nature neuroscience

ARTICLES

https://doi.org/10.1038/s41593-019-0561-6

The neural basis of predictive pursuit

Seng Bum Michael Yoo[1*], Jiaxin Cindy Tu[1,4], Steven T. Piantadosi[2] and Benjamin Yost Hayden[1,3]

It remains unclear whether and, if so, how nonhuman animals make on-the-fly predictions during pursuit. Here we used a novel laboratory pursuit task that incentivizes the prediction of future prey positions. We trained three macaques to perform a joystick-controlled pursuit task in which prey follow intelligent escape algorithms. Subjects aimed toward the likely future positions of the prey, which indicated that they generate internal predictions and use these to guide behavior. We then developed a generative model that explains real-time pursuit trajectories and showed that our subjects use prey position, velocity and acceleration to make predictions. We identified neurons in the dorsal anterior cingulate cortex whose responses track these three variables. These neurons multiplexed prediction-related variables with a distinct and explicit representation of the future position of the prey. Our results provide a clear demonstration that the brain can explicitly represent future predictions and highlight the critical role of anterior cingulate cortex for future-oriented cognition.

Many foragers pursue fleeing prey. The ability to effectively pursue prey is therefore a critical element in our behavioral repertoires[1,2]. To effectively pursue, a forager needs to perform a series of computations; that is, it must maintain a representation of its current position relative to that of the prey, then compute the best path to capture the prey and then execute that path. Because the ability to perform such computations can determine foraging success, pursuit has likely been an important driver of our cognition and its underlying brain systems[3–5].

One way to improve pursuit effectiveness is to predict the future position of the prey and head toward the predicted position[6]. Estimating future positions can be done using the basic Newtonian variables of the prey (most importantly, its current position, velocity and acceleration) and can be improved using additional (potentially even recursive) variables, such as predictions about the likely evasive strategy of the prey in response to the future path of the predator. By using such information, the forager may be able to formulate a representation of the predicted future position of the prey. The ability of nonhuman animals to actively predict positions of prey during pursuit is poorly understood. Nonetheless, predictive pursuit is an important part of the repertoire of many species.

Prediction is important for many cognitive and behavioral processes, not just foraging. These include motor control, economic decision-making and abstract long-term planning[7–11]. There is some evidence to indicate that foraging animals can predict the long-term future; that is, they may be able to travel mentally in time and see themselves in the future[12,13]. However, observations about animal prediction tend to be limited to a small number of highly adapted species in unique contexts. And, while future planning of movements is relatively well studied, the ability to predict future positions of prey during dynamic behavior with rapidly changing goals—which feed into but are distinct from motor plans—is not. In the context of pursuit, a critical question is whether future-predicting foragers maintain a specific representation of potential future prey positions and whether those representations (assuming they exist) make use of specialized processes.

Although the neural bases of predictive pursuit remain unclear, we can draw some inferences about its likely neuroanatomy. In particular, the dorsal anterior cingulate cortex (dACC) has been implicated in prediction, prospection and related processes[14–17]. For example, neuroimaging studies indicate that the human dACC is a key region for economic prediction[18], for prospective reasoning[8] and for more open-ended prospective processes[18]. The dACC is well positioned for this role as it receives broad inputs from limbic and cognitive systems, integrates these and generates high-level control signals that regulate behavior in an abstract and high-level way[16,19–21].

Here, we examined the future predicting abilities of rhesus macaques using a novel virtual pursuit task. Subjects used a joystick to move an avatar in an open two-dimensional (2D) field displayed on a computer screen. Subjects, controlling the avatar, pursued a fleeing prey item that used an artificial intelligence algorithm to avoid predation. By examining the properties of a generative model fit to our data, we found that our subjects moved toward extrapolated future positions of prey rather than just pointing toward the present positions of the prey. Our subjects made their predictions based on three Newtonian variables associated with the current state, but not other factors that could further improve predictions (such as the effect of the movements of the subject on the future position of the prey). We also found that neurons in the dACC were selective for those three Newtonian variables (and not others), which indicates that responses in this region provide sufficient information to generate the types of predictions our subjects made. Finally, we found that dACC neurons used a spatial code to explicitly represent the predicted future position of the prey, and that this future representation is multiplexed with the representation of current Newtonian variables.

