projection to the basal ganglia, reward information was already integrated. We termed these modulations ‘reward-related’ since reward was the task variable we manipulated. However, it is possible that these modulations reflected effects downstream to the representation of reward expectation. In the extreme case, all reward-related activity could be linked to the motor command. In the current task and recording site in particular, the reward-related modulations could reflect the shorter latency (Fig. 1C). We therefore used trial-by-trial variability in behavior to test whether reward modulations could be explained by saccade latency. Within each direction and reward size condition, we divided the trials into those with short and long saccade latencies. This division had little impact on either the population averages or single-neuron activity (Fig. S5). We then directly compared the short latency trials in the small-reward condition with the long latency trials in the large-reward condition. Selecting this set of trials reversed the behavioral effect (Fig. S5A, purple dots). However, the neural effect persisted, as we found for the full trial set. Specifically, the tuned neurons responded more strongly to large compared to small rewards in their preferred direction (Fig. S5C, p=10^-7 Wilcoxon signed-rank test). In addition, the sign of the modulations in the untuned neurons that coded reward size remained the same (Fig. S5D, 69/70 neurons) indicating that although the behavior flipped in sign (large shorter than small, to small shorter than large) in this control, the neural activity did not.

Similar controls for the amplitude of the saccade, peak velocity showed that these parameters had no detectable impact on FEF activity in the cue (not shown). In addition, removing trials with microsaccades during the cue[38] had no impact on the reward modulations (not shown). Thus, as we[17] and others [39] have shown in the FEF, reward-related activity could not be explained solely by coding of the eye movement.

During movement, reward modulations averaged out at the population level but remained coded beyond eye movements

The analysis above indicated that in visually tuned neurons, reward potentiated directional tuning. We next tested whether this potentiation extended to movement. In the movement epoch we could control directly for putative differences in the reward effect on saccade latency (Fig. 1C) by aligning the neural activity on the time of the saccade. In all populations, the overall directional tuning was only slightly modulated by reward (Fig. 5A-C), as indicated by the similarity of the population response in the PD at the time of the saccade. Neuron-by-neuron analysis revealed that this lack of modulation was the result of the cancelling out of opposite modulations of reward rather than the elimination of reward modulation by single neurons. This was indicated by the distribution of neurons with significant coding of reward along both sides of the diagonal in Figure 5E-G. Overall, the fraction of neurons that coded reward significantly around the saccade (-100-300 ms) was much greater than chance (142 out of 512 neurons, p ~ 0, chi-square test).

A similar averaging out was observed during pursuit (Fig. S6) and in neurons that were not directionally tuned during saccade and pursuit (Fig. 5D, H). These results extend our previous findings on reward modulation during pursuit[17] where we found a general inconsistency in the coding of reward and direction during movement in the FEF. This coding enables commonly assumed readouts such as the population vector, with minimal interference from reward[17], thereby permitting simple population readouts that disentangle the multiplexed direction from reward information.

FEF inputs to the basal ganglia code the selected target before movement

Several theories have posited that the basal ganglia implement action selection[40]. The effect size analysis above indicated that in the choice task, the inputs to the basal ganglia had already differentiated between selection conditions. Specifically, activity was different if the large reward target appeared on the left or right side (dashed line in Fig. 2H). We next performed an analysis that went beyond effect size to better understand how the FEF codes choice-related activity.

For each neuron we defined the preferred choice condition as the choice condition with the larger response, and the condition with the smaller response as the null choice condition. In the choice task we contrasted two directions. As a result, the preferred choice condition could differ from the PD in the tuning task which was based on four directions (see Methods). We averaged the activity across neurons in the preferred and null choice conditions (Fig. 6A-C). We controlled for statistical biases by using half of the data from each neuron to determine the preferred choice condition and the other half to calculate the average activity. We found that in all populations, soon after cue presentation, the FEF activity differentiated between choice conditions (Fig. 6A-C, solid versus dashed black line).

To further study neural representations during the choice task, we ran analyses comparing choice vs. single-target trials[16,41]. The monkeys almost always selected the larger reward target (Fig. 1E). In addition, saccade latency was very similar (although not identical) in the large-reward and choice trials (Table S1). Given these findings, we next examined how the choice process was reflected in FEF activity and in the information conveyed to the basal ganglia. We tested whether FEF activity and its inputs to the basal ganglia would resemble the responses to the selected target alone, while ignoring the other target. These representations were expected if the FEF is situated near the completion of the choice process and contrast with other representations, which are expected at early or intermediate stages of the choice process (see below).

In all populations, the activity in the choice condition was almost identical to the activity in the corresponding large reward condition; i.e., activity in the choice trials was similar to trials where the monkeys only observed the large reward target. The similarity between choice and large reward is manifested in the similarity of the black and blue lines in the population average in Figure 6A-C. In contrast, activity during choice trials differed from that observed when only the small-reward target was shown. As depicted in Figure 6A–C, responses in the preferred choice condition were stronger than in the single small-reward trials, whether the small-reward target appeared in the same location as the small target in the choice trials (solid black vs. dashed red lines) or when it appeared in location of the large reward target in the choice trials (solid black vs. solid red lines). Quantification using PSTH distances showed that the distance between the choice and the large reward target appearing alone was the smallest (Fig. 6D, E, for preferred and null choice conditions).

This representation in the selection task aligns with expectations associated with the coding of the selected target and its reward value[41], but not with other coding schemes typically associated with intermediate processing stages. The difference between choice conditions suggests that the neurons did not code the only the large reward value, regardless of its location[42]. The results also showed that neurons did not code the reward value in a specific location. If this were the case, we might expect activity in the preferred choice condition to resemble the response to the large reward appearing alone, since the larger reward target lies within the neuron's receptive field. However, this interpretation falls short, because in the null choice condition, where the small reward target is within the receptive field, we would expect activity to resemble the response to the small reward target appearing alone, this was not observed (Fig. 6A-C, dashed black line versus solid red and Fig. 6E). Instead, the activity in the null choice condition more closely matched that of the large reward condition (Fig. 6A-C, dashed black line versus dashed blue line). Finally, the coding of reward in the single-target trials by these same neurons (Fig. 6A–C, blue versus red traces) indicates that they do not merely encode which target was selected, but also incorporate value signals into their activity. Thus, overall, the activity in the selection trials is best interpreted as a post-selection evaluation of the selected target (see Discussion).

In the effect size analysis, we found slightly stronger coding of choice conditions in the tagged neurons, but this effect was not significant (dashed line in Fig. 2H). Therefore, in the analyses above, where we compared the choice and single target conditions, we did not include statistical comparisons between populations. However, as expected, the trends were consistent across analyses, with tagged neurons showing larger modulation.

Discussion

We used optogenetics to identify corticostriatal neurons. We then examined the functional properties of these neurons in three eye movement tasks. We found that inputs to the basal ganglia already contained considerable information about the reward (Fig. 2, 3 and 4) and choice (Fig. 6). We also found an unexpectedly strong condition-independent modulation in the inputs to the basal ganglia (Fig. 2 and 3). Below we discuss the functional implications of these results.

Optogenetics in monkeys

Optogenetics is widely used in rodent research, but many studies on non-human primates continue to rely exclusively on electrophysiological methods. One major hurdle is that optogenetics in monkeys has been ineffective in driving movement. Our findings, along with other recent studies[27,43], indicate that this limitation is not due to a lack of viral expression. We showed that optogenetics can be successfully applied to study neural pathways functionally, thus demonstrating its utility beyond causal manipulations. Our success in the expression of AAVretro in targeting cortico-striatal neurons in monkeys[44] provides the technical foundation for studying the functional properties of the cortico-striatal neurons. The high yield of neurons we identified suggests that this approach can be applied to studying other pathways as well[45]. Thus, our findings make the case for expanding the use of optogenetics in non-human primates for functional circuit dissection.

Interpreting the difference in effect size between groups

Ideally, we would want to split neurons into those that project to the basal ganglia and those that do not. However, untagged neurons cannot be classified as non-projecting. The functional difference between the tagged and untagged neurons suggests that some of the untagged neurons did not project to the basal ganglia. Therefore, the observed functional differences between tagged and untagged neurons are likely to be an underestimation of the actual differences between neurons projecting and not projecting to the basal ganglia.