Results

Behavioral results. Three macaques (*Macaca mulatta*, subjects K, H and C) used a joystick to control the position of an avatar (a yellow or purple circle) that was moving continuously and smoothly in a rectangular field on a computer screen (Fig. 1; Supplementary Video 1; Methods). During each trial, subjects had up to 20 s to capture a prey item (a fleeing colored square) to obtain a juice reward. Prey avoided the avatar with a deterministic strategy that combined repulsion from the current position of the subject with repulsion from the walls of the field. The prey item was drawn randomly from

[1]Department of Neuroscience, Center for Magnetic Resonance Research, Minneapolis, MN, USA. [2]Department of Psychology, University of California Berkeley, Berkeley, CA, USA. [3]Center for Neuroengineering, University of Minnesota, Minneapolis, MN, USA. [4]Present address: Department of Neuroscience, Washington University in St Louis, St Louis, MO, USA. *e-mail: sbyoo.ur.bcs@gmail.com

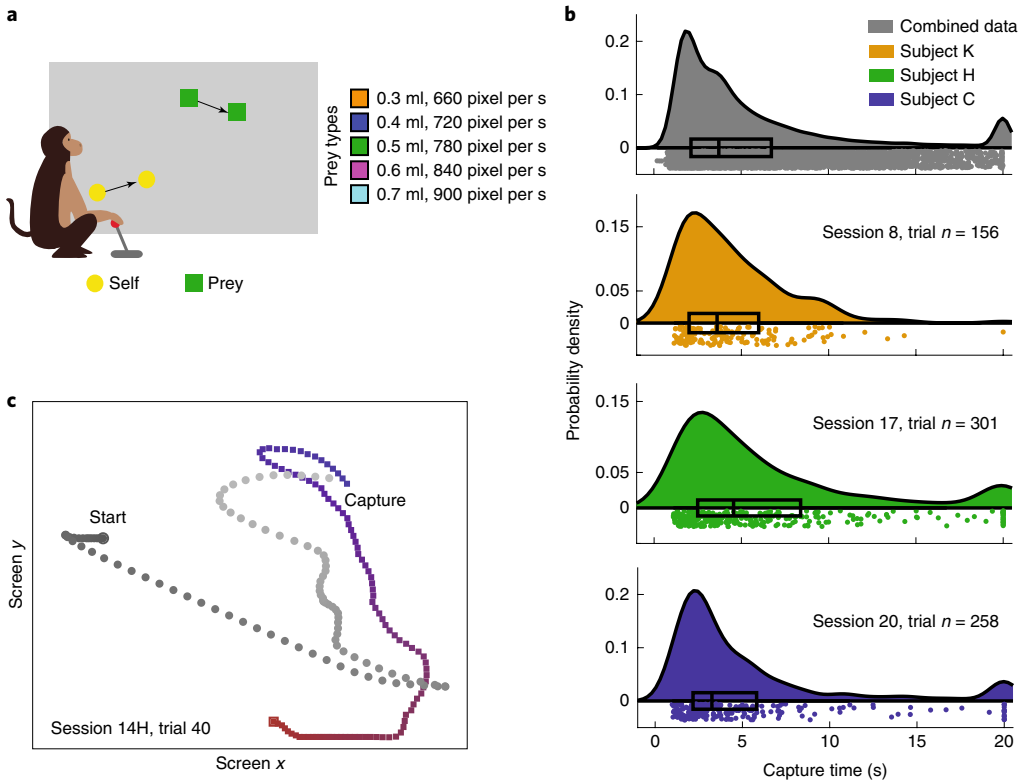

Fig. 1 | Experimental paradigm and behavioral results. a, Cartoon of the virtual pursuit task. The subject uses a joystick to control an avatar (circle) and to pursue prey (square) on a computer screen. **b**, Raincloud plot showing the capture times of each subject in an example session (limit was 20 s). The box plot indicates second and third quartiles of the data, and the midline indicates the median of the data (K: 3.36 s, H: 3.73 s, C: 3.93 s). The circles under the probability density functions indicate individual data points. **c**, Avatar and prey trajectories on example trials. Gray represents the path of the avatar; red–blue represent the path of the prey. The color gradient indicates time progression through the trial.

a set of five, identified by color, that differed in maximum velocity and associated reward size.

All subjects showed stable behavior within 12 2-h training sessions that followed a longer training period on joystick use (Supplementary Figs. 1 and 2). All data presented here were collected after the training sessions (the number of trials was 3,229 for K, 3,890 for H and 2,512 for C). Subjects successfully captured the prey in over 95% of trials and, on successful trials, did so in an average of 5.04 s (K = 4.26 s, H = 5.32 s and C = 5.54 s) and a median of 3.62 s (K = 3.36 s, H = 3.73 s and C = 3.93 s). The performance of the subjects varied lawfully with prey type, which indicates that the performance exhibited sensitivity to the manipulation of the reward and/or difficulty (Supplementary Fig. 1).

Behavioral evidence of future state prediction. For analysis purposes, we split all data into 1-s segments (Supplementary Fig. 3). Within each segment, we calculated the error (sum of squares) between the model (see below) and the behavior at each frame (that is, each 16.67 ms). For each segment, we computed the minimum point on a 201 × 201 matrix of intensities for each parameter pair (force by time; Fig. 2, see below). We then averaged over all segments and all trials separately for the three subjects.

We developed a generative model of behavior (see Methods). We used the variable τ (tau) to refer to the prediction parameter for each subject. The variable τ comes from the model and refers to a fit scalar variable, which is multiplied by the future position (see the equations in the section "Behavioral model" in the Methods). In practice, it can be interpreted as the distance into the future that the subject prospects to guide his behavior (Fig. 2a). The variable τ can have positive, negative or zero values. A positive value for τ indicates that the subject points toward the expected future position of the prey; that is, the strategy reflects the prediction. A zero τ indicates that the subject points the joystick directly at the current position of the prey. A negative value for τ indicates that the subject points the joystick toward where the prey was in the recent past. Note that all of these strategies (within limits) are capable of eventually catching all prey, since the avatar of the subject is, by design, faster than the prey. The scalar parameter κ (kappa) reflects the amount of force applied toward the direction of the predicted position. Thus, a negative value indicates that force is exerted away from (180° opposite) the position of the prey, whereas a positive value indicates that force is exerted toward it.

We also added an inertia term to the model. Specifically, we computed an inertially biased path for each 16.67 ms of the frame. The biased path is a vector sum of the computed the best predicted direction and the previous direction ($P_{\text{subject}}(t) - P_{\text{subject}}(t-1)$). In our implementation, these two terms have equal weighting. Note that in practice, their relative weighting may nonetheless vary because the force term (κ, which is fit in the model), affects the weight of the new direction relative to the past direction. This approach for implementing inertia is designed to intuitively align with how inertia works (see Methods; Supplementary Figs. 2 and 4).

We called our first model the 'physics-variable-based prediction model' (PVBP). It assumes that prediction by the subjects is derived from the current position, velocity (that is, both speed and direction) and acceleration (which includes both the direction and the magnitude of acceleration) of the prey, as well as further derivatives (Supplementary Fig. 5). For all three subjects, the best fitting τ was positive, which indicates that they point the joystick toward the future position of the prey. For ease of interpretation, we translated τ