The largest effect size for the tagged neurons indicated that the signal-to-noise of these neurons in coding of the task variables was the largest. This increase in signal-to-noise could arise from the convergence of inputs from neurons with a similar signal, which would result in a sharpening of the response[31]. Thus, the difference in effect size between the tagged and nearby untagged neurons (Fig. 2, 3) suggests that in the transformation within the cortex, inputs with similar reward information or condition-independent modulations tend to converge on neurons that convey the cortical output to the basal ganglia.

The difference in effect size between the tagged and the area-untagged neurons could be due to several circuit motifs. One possibility is that the area-untagged neurons were located in FEF regions that are not directly connected to the basal ganglia and that these regions receive less reward and condition-independent information. A second possibility is that these neurons are part of the same microcircuitry as the other neurons but are in different layers. If this is the case, our results may point to an intra-layer computation[46,47] in which reward information and condition-independent modulations are refined between cortical layers. Consistent with this possibility, we found that tagged neurons were often confined to a thin layer surrounded by untagged sites (Fig. S2). However, it is experimentally difficult to identify layers with standard electrophysiological methods because the FEF lies within the arcuate sulcus.

Reward and choice in the FEF

Our results confirm previous studies that have found reward and choice signals in the FEF[15–17,48,49]. In most task epochs, we found that many neurons differentiated between the small and large reward conditions. One of the major issues in studies on reward-related activity is how to link it to specific cognitive or behavioral processes such as attention, reward expectation, or motor planning. Fully dissociating reward from all other cognitive and motor variables is challenging and perhaps not even conceptually possible[50]. A common finding in motor systems is that behavior under task control does not fully explain reward modulations. Consistent with these findings, our analyses (Fig. S5) showed that the reward-related signal could not be explained as a modulation of eye movements.

Further insights into how reward modulates FEF activity did arise from the patterns of interaction with the coding properties of the neurons. In the movement epoch, the reward did not consistently modulate directional tuning (Fig. 4 and 5). This coding ensures that commonly suggested readouts of movement, such as the population vector, can reliably read out the movement but are weakly modulated by reward. Similarly, population readouts that aggregate the single-neuron reward modulations[51] average out movement-related activity. By contrast to the movement epoch, in the cue epoch of the saccade task, the reward-related activity potentiated directional tuning (Fig. 4). Thus, in the cue epoch, readouts such as the population vector would contain reward-related information in addition to the stimulus direction.

The lack of activity related to the small reward target in the choice task is expected at the outcome of the selection process. The task we used was quite simple in terms of sensory discrimination and reward valuation. It is therefore possible that intermediate representations, such as the accumulation of evidence[52] or deliberation between reward values[53], could emerge in other tasks, thus making it interesting in future works to study FEF inputs to the basal ganglia in such contexts.

As with the single-target condition, it is difficult to pinpoint the exact content of the choice-related modulations. The loose relationship between activity and the specifics of the upcoming eye movement (Fig. 5S), along with the dissociation we previously observed between choice and movement details[16], suggests that during the cue epoch, choice-related activity does not directly encode the eye movement command. The coding of reward in the single task indicates that the FEF is not just coding which target was selection. Instead, activity is consistent with coding of the selected target and the reward value associated with it; i.e., the post-decision evaluation of the upcoming reward.

Sources for reward information to the basal ganglia

A critical question that arises from our findings is how cortical and dopaminergic sources of reward interact within the basal ganglia. One aspect of this question is how the instantaneous activity in the basal ganglia is affected by the convergence of two sources of reward. The classical model of the basal ganglia assumes that reward information in D1 and D2 receptor - expressing neurons is oppositely modulated by the dopamine signal[61,62]. Our results point towards another source of reward information, which might be even more important for determining how reward modulates striatal activity.

In addition to instantaneous activity, reward information in the basal ganglia is thought to drive learning. An important question is how the two sources of reward interact during the learning of reward value. Dopamine has been shown to initially respond to the reward outcome. During learning, the response shifts to the cues that predict reward[4]. Here, we did not aim to characterize how cortical inputs code reward-related activity as it emerges during learning. Unlike dopaminergic neurons, reward was encoded throughout the trial, including the cue epoch, movement, and outcome phases (Fig. 2I and 3G,J). Further work that uses either learning paradigms or probabilistic rewards could further dissect the functional properties of the cortical input. Comparing the dynamics of neural changes during learning in dopamine and cortical inputs to the striatum would provide important constraints on how reward is distributed along the different inputs of the basal ganglia.

Methods

Animals and ethics

We collected neural and behavioral data from two female (monkey K and monkey H) *Macaca fascicularis* monkeys (4-5 kg). All procedures were approved in advance by the Institutional Animal Care and Use Committees of the Hebrew University of Jerusalem.

Surgical procedures

We first implanted head holders to restrain the monkeys' heads in the experiments. We performed a second surgery to place a round recording cylinder with an inner diameter of 19 mm over the FEF. The center of the cylinder was placed above the skull at 19 mm anterior and 15 mm lateral to the stereotaxic zero, and tilted 25°with respect to the medial plane to allow for injections to the caudate nucleus. The localization of the FEF and caudate was based on stereotactic coordinates and then confirmed with post-implant MRI (Figs. S1 and S2) scan and electrophysiological markers (see below).

Experimental setup

After the monkeys recovered from head holder surgery, they were trained to sit calmly in a primate chair (Crist Instruments) and consume liquid food rewards (baby food mixed with water and infant formula) from a tube set in front of them. We trained the monkeys to track spots of light that moved across a video monitor placed in front of them. Visual stimuli were displayed on a CRT monitor with a refresh rate of 85 Hz (55 cm from the eyes of the monkeys). The stimuli appeared on a dark background in a dimly lit room. A computer performed all real-time operations and controlled the sequences of target motions. The position of the eye was measured with a high temporal resolution camera (1 kHz, Eyelink 1000 plus, SR research) and collected for further analysis.

Experimental design

The monkeys were engaged in saccade, pursuit, and choice tasks. In the saccade task each trial started with a bright white circular 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 at a position 10° eccentric to the fixation target in one of the four cardinal directions. We defined this time point as the cue onset. The color of the target indicated that upon successful completion of the trial, the monkey would receive either a small or large reward (0.07 or 0.15 ml of food). For monkey K we used blue and red targets for the large and small rewards, and reversed the associations for monkey H. The monkeys were required to maintain fixation on the central target. Moving the eye towards the eccentric target would abort the trial. After a variable delay of 500-700 ms, the fixation target disappeared, which served as the Go signal for the monkeys to move their eyes to the peripheral target. The monkeys had to acquire fixation in 750 ms and then after an additional 200-300 ms they received the reward.

The structure of the pursuit trials was similar to the saccade trials but with the following differences. A colored target replaced the white target in the center of the screen. After the cue epoch, the target jumped 4° in one of the four cardinal directions and then moved in the opposite direction towards and through the center of the screen at 20°/s (step-ramp[63]). The target moved for 750 ms, then stopped and stayed still for an additional 200-300 ms.

The structure of the choice trials was also similar to the saccade task but instead of a single target, two colored eccentric targets appeared at the cue. The target appeared either along the horizontal or vertical axis 10° eccentric to the fixation target. When the central target disappeared, the monkey was free to select the target to which to saccade. Online we detected the saccade as an eye movement that exceeded 80°/s. The target that was closer to the eye at the end of the saccade remained and the other target disappeared. The monkeys were then required to continue to fixate on this target and received a reward based on the color of the target. In all tasks we enforced a precision window of 3-5°x3-5° around the fixation target with grace periods after the Go signal to allow the monkeys to acquire the target. Trials in which the monkeys failed to meet the fixation requirements were removed from the analysis. In each recording session, we interleaved saccade, pursuit, and choice trials such that they contained 80 pursuit and 80 saccade trials (4 directions X 2 reward conditions X10 repetitions), and 80 selection trials (2 directions X 2 target configuration conditions X 20 repetitions).