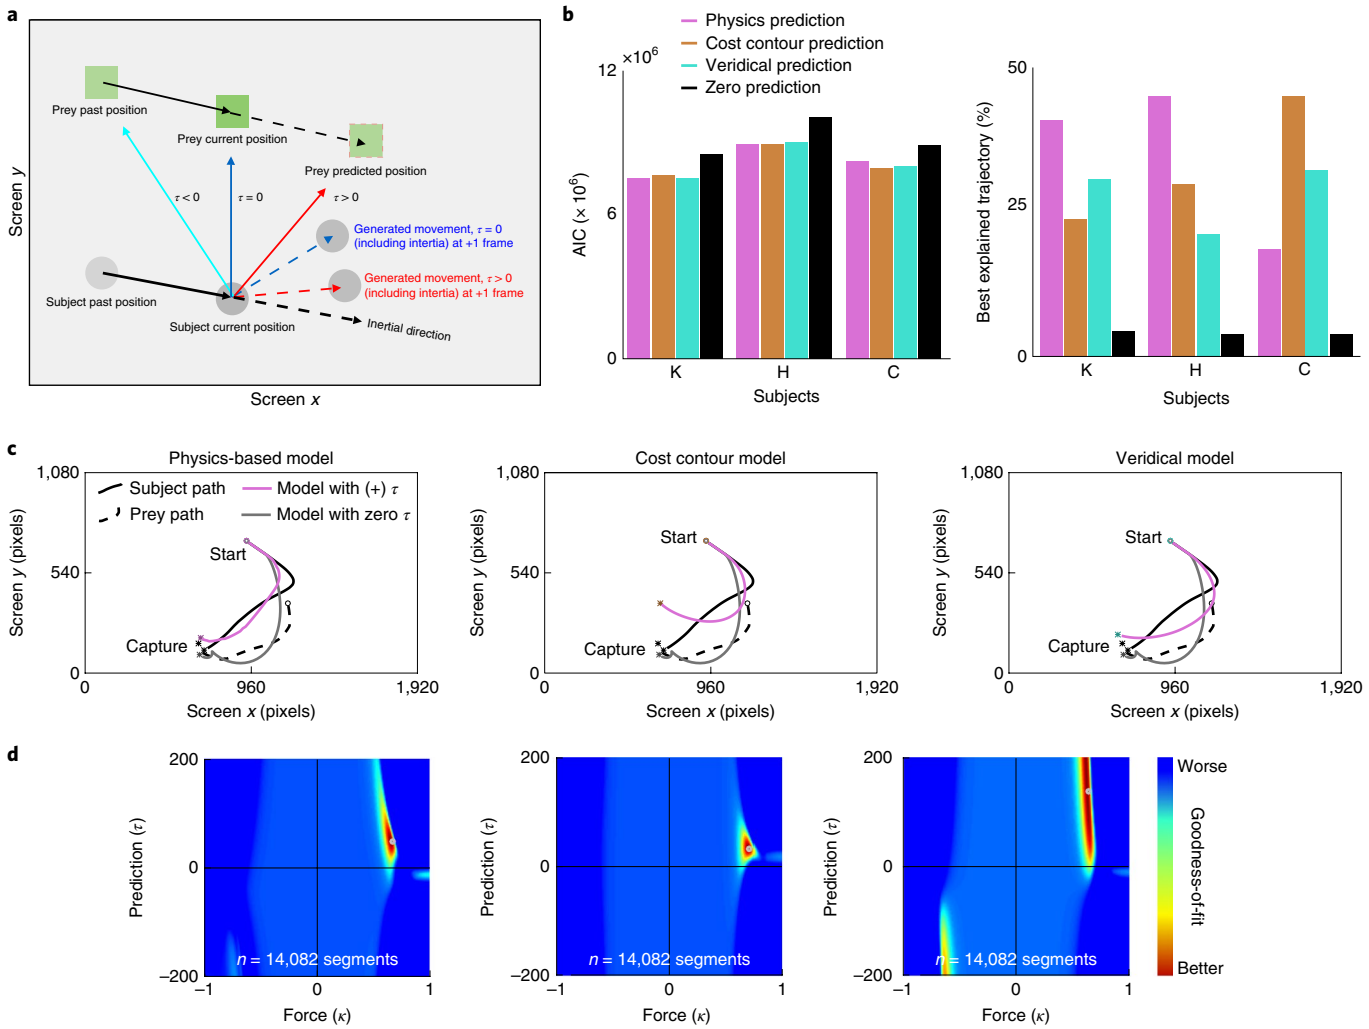

Fig. 2 | Model description and fitting results. a, Cartoon of the model for generating the future position based on prediction. The solid black arrow indicates movement from a previous time frame to the current one. The resulting movement (red and blue dashed arrows) vectors are constrained to a maximum speed and inertia (black dashed arrow). **b**, Fitting results based on AIC (left) across all trajectories and the percentage of trials best explained by each model (right). To calculate the AIC, we summed the log-likelihood across the entire dataset from each subject individually and used the quantity (twice the number of the segments) as the number of free parameters. The quantity of parameters was as follows: subject K = 28,164; subject H = 35,308; subject C = 20,720. The predictive models provide better fits than the zero-prediction ones. **c**, Example trajectories and corresponding fit trajectories generated by the predictive and nonpredictive models. **d**, Heatmap plots of model performance explaining the pursuit segment of the subject across the parameter space from a single subject (subject K) for the physics-based model (left), the cost-contour model (middle) and the veridical model (right). The small gray circle at the peak indicates the best parameter combination that explains that behavior of the subject; that is, the one that generates the smallest distance between the actual segment and the model-predicted segment.

into time units by calculating the distance between the current position and estimated position, then divided that quantity by the average velocity of the prey across the session. The results of this calculation indicated that subjects K, H and C pointed the joystick toward the position that the prey would occupy in an average of 800 ms, 767 ms and 733 ms in the future, respectively. In the context of the task, these numbers are substantial: they reflect 18.78%, 14.42% and 13.23% of the average trial duration for K, H and C, respectively.

To determine whether the positive prediction parameter τ is significantly greater than zero, we performed a bootstrap (randomization test making use of resampling with replacement) of heatmap slices from each segment (individual heatmap from 500 segments). This resampling was performed 500 times, and the resulting heatmaps were added. Then the τ and κ values that best explained each segment (that is, the ones resulting in the lowest cost) were selected in each resampling. We confirmed that the estimated values of τ and κ were both greater than zero more than 99% of the time (that is, $P < 0.01$).