Virus injection

We injected AAVretro-hSyn-ChR2-GFP virus (ELSC virus core) into the striatum. We first injected the virus into the striatum of mice and confirmed the expression with histology and recorded LFP and single units in the cortex in response to light stimulation. We then injected the same virus to the monkeys. We made efforts to confirm the injection site without histology, as we aimed to release the monkeys to a sanctuary at the end of the experiments. Before the injections we mapped the recording chamber, and then lowered single electrodes (Thomas recordings Mini matrix) to identify the striatum. Identification was based on well-known physiological and anatomical characteristics, including the depth of the electrode compared to the MRI scan and atlas, the shape of the extracellular spike associated with medium spiny neurons (MSN), the characteristic low firing rate of the MSN, the presence of tonically active neurons (TANs) with their broad extracellular spike, and a gap in the recorded neurons after penetration of the cerebral cortex.

After a few weeks of mapping, we injected the virus. We loaded low-viscosity silicone fluid (polydimethylsiloxane, viscosity: 1 cSt at 25°C, Merck) into a 1 ml syringe and a thick-walled plastic tubing (inner diameter: 0.25 mm, length: 50 cm, Thomas Recordings injection system). The fluid fluoresced in color to allow us to identify the boundary between the silicon and the virus to confirm the injections. After loading the virus, we lowered a recording electrode and an aligned injection micropipette (Thomas Recordings, 35 gauge, lowered with Mini matrix). Based on the extracellular activity we identified the upper edge of the striatum. We then searched for the lower edge of the striatum and then injected 1-2 mm above the lower edge to 1-2 mm before the upper edge of the striatum 1 mm apart. In each site we injected 2-3 µl of the virus (titer $2.3*10^{12}$) at a speed of 15nl/s. We then waited 5 min before we pulled the electrode 1 mm upwards to the next site. In the upper site of the striatum, we waited 30 min before extracting the electrode from the brain. In each penetration we injected 1-4 sites depending on the thickness of the striatum. We aimed the injections towards the eye movement areas in the caudate (Fig. S1). Overall, we injected in 10 and 9 penetrations in monkeys K and H.

To further verify the injection location, for several injections we added Manganese (Mn^{+2}, 0.01M) to the injected solution. Then, a few hours after the injection, we scanned the monkeys in an MRI using contrasts that make Mn^{+2} visible (Fig. S1, sequence details: T1, flash, TR=500, TE=3.7, FA=30). In the first monkey we performed the scan in which we injected only Mn^{+2} to validate our injection system before injecting the virus. After validation in the second monkey, we added Mn^{+2} to the same solution we used to inject the virus.

Neural recordings and spike sorting

On each recording day, we lowered an optrode (an optical fiber connected to a tetrode, OD=356 µm, Thomas Recordings) to the FEF. When we reached the site with the neurons, we applied light stimulation (LED light, wavelength: 470nm) protocol that included 10 repetitions of 10 pulses (5 ms durations) of stimulation at 10, 20 and 50 Hz (overall 10X10X3 = 300 pulses). Online we monitored the activity and searched for responses of single units or local field potential (LFP) to the light stimulation. When we observed a response, we applied the behavioral protocol. Often, we found a relatively thin layer (typically less than 2 mm, Fig. S2) in which neurons or the LFP responded to the stimulation, in contrast to widespread activity that would be expected if the virus had leaked to the recording site. For a more complete comparison we also recorded neural activity in areas where we did not observe responses to the stimulation. We sorted the spikes offline (Plexon offline sorter) and only well-isolated units were included in the analysis.

To confirm that the neurons that responded to light stimulation were the same neurons that were recorded during the task, we calculated the correlation and distance between the spike waveform of all the cells that were recorded during the task session and the light-evoked response waveform (Fig. S3). 86% (93/108) of the neurons identified as connected met previously used criteria for ensuring that light-evoked waveforms were similar to the task waveform (correlation coefficient > 0.9)[28].

Quantitative and statistical analyses

To detect responses to the optical stimulation we calculated the spike rate in the 20 ms after the stimulation in 1 ms bins (in units of spikes/ms). We then subtracted the rate before the stimulation from the rate after the stimulation to calculate the rate of increase of the spikes in each ms after the stimulation. We excluded the 50 Hz stimulation condition in this analysis since the inter-stimulus-interval was too short to establish the pre-stimulus activity. Neurons were characterized as responding to the stimulation if the cumulative distribution of the increase in rate exceeded 0.5. This criterion corresponded to an additional spike above baseline within 20 ms of stimulus onset in 50% of the trials, and was much more stringent than a statistical test for an increase in rate. We used this stringent criterion since even a small number of errors in the spike sorting in which a non-responding neuron was contaminated with a responding neuron could lead to a significant increase. Finally, we confirmed that only neurons were included in which the response had a latency of less than 5 ms as determined by an increase in rate of at least 0.2 in the first 5 ms after the stimulation. This additional screening only removed two neurons and was practically negligible. The criterion was conservative in comparison to the classification of a human observer. All the neurons we classified as responding were also classified as responding by a human observer (N=108), but the human observer classified 57 other neurons as responding. Repeating the analysis based on a human observer classification did not alter our conclusions. Overall, 73, 163 and 50 neurons were classified as tagged, untagged and area-untagged in monkey H, and 35, 103 and 88 were classified as tagged, untagged and area-untagged in monkey K.

To calculate the coding of a variable throughout an epoch, we calculated the partial ω^2 in ANOVA models. The ANOVA model could include three variables: 1) Direction, corresponding to the direction of movement or visual cue, 2) Reward size, which was either large or small, and 3) Time, corresponding to the time bin relative to the event, calculated in bins of 250 ms. We included a single time bin before the event and up to four time bins after the event. In the saccade task, we used a 3-way ANOVA in all task epochs (cue, movement, and reward outcome). In the pursuit task, we used only reward and time in the cue epoch, since the direction was not cued in this task, and the full 3-way ANOVA in the movement and reward outcome epochs. In the choice task, we used a 2-way ANOVA model with time and direction, where direction was defined as whether the large reward appeared on one side (e.g., right) or the other (e.g. left).

Note that the variables we used in the effect size analysis were defined by the task design and not by behavior, but since the monkeys were well- trained, there was mostly full correspondence between the two. Specifically, in the saccade and pursuit tasks, the monkeys always moved to the target, so that the target and movement direction were the same experimental condition. A possible deviation could arise during the choice task, but since the monkeys almost always chose the larger reward (Fig. 1E), splitting the task trials by choice direction or by the location of the large reward target was in fact the same. For this reason, we could not calculate the reward effect size in the choice trials since the monkeys very rarely selected the small reward.

To quantify the contribution of each task variable we used ω^2, a common effect size measure employed in ANOVA designs that is often preferred over other effect size measures since it is unbiased. We calculated the number of spikes in 250 ms bins, and then calculated the ω^2 as follows:

$$\omega^2 = \frac{SS_{effect} + SS_{effect\ x\ time} - \frac{df_{effect} + df_{effect\ x\ time}}{df_{error}} \cdot SS_{error}}{SS_{effect} + SS_{effect\ x\ time} + \frac{\left(N - df_{effect} - df_{effect\ x\ time}\right)}{df_{error}} \cdot SS_{error}} \text{ (equation 1)}$$

where SS_{effect} is the ANOVA sum of squares for the effect of a specific variable, $SS_{effect\ x\ time}$ is the ANOVA type II sum of squares for the interaction of a specific variable with time, SS_{error} is the sum of squares of the errors after accounting for all experimental variables, df_{effect}, $df_{effect\ x\ time}$ and df_{error} are the degrees of freedom for a variable, the interaction of the variable with time and the error, respectively. We included the interaction term since it quantifies the time-varying coding of the variable. N is equal to the number of trials x the number of time bins. To calculate the time effect size, we did not include interaction term or the corresponding degree of freedom.

For the cue and outcome epochs we used a time window from 250 before to 500 ms after the event onset. For movement, we used a time window of 0-1000 ms after the Go signal. Changing the bin size or the time window yielded highly correlated effect sizes. We also quantified all possible interactions (Reward x Direction, Reward x Direction x Time) but these did not yield any consistent results and were often small in comparison to other values. We used the partial effect size since it enables a better comparison between neurons that responded or did not respond to other experimental variables. Using the full effect size did not alter any of our conclusions. To test for differences between the average effect size across populations we used a permutation test. We shuffled the group labels, and calculated the difference between the average effect size of the shuffled populations. We repeated this 1000 times and if the actual average difference was larger than 95% of the shuffles, the difference was considered significant.

To determine the preferred direction (PD) of each neuron, we calculated the direction of the center of mass of the responses and assigned the PD to the cardinal direction that was closest to the calculated direction. We used both large and small reward trials to calculate a single PD for each cell in each task and task epoch. We used a window of 50-500 ms after the cue. When aligning the saccade, we used a window from 100 ms before to 300 ms after the saccade. We used a shorter time before the saccade to minimize the effect of the cue on the tuning of the saccade. For pursuit, we used a window of 50-500 ms after motion onset. We confirmed that the results were insensitive to the time windows. We used these same time windows to determine whether a neuron responded significantly to the reward or directions. For each trial, we counted the number of spikes in the window and performed a 1-way (Direction or Reward) or a 2-way (Direction and Reward) ANOVA to test for significance of the response to a variable.

We calculated the peristimulus time histogram in bins of 1ms and smoothed with a 30 ms SD Gaussian. We calculated the population PSTH by averaging the responses from all neurons in the PD or Null direction. Since we were interested in the effect of the reward on tuning and we used both rewards for calculating the PD, the effect of reward on tuning was not biased by this analysis. To avoid bias in the analysis of choice conditions, we used half of the choice trials to test whether a neuron's response was significantly different across the choice conditions and defined the preferred choice direction based on which direction had the larger firing rate. We then used the other half of the trials to plot the population PSTHs and perform comparisons with the single target trials. This procedure insured that the population responses were not statistically biased by the selection of the preferred direction.

All statistical tests were two-sided. The distributions of effect sizes tended to be non-normal and long-tailed. We therefore used permutation methods to compare the mean of two populations. The population labels of the effect sizes were permuted, and the difference between group means was calculated for each permutation. The p-value was defined as the probability of observing a statistic more extreme than that obtained for the real sample. The statistics were recalculated 10,000 times, as we observed that the p-values were stable at this number of repetitions.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT to edit the English. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Author contributions

AHL. and MJ conceived and planned the experiments. AHL and YH collected the data. TK performed preliminary experiments in mice. AHL and MJ, analyzed the data. AHL, EL and MJ interpreted the data. AHL. and MJ wrote the manuscript. All authors read, commented on, and approved the final manuscript.

Data availability

The data supporting the findings of this study are available at Zenodo (DOI: 10.5281/zenodo.18467078). The repository contains the curated and processed datasets required to reproduce analyses and figures reported in the paper. The raw data underlying these datasets can be provided by the corresponding author upon request.

Code availability

The code used for data analysis is available from the corresponding authors upon request.

Competing interests

The authors declare no competing interests.

Acknowledgments

This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 755745).

1. Sutton, R. & Barto, A. Reinforcement learning: An introduction. (2018).
2. Takahashi, Y., Schoenbaum, G. & Niv, Y. Silencing the critics: understanding the effects of cocaine sensitization on dorsolateral and ventral striatum in the context of an actor/critic model. *Front Neurosci* 2, 86–99 (2008).
3. Doya, K. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol* 10, 732–739 (2000).
4. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science (1979)* 275, 1593–1599 (1997).
5. Joshua, M., Adler, A., Mitelman, R., Vaadia, E. & Bergman, H. Midbrain Dopaminergic Neurons and Striatal Cholinergic Interneurons Encode the Difference between Reward and Aversive Events at Different Epochs of Probabilistic Classical Conditioning Trials. *Journal of Neuroscience* 28, 11673–11684 (2008).
6. Matsumoto, M. & Hikosaka, O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459, 837–841 (2009).
7. Lammel, S., Ion, D. I., Roeper, J. & Malenka, R. C. Projection-Specific Modulation of Dopamine Neuron Synapses by Aversive and Rewarding Stimuli. *Neuron* 70, 855–862 (2011).
8. Kim, H. R. *et al.* A Unified Framework for Dopamine Signals across Timescales. *Cell* 183, 1600-1616.e25 (2020).
9. Mohebi, A., Wei, W., Pelattini, L., Kim, K. & Berke, J. D. Dopamine transients follow a striatal gradient of reward time horizons. *Nature Neuroscience 2024 27:4* 27, 737–746 (2024).
10. Engelhard, B. *et al.* Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature 2019 570:7762* 570, 509–513 (2019).
11. Salinas, A. G. *et al.* Distinct sub-second dopamine signaling in dorsolateral striatum measured by a genetically-encoded fluorescent sensor. *Nature Communications 2023 14:1* 14, 1–16 (2023).
12. Turner, R. S. & DeLong, M. R. Corticostriatal Activity in Primary Motor Cortex of the Macaque. *Journal of Neuroscience* 20, 7096–7108 (2000).
13. Musall, S. *et al.* Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making. *Nature Neuroscience 2023 26:3* 26, 495–505 (2023).
14. Roesch, M. R. & Olson, C. R. Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *J Neurophysiol* 90, 1766–1789 (2003).
15. Ding, L. & Hikosaka, O. Comparison of Reward Modulation in the Frontal Eye Field and Caudate of the Macaque. *Journal of Neuroscience* 26, 6695–6703 (2006).
16. Raghavan, R. T. & Joshua, M. Dissecting patterns of preparatory activity in the frontal eye fields during pursuit target selection. *J Neurophysiol* 118, 2216–2231 (2017).
17. Lixenberg, A. & Joshua, M. Encoding of Reward and Decoding Movement from the Frontal Eye Field during Smooth Pursuit Eye Movements. *Journal of Neuroscience* 38, 10515–10524 (2018).
18. Bruce, C. J. & Goldberg, M. E. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53, 603–35 (1985).
19. Schall, J. D. & Hanes, D. P. Neural basis of saccade target selection in frontal eye field during visual search. *Nature* 366, 467–469 (1993).

20. Parthasarathy, H. B., Schall, J. D. & Graybiel, A. M. Distributed but convergent ordering of corticostriatal projections: Analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *Journal of Neuroscience* 12, 4468–4488 (1992).
21. Cui, D. M., Yan, Y. J. & Lynch, J. C. Pursuit subregion of the frontal eye field projects to the caudate nucleus in monkeys. *J Neurophysiol* 89, 2678–2684 (2003).
22. Kate, M. *et al.* Eye movements in monkeys with local dopamine depletion in the caudate nucleus. I. Deficits in spontaneous saccades. *Journal of Neuroscience* 15, 912–927 (1995).
23. Bruce, C. J., Goldberg, M. E., Bushnell, M. C. & Stanton, G. B. Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol* 54, 714–734 (1985).
24. Schall, J. D. The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357, 1073–1082 (2002).
25. Lemon, R. Methods for Neuronal Recording in Conscious Animals. *IBRO Handbook Series* 4, 56–60 (1984).
26. Tervo, D. G. R. *et al.* A Designer AAV Variant Permits Efficient Retrograde Access to Projection Neurons. *Neuron* 92, 372–382 (2016).
27. Cushnie, A. K. *et al.* Using rAAV2-retro in rhesus macaques: Promise and caveats for circuit manipulation. *J Neurosci Methods* 345, 108859 (2020).
28. Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B. & Uchida, N. Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature 2012 482:7383* 482, 85–88 (2012).
29. Lima, S. Q., Hromádka, T., Znamenskiy, P. & Zador, A. M. PINP: A New Method of Tagging Neuronal Populations for Identification during In Vivo Electrophysiological Recording. *PLoS One* 4, e6099 (2009).
30. Stauffer, W. R. *et al.* Dopamine Neuron-Specific Optogenetic Stimulation in Rhesus Macaques. *Cell* 166, 1564-1571 e6 (2016).
31. Larry, N., Zur, G. & Joshua, M. Organization of reward and movement signals in the basal ganglia and cerebellum. *Nature Communications 2024 15:1* 15, 1–13 (2024).
32. Olejnik, S. & Algina, J. Measures of effect size for comparative studies: Applications, interpretations, and limitations. *Contemp Educ Psychol* 25, 241–286 (2000).
33. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science (1979)* 348, 1352–1355 (2015).
34. Buzsáki, G. & Mizuseki, K. The log-dynamic brain: how skewed distributions affect network operations. *Nat Rev Neurosci* 15, 264–278 (2014).
35. Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* 3, (2016).
36. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370–373 (2003).
37. Thompson, K. G., Biscoe, K. L. & Sato, T. R. Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field. *Journal of Neuroscience* 25, 9479–9487 (2005).
38. Joshua, M., Tokiyama, S. & Lisberger, S. G. Interactions between target location and reward size modulate the rate of microsaccades in monkeys. *J Neurophysiol* jn.00401.2015 (2015) doi:10.1152/jn.00401.2015.