The distance into the future that our subjects predicted did not detectably depend on the speed of the prey (linear regression between reward/speed and mean τ for K: $\beta = 3.0316$, $P = 0.1110$; for H: $\beta = 4.5798$, $P = 0.1791$; for C: $\beta = 7.1007$, $P = 0.0957$; the term β refers to the regression coefficient for speed against neural activity). We next asked whether taking more complex paths (ones with more turns versus more straight paths) affected prediction span. The complexity of the prey path (as measured by the path curvature, which was estimated using the average angle method) affected prediction. Specifically, subjects predicted less far into the future when the prey path had more curves (for K: $\beta = -0.0687$; for H: $\beta = -0.0567$; for C: $\beta = -0.0898$, with $P < 0.0001$ for each subject). Thus, subjects had the ability to dynamically adjust their own prediction in light of changing circumstances.

Alternative models do not predict trajectories as well as physics-based prediction. We next compared the physics-based model to two other models implementing different prediction algorithms (Fig. 2b). First, the veridical prediction (VP) model assumes that the subjects will make perfect predictions that incorporate all game dynamics, including repulsion of prey from the walls and the avatar of the subject. This means that a subject that makes a VP takes into account the effect his own movements will have on the strategy of the prey. Second, the cost-contour map prediction (CCMP) model is the same as the VP model, but excludes repulsion from the avatar, which means that the prediction model of the subject for the prey would not take into account their own motion. We compared the performance of each model by computing the sum of squares error between the prediction trajectory and the observed trajectories over all time bins.

Using the Akaike information criterion (AIC), we found that the PVBP model fit better than the other two models in our well-trained subjects (K: 7.529×10^6, for subject K, second best was VP: 7.542×10^6; H, PVBP: 8.923×10^6; for subject H, second best was CMPP: 8.950×10^6; Fig. 2d). We fit each segment with distinct τ and κ parameters, and we fit these same two parameters for each of our three models. As a consequence, the comparison of models can be directly done without concern of potential bias toward any specific model. In other words, by fitting each of the three models subject to identical constraints, we ensured a fair comparison across models. For the less well-trained subject, C, the CCMP model most accurately explained the trajectories (7.955×10^6).