39. Glaser, J. I. *et al.* Role of expected reward in frontal eye field during natural scene search. *J Neurophysiol* 116, 645–657 (2016).
40. Gurney, K., Prescott, T. J. & Redgrave, P. A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol Cybern* 84, 401–410 (2001).
41. Tanaka, M. & Lisberger, S. G. Role of arcuate frontal cortex of monkeys in smooth pursuit eye movements. II. Relation to vector averaging pursuit. *J Neurophysiol* 87, 2700–2714 (2002).
42. Padoa-Schioppa, C. & Assad, J. A. Neurons in the orbitofrontal cortex encode economic value. *Nature 2006 441:7090* 441, 223–226 (2006).
43. Tremblay, S. *et al.* An Open Resource for Non-human Primate Optogenetics. *Neuron* 108, 1075-1090.e6 (2020).
44. Cushnie, A. K. *et al.* Using rAAV2-retro in rhesus macaques: Promise and caveats for circuit manipulation. *J Neurosci Methods* 345, 108859 (2020).
45. Yu, X. *et al.* Retrograde optogenetics reveals sensorimotor convergence within a corticotectal pathway of non-human primates. *Current Biology* 0, (2025).
46. Sajad, A., Godlove, D. C. & Schall, J. D. Cortical microcircuitry of performance monitoring. *Nature Neuroscience 2019 22:2* 22, 265–274 (2019).
47. Levy, S. *et al.* Cell-Type-Specific Outcome Representation in the Primary Motor Cortex. *Neuron* 107, 954-971.e9 (2020).
48. Hauser, C. K., Zhu, D., Stanford, T. R. & Salinas, E. Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. *Elife* 7, (2018).
49. Schall, J. D., Hanes, D. P., Thompson, K. G. & King, D. J. Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. *J Neurosci* 15, 6905–18 (1995).
50. Maunsell, J. H. R. Neuronal representations of cognitive state: Reward or attention? *Trends Cogn Sci* 8, 261–265 (2004).
51. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84 (2013).
52. Ding, L. & Gold, J. I. Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral Cortex* 22, 1052–1067 (2012).
53. Balewski, Z. Z., Knudsen, E. B. & Wallis, J. D. Fast and slow contributions to decision-making in corticostriatal circuits. *Neuron* 110, 2170-2182.e4 (2022).
54. Matsumoto, M. & Hikosaka, O. Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* 447, 1111–1115 (2007).
55. Mohebi, A. *et al.* Dissociable dopamine dynamics for learning and motivation. *Nature* 570, 65–70 (2019).
56. Niv, Y., Daw, N. D., Joel, D. & Dayan, P. Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology (Berl)* 191, 507–520 (2007).
57. Lauwereyns, J., Watanabe, K., Coe, B. & Hikosaka, O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 418, 413–417 (2002).
58. Semedo, J. D. *et al.* Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communications 2022 13:1* 13, 1–14 (2022).

59. Reinhold, K. *et al.* Striatum supports fast learning but not memory recall. *Nature 2025 643:8071* 643, 458–467 (2025).
60. Znamenskiy, P. & Zador, A. M. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature 2013 497:7450* 497, 482–485 (2013).
61. DeLong, M. R. Primate models of movement disorders of basal ganglia origin. *Trends Neurosci* 13, 281–285 (1990).
62. Surmeier, D. J., Ding, J., Day, M., Wang, Z. & Shen, W. D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons. *Trends Neurosci* 30, 228–235 (2007).
63. Rashbass, C. The relationship between saccadic and smooth tracking eye movements. *J Physiol* 159, 326–338 (1961).

Figure captions

Figure 1: Task behavior and opto-tagging of FEF neurons. A. Schematics of the three tasks. The series of black squares represent the consecutive visual stimuli presented on the screen. Dots show stationary targets and arrows indicate moving targets. The size of the drop corresponds to the reward size. **B.** Example of traces of the eye position in the horizontal direction; positive and negative values correspond to movement to the right and left. Each thin line shows a single trial; traces correspond to the task schematics above them. **C-E:** Day-by-day analysis of the behavior of the monkeys on the saccade (**C**) and pursuit tasks (**D**). Each dot shows the saccade latency (**C**) or average eye velocity 200 -250 ms after motion onset (**D**) averaged across trials and direction of movement from a single day. Horizontal and vertical axes show the large and small reward conditions. **E.** Fraction of trials in which the large reward was selected as a function of the recording session. Only days in which we recorded neural data are included. **F.** Schematics showing the injection in the striatum and recording in the FEF. **G**. Example of an extracellular recording during light stimulation. Blue dots and the blue line show the time of the stimulation (5 ms duration). **H.** Raster of the response of a single neuron to 100 light stimulations. Each dot shows the time of a single spike aligned to the time of the stimulation. Only trials from the 10 Hz stimulation are shown, to be able to depict the pre-stimulus activity without confounding with the responses to the previous stimulation.

Figure 2: Effect size analysis during the cue epoch. A-E. Example of responses of one neuron during the cue epoch on the saccade task. Blue and red traces correspond to the large and small reward conditions. The luminance (light to dark) corresponds to the direction of movement. **A** and **B** show the raster and PSTH for all conditions separately**. C-E** show traces in which the PSTHs were averaged across the reward (**C**), direction (**D**) and all (**E**) conditions. The ω^2 in the figure shows the effect size associated with each of the variables. **F**. Schematics demonstrating the calculation of the effect size. The outer square represents the total trial-by-trial variability, the intermediate square shows the variability explained by the experimental variables and the inner square shows the variability explained by one specific variable. The lower schematics demonstrate the calculation of the effect size. **G.** Example of the distribution of effect size. Each dot shows the effect size for the reward (horizontal) and time (vertical) variables for a single neuron. Different colors correspond to the different populations. **H-J**. Summary of the effect sizes during the cue for the three populations. Different plots show the effect size for the direction (H), reward (I) and time (J) variables. Vertical values show the averages, and error bars show the SEM. Lines above show the results of the permutation test for average differences across all tasks (ns – not significant, * $p < 0.05$, ** $p<0.01$, *** $p< 0.001$). P-values in **H–J,** corresponding to the order of the significance markers from left to right are 0.84, 0.29, 0.37, 0.005, 0.0002, 0.044, 0.0035, 0.0002, 0.08. In **H–J**, the saccade and pursuit groups included 108, 266, and 138 neurons in the tagged, untagged, and area-untagged groups. In the choice traces, the number of samples was doubled because we separated the horizontal and vertical choice tasks.

Figure 3: Effect size analysis during the movement epoch. A. Example of responses of a neuron during the movement epoch on the saccade task. Blue and red traces correspond to the large and small reward conditions. The luminance (light to dark) corresponds to the direction of movement. **A** and **B** show the raster and PSTH for all conditions separately. **C-E** show traces in which the PSTHs were averaged across the reward (**C**), direction (**D**) and all (**E**) conditions. The ω^2 in the figure shows the effect size associated with each of the variables. **F-K**. Summary of the effect sizes during movement (**F-H**) and outcome (**I-K**) for the three populations and tasks. Different plots show the effect size for the direction (**F,I**), reward (**G,J**) and time variables (**H,K**). Vertical values

show the averages, and error bars show the SEM. Lines above show results of the permutation test for average differences across all tasks (ns – not significant, * $p < 0.05$, ** $p<0.01$, *** $p< 0.001$). P-values in **F–H,** corresponding to the order of the significance markers from left to right are 0.18, 0.93, 0.17, 0.018, 0.0002, 0.08, 0.018, 0.0002 and 0.068. P-values in **I–K,** corresponding to the order of the significance markers from left to right are 0.042, 0.97, 0.027, 0.03, 0, 0.03, 0.002, 0.006, 0.97. In **F–K**, the saccade and pursuit groups included 108, 266, and 138 neurons in the tagged, untagged, and area-untagged groups. In the choice traces, the number of samples was doubled because we separated the horizontal and vertical choice tasks.

Figure 4: Reward modulation of directional tuning in the cue epoch. A-C. Average PSTH of the directionally tuned neurons for the three populations in the cue epoch on the saccade task. Light and dark colors show activity in the PD and the direction opposite to the PD (null). Blue and red show the activity in the large and small reward conditions. Bands show the SEM. **D.** Summary of the population average reward modulations (large - small). Solid line shows the modulation in the PD for neurons tuned to direction in the saccade task. Dashed line shows the average across all conditions for neurons that were not directionally tuned during the saccade. Gray line shows the average activity for all neurons during pursuit (tuning was not tested experimentally). Error bars show the SEM. The numbers of tuned and untuned neurons in each group for the saccade task are specified in E–J. In the pursuit task, the numbers of neurons were 108, 266, and 138 for the tagged, untagged, and area-untagged groups. **E-J.** Scatterplot showing the reward modulations of single neurons during the cue. Values show the average activity between 0 and 500 ms after cue appearance for the large (horizontal) and small (vertical) reward conditions. The color of the spot indicates whether the difference between reward conditions was significant (open circles, $p <0.05$, Wilcoxon rank-sum test) or not significant (gray spots, $p > 0.05$). Different columns correspond to the three populations of neurons. The N in the plots is the number of the neurons shown in the plot out of the total number of neurons in each population. Top row (**E-G**) shows activity in the PD of neurons that were directionally tuned. Bottom row (**H-J**) shows the average activity across all directions for neurons that were not directionally tuned.

Figure 5: Reward modulation of directional tuning aligned to movement. A-C. Average PSTH of the directionally tuned neurons for the three populations aligned to the saccade. Light and dark colors show activity in the PD and the direction opposite to the PD (null). Blue and red show the activity in the large and small reward conditions. Bands show the SEM. **D, H.** Summary of the population average reward modulations (large - small) for the saccade (**D**) and pursuit tasks (**H**). Solid line shows the PD for neurons directionally tuned during the tasks. Dashed line shows the average across all conditions for neurons that were not directionally tuned during movement. Error bars show the SEM. The numbers of tuned and untuned neurons in each group for the saccade task are shown in E–G. For the pursuit task, the numbers of tuned and untuned neurons were 39 and 69, 90 and 176, and 62 and 76 in the tagged, untagged, and area-untagged groups. **E-G.** Scatterplot showing the reward modulations of single neurons aligned to the saccade in the PD of the directionally tuned neurons. Values show the average activity from 100 before to 300 ms after the saccade in the large (horizontal) and small (vertical) reward conditions. The color of the spot indicates whether the difference between reward conditions was significant (open circles, $p <0.05$, Wilcoxon rank-sum test) or not significant (gray spots, $p > 0.05$). Different columns correspond to the three populations of neurons. Only directionally tuned neurons are shown, to contrast the results for this population with the cue results. The N in the plots is the number of the neurons shown in the plot out of the total number of neurons in each population.

Figure 6: Choice signals in the FEF. A-C. Population average of neurons that significantly differentiated between choice conditions (N= 26, 83 and 36 for tagged, untagged and area-untagged groups). Colored-coded lines correspond to the choice trials (black), large reward trials (blue) and small reward trials (red). Solid and dashed lines in the choice trials show activity in the preferred and null choice conditions. The line style (solid or dashed) in single-target trials indicates whether the target appeared at the location of the large reward on the preferred (solid) or null (dashed) choice trials. Because the monkeys rarely made errors on the choice trials, the line style also corresponds to the movement direction, i.e. monkey moved in the same direction in choice and single-target trials with the same line style. Different rows show different FEF populations. Gray bands show the SEM. **D, E.** The average distance (mean absolute value) between the choice in the preferred (D) or null (E) condition and single target conditions for all neurons (both significant and non-significant; N = number of neurons × 2, i.e., 216, 532, and 276 for tagged, untagged, and area-untagged groups). To allow for paired comparisons, the distance from choice to the corresponding single large condition was subtracted for each cell before the averages and SEM (bars) were calculated.

Editorial summary: Characterizing the information transmitted between structures is key to studying function. Using optogenetically identified neurons in monkeys, the authors show that cortical projections to the basal ganglia encode expected reward and selected actions.

Peer review information: *Nature Communications* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

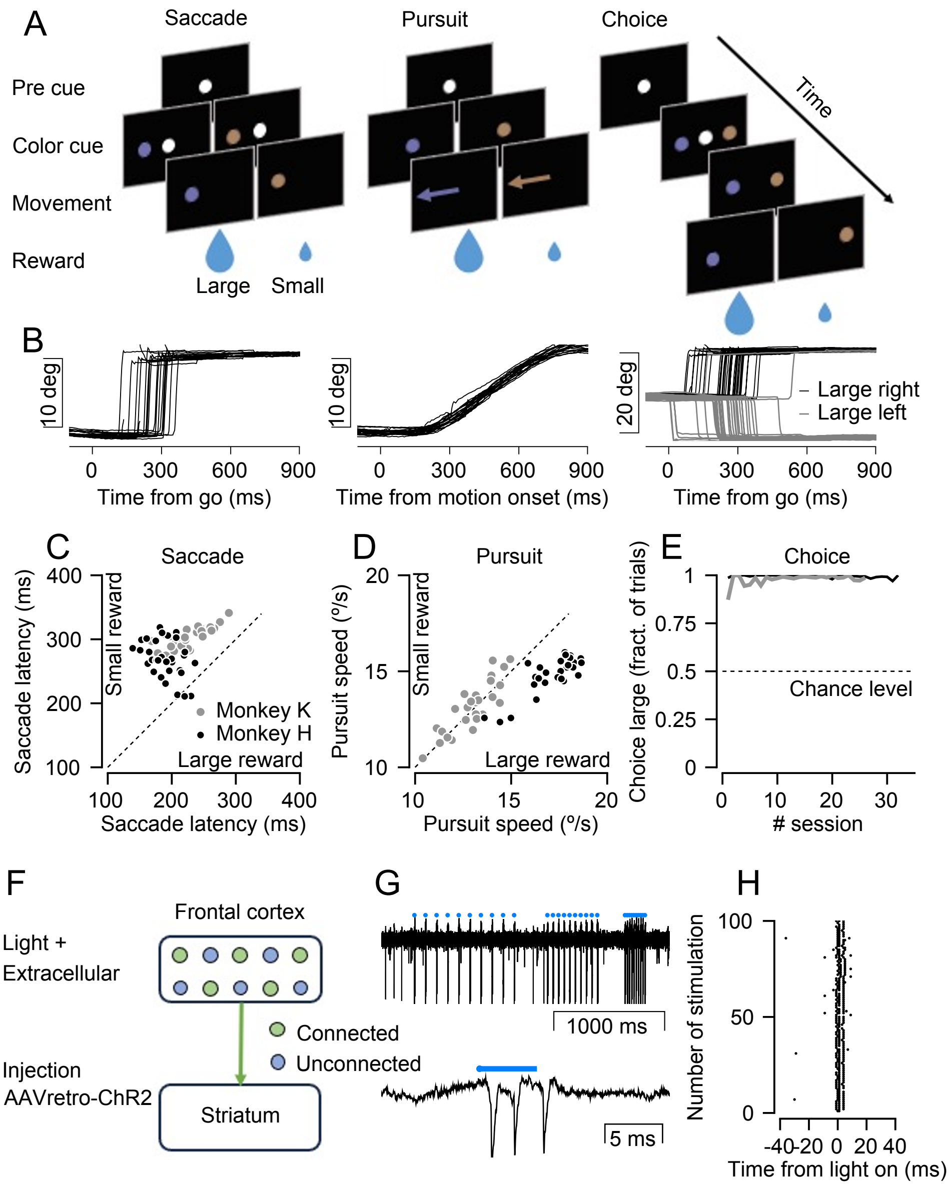
A
Saccade
Pursuit
Choice
Pre cue
Color cue
Movement
Reward
Large
Small
Time
B
10 deg
10 deg
20 deg
Large right
Large left
0
300
600
900
Time from go (ms)
Time from motion onset (ms)
Time from go (ms)
C
Saccade
Saccade latency (ms)
400
300
200
100
Small reward
Large reward
Monkey K
Monkey H
100
200
300
400
Saccade latency (ms)
D
Pursuit
Pursuit speed (°/s)
20
15
10
Small reward
Large reward
10
15
20
Pursuit speed (°/s)
E
Choice
Choice large (fract. of trials)
1
0.75
0.5
0.25
0
Chance level
0
10
20
30
session
F
Frontal cortex
Light +
Extracellular
Connected
Unconnected
Injection
AAVretro-ChR2
Striatum
G
1000 ms
5 ms
H
Number of stimulation
100
50
0
-40
-20
0
20
40
Time from light on (ms)