We speculated that the speed of the prey might influence strategy. Indeed, we found that all three subjects used PVBP more frequently when the speed of the prey was faster (Supplementary Fig. 6). Note that this observed link between speed and the fit of the PVBP model occurred even in our third (less fully trained) subject ($P<0.001$, logistic regression; Supplementary Fig. 6). In any case, the classification of strategies using our models appears to be robust: the same results were obtained using a different method. Specifically, we fit all individual segments to the best model and computed the model that fit the most overall number of segments (Fig. 2d; Methods).

Doing the fitting this way may seem excessively flexible. That is, using two times the number of segments might allow us to fit only noise. Consider, for example, the case of fitting nine data points with ninth-order polynomial curve. Conversely, the extra freedom may allow us to better fit signal, or, of course, it may fit both noise and signal. The key question, then, is whether using a large number of parameters makes the fit better despite the possibility of fitting noise. To answer this question, we directly compared the two approaches (Supplementary Fig. 3). Specifically, we compared a model assigning two parameters globally versus one applying two parameters for each 1-s segment (that is, monkey K: 28,164; monkey H: 35,308; monkey C: 20,720 parameters; Supplementary Fig. 3). We then used AIC to compare models. We found that the second-by-second fitting resulted in lower AIC values, which implies a better fit, than the two-parameter counterpart. Specifically, in this figure, for all individual subjects, the change in AIC (AIC for global parameter model minus AIC for second-by-second model) was positive, which implies that the model fit by the second-by-second model explains the data better than the counterpart, for the best physics-based model explaining the behavior of the subjects.

Overall, the model comparison results showed that subjects predict the upcoming position of the prey using Newtonian physics but ignore the walls and their own influence on the prey. That is, subjects use a simplified approximation of the structure of the game to make future predictions; presumably this simplified one is sufficient to generate good predictions with lower mental effort costs. Indeed, the correlation between speed of prey and the reliance of the subjects on physics-based prediction (a result confirmed with two different analytical approaches) suggests that prediction might have a computational cost.

Prediction-related information encoded in the dACC. Based on its role as a nexus for motivational, cognitive and motor information[15,20], and its demonstrated role in human prospection[8,22], we hypothesized that the dACC would be critical for predictive pursuit (Fig. 3a). We fit a statistically unbiased linear–nonlinear generalized linear model (GLM)[23,24] to responses of 150 well-isolated dACC neurons (for K, $n=31$; for H, $n=119$). For this analysis, we focused on the entire trial period rather than preselecting epochs.

Position, velocity and acceleration of the prey were all encoded by significant proportions of neurons (Fig. 3; position: 62.00%, $n=93/150$; speed: 35.33%, $n=53/150$; 36.67%; direction: $n=55/150$, acceleration: 24.67%, $n=37/150$, $P<0.001$ in all cases, two-way binomial test). The model fit shown in magenta is the shape of reconstructed filter (see examples in Fig. 3c,d). According to the GLM, jerk, the derivative of acceleration, is not encoded (Supplementary Fig. 5). Jerk also did not measurably affect the neural responses of the subject (it only modulated 2.00% of cells, $n=3/150$, $P=0.1288$, two-way binomial test). Together, these results indicate that dACC ensembles carry the major raw ingredients that our subjects use to predict prey positions.

We wondered whether ostensible coding for prey variables could be the by-product of coding for self-position, since self-position and prey position do tend to be correlated. We therefore repeated our GLM analyses but included self-position, self-direction and self-speed as explanatory factors and considered variance explained by prey parameters only after accounting for these variables. Doing this, the proportion of neurons selective for the position information of the prey remained significant ($n=108/167$, 64.50%; $P<0.001$, two-way binomial test), as did neurons selective for prey speed ($n=31/167$, 22.16%; $P<0.001$, two-way binomial test) and direction ($n=20/167$, 11.98%, $P<0.001$, two-way binomial test). Note that for this analysis, we included a small subset of neurons, $n=17$, that were excluded from the previous analysis. Results did not differ if we used the same slightly smaller set of neurons.

Neurons in the dACC encode future position. We next asked whether dACC neurons encode the future position of the prey. For each neuron, we refit the GLM using an additional parameter: the position of the prey at time t in the future. We selected the time t ($t=833$ ms) that was most similar to the value of τ resulting from our generative model; that is, the one indicating the most likely time span of prediction (733, 766 and 800 for K, C and H, respectively) subject to the additional constraint of being a multiple of 166.67 (that is 10 frames). Note that although this value was chosen in advance, it aligns with the empirically derived measure of peak future position coding (Fig. 4d, see below).

Our analysis approach deals with the problem of correlation between the set of current Newtonian variables (including current position) and future position by assigning all explanatory power to the set of current variables first, and only counting as significant any additional variance explained by future position (see Methods). Despite this conservative criterion, we found that responses of 24.67% of dACC neurons were selective for the future position of the prey at time t ($n=37/150$).

Visual inspection of the filters of the neurons showed that their selectivity is complex (examples are shown in Fig. 4c). That is, they are positionally tuned, but, unlike place cells, have nonpoint-like shapes. They contained multiple peaks and there did not appear to be smooth gradients. Instead, they appeared to be heterogeneously spatially tuned. In this manner, they resemble recently identified non-grid-like space-selective cells in the entorhinal cortex[23]. Notably, conventional methods for detecting place/grid-like cells will greatly underestimate the proportion of such tuning.