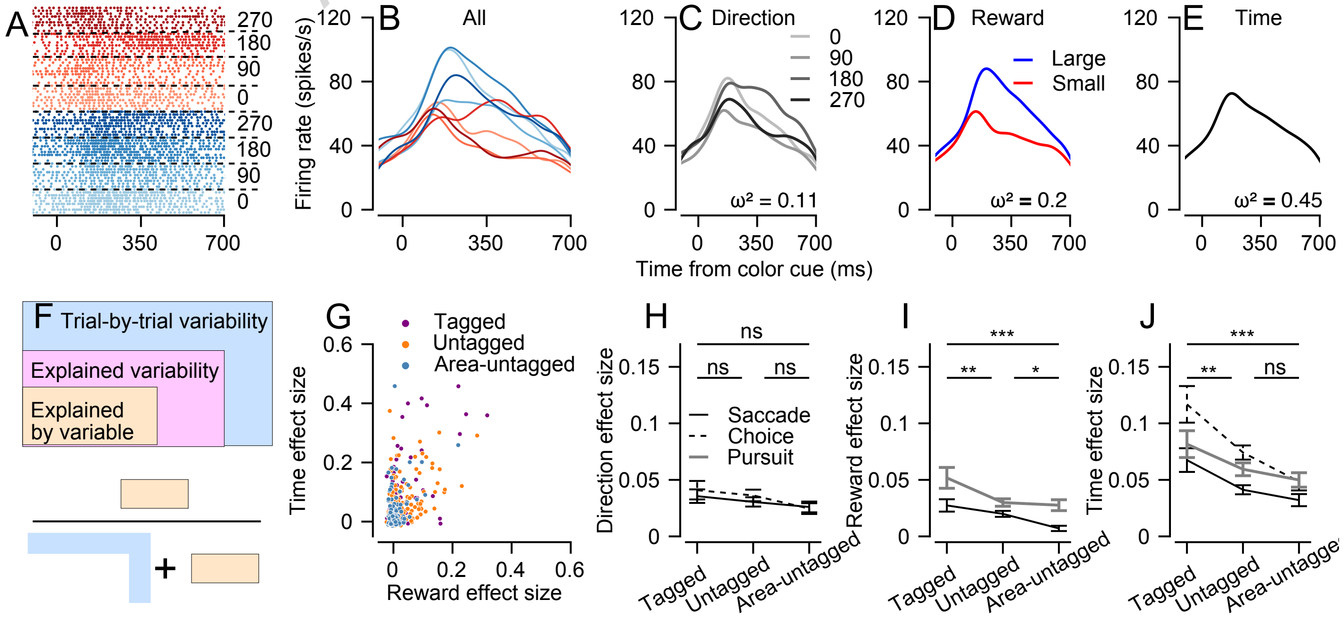
A
270
180
90
0
270
180
90
0
0
350
700
B
All
Firing rate (spikes/s)
120
80
40
0
0
350
700
C
Direction
0
90
180
270
ω² = 0.11
Time from color cue (ms)
D
Reward
Large
Small
ω² = 0.2
E
Time
ω² = 0.45
F
Trial-by-trial variability
Explained variability
Explained by variable
+
G
Tagged
Untagged
Area-untagged
Time effect size
Reward effect size
0.6
0.4
0.2
0
H
ns
ns
ns
Saccade
Choice
Pursuit
Direction effect size
0.15
0.1
0.05
Tagged
Untagged
Area-untagged
I
***
**
*
Reward effect size
J
***
**
ns
Time effect size

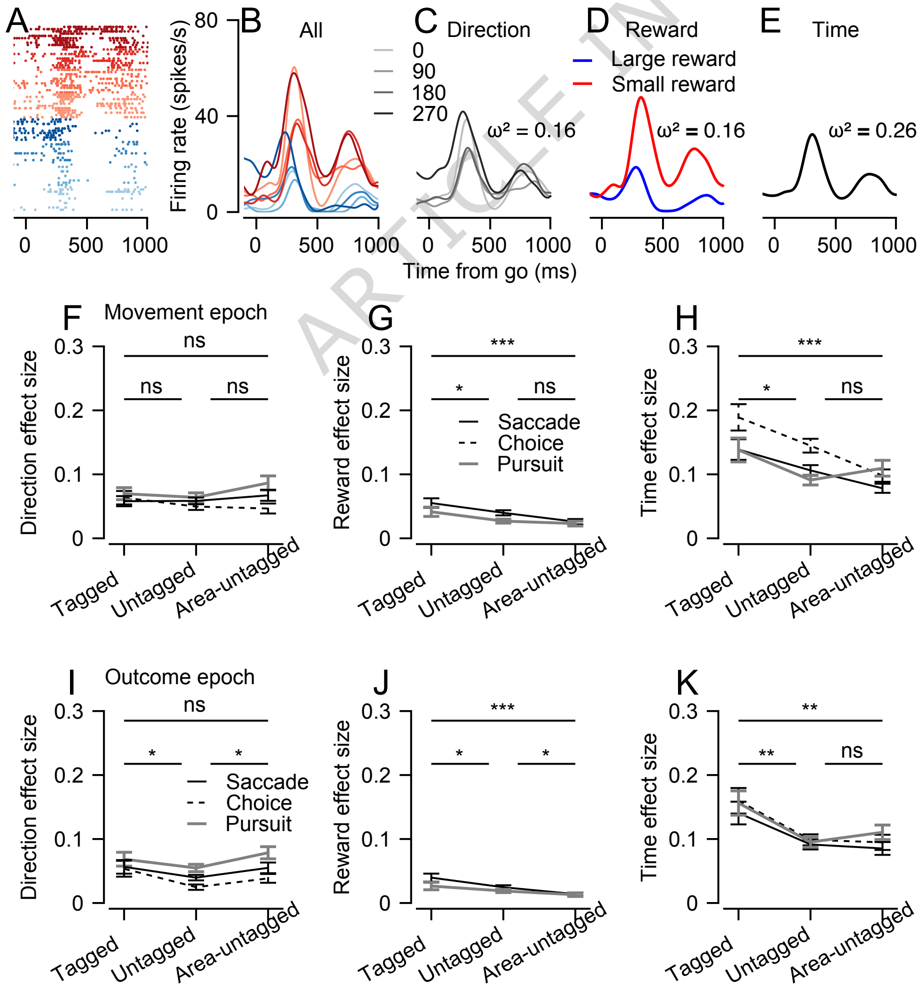
A
B
All
C
Direction
D
Reward
E
Time
Firing rate (spikes/s)
80
40
0
0
90
180
270
Large reward
Small reward
ω² = 0.16
ω² = 0.16
ω² = 0.26
0
500
1000
Time from go (ms)
F
Movement epoch
G
H
Direction effect size
Reward effect size
Time effect size
0.3
0.2
0.1
0
ns
*
**
***
Saccade
Choice
Pursuit
Tagged
Untagged
Area-untagged
I
Outcome epoch
J
K