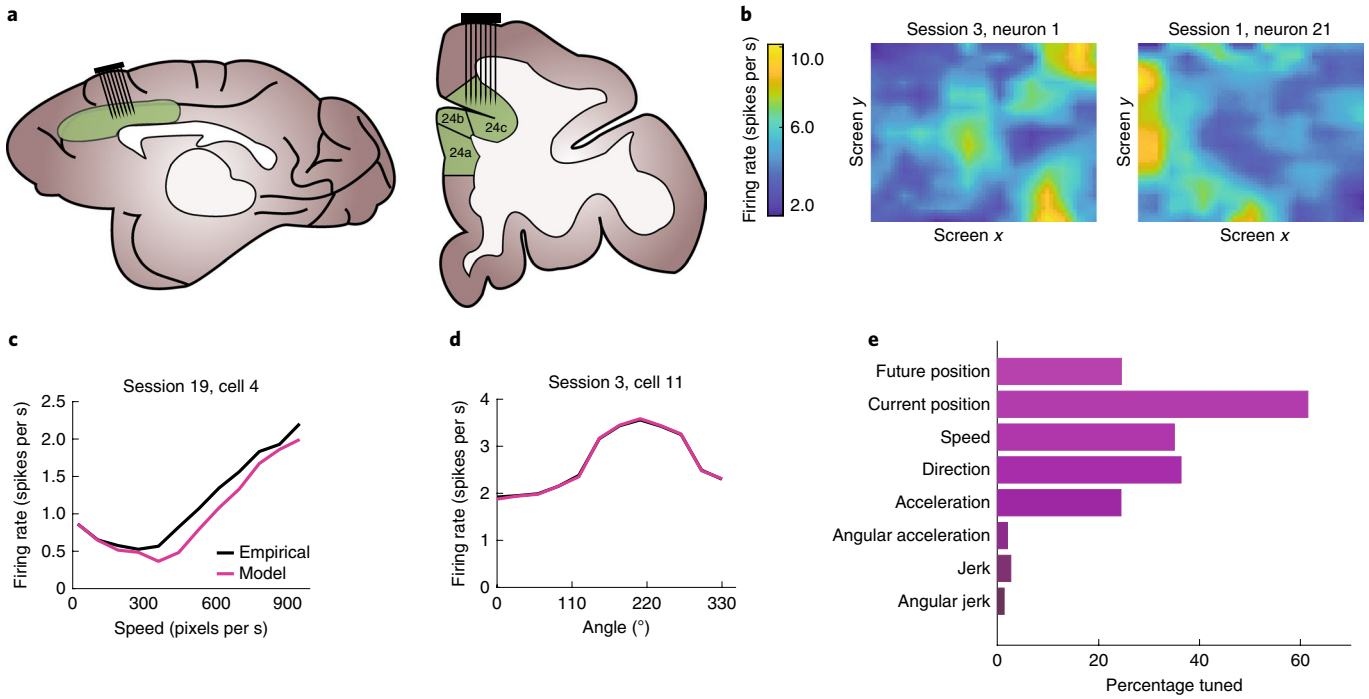

Fig. 3 | Basic neural results. a, Cartoon showing the location of the recorded brain areas in the dACC; sagittal (left) and coronal (right) views. **b**, Filters (tuning surfaces) of two example neurons showing selectivity for the current position of the prey. **c**, Example neuron showing tuning for speed, with the corresponding model fit. **d**, Example neuron showing tuning for prey direction (black line), with the corresponding model fit (pink line). **e**, Preponderance of tuning for the Newtonian physics variables tested. Tuning for future position is counted only if the neuron is selectively tuned for the future position above and beyond the current position.

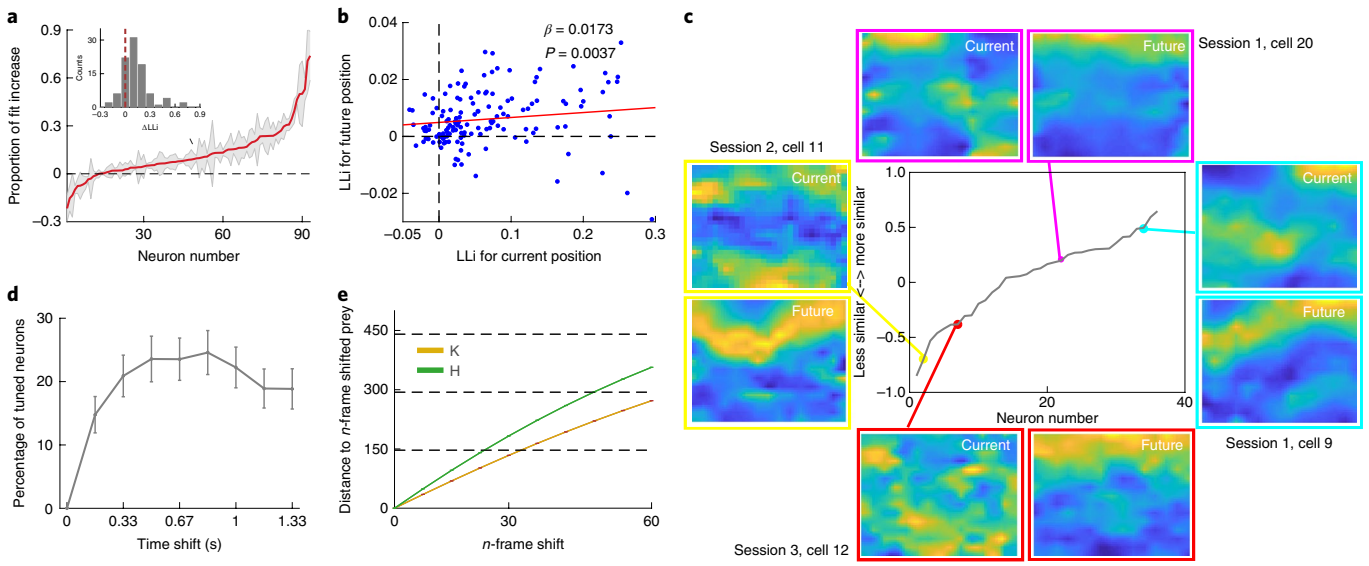

Fig. 4 | Properties of future position selectivity. a, The proportion of variance explained by including a future position as a regressor for spatial tuning for each neuron. Only neurons that are selective for the current position are shown. Neurons are sorted according to the amount of additional variance explained by the future position. Inset: histogram of LLi (a measure of explanatory power) observed after adding the future position of the prey to the GLM. **b**, LLi values for current and future position are correlated on a cell-by-cell basis. The red solid line indicates the linear regression line; $n = 150$ neurons; significance: two-tailed t-test. **c**, Example filters from neurons that were significantly tuned for both current and future prey position. The