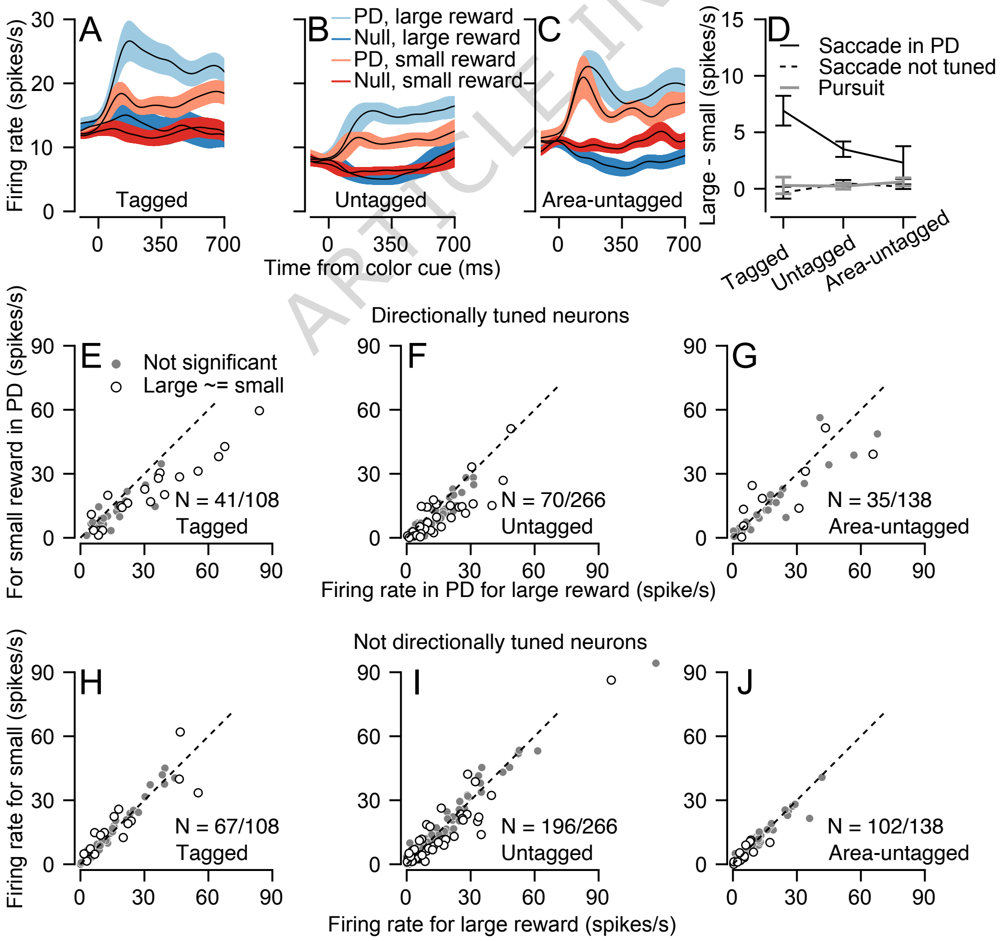
A
B
C
D
Firing rate (spikes/s)
30
20
10
0
Tagged
Untagged
Area-untagged
PD, large reward
Null, large reward
PD, small reward
Null, small reward
0
350
700
Time from color cue (ms)
Large - small (spikes/s)
15
10
5
0
Saccade in PD
Saccade not tuned
Pursuit
Tagged
Untagged
Area-untagged
Directionally tuned neurons
E
F
G
For small reward in PD (spikes/s)
90
60
30
0
Not significant
Large ~= small
N = 41/108
Tagged
N = 70/266
Untagged
N = 35/138
Area-untagged
Firing rate in PD for large reward (spike/s)
Not directionally tuned neurons
H
I
J
Firing rate for small (spikes/s)
N = 67/108
Tagged
N = 196/266
Untagged
N = 102/138
Area-untagged
Firing rate for large reward (spikes/s)

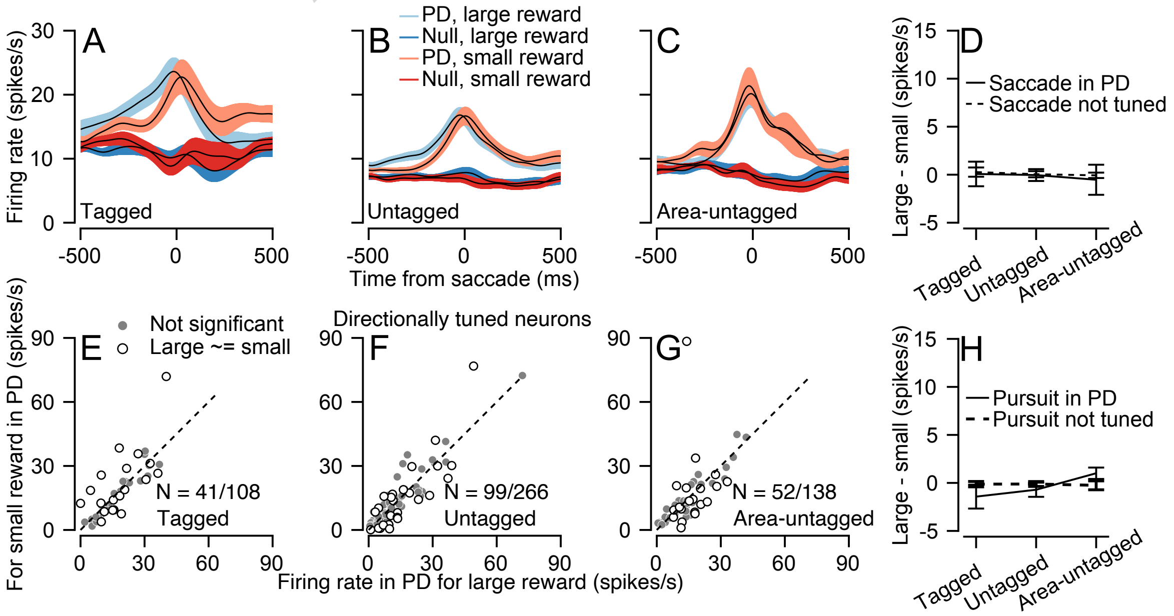
A
B
C
D
Firing rate (spikes/s)
30
20
10
0
Tagged
Untagged
Area-untagged
PD, large reward
Null, large reward
PD, small reward
Null, small reward
-500
0
500
Time from saccade (ms)
Large - small (spikes/s)
15
10
5
0
-5
Saccade in PD
Saccade not tuned
Tagged
Untagged
Area-untagged
E
F
G
H
For small reward in PD (spikes/s)
90
60
30
0
Not significant
Large ~= small
Directionally tuned neurons
N = 41/108
Tagged
N = 99/266
Untagged
N = 52/138
Area-untagged
Firing rate in PD for large reward (spikes/s)
Pursuit in PD
Pursuit not tuned

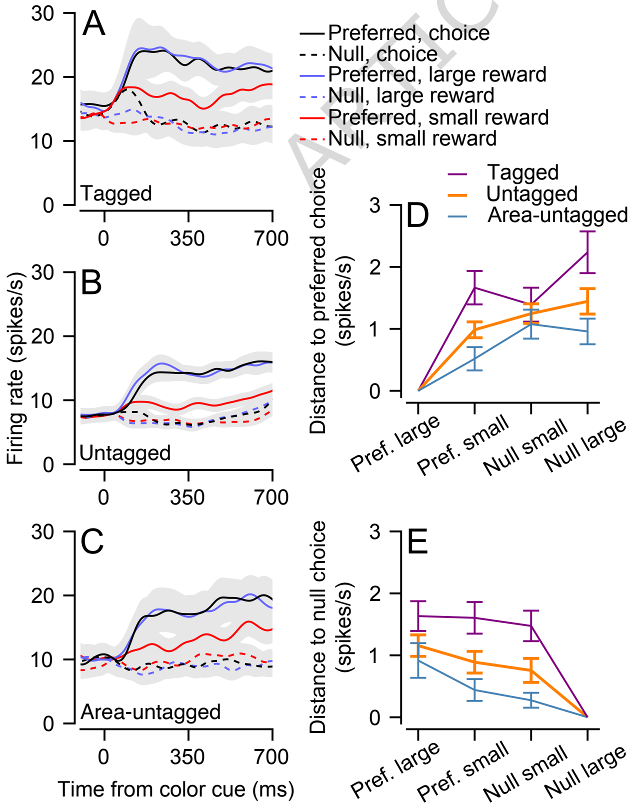
A
Tagged
B
Untagged
C
Area-untagged
Firing rate (spikes/s)
Time from color cue (ms)
0
350
700
0
10
20
30
Preferred, choice
Null, choice
Preferred, large reward
Null, large reward
Preferred, small reward
Null, small reward
D
Tagged
Untagged
Area-untagged
Distance to preferred choice (spikes/s)
E
Distance to null choice (spikes/s)
0
1
2
3
Pref. large
Pref. small
Null small
Null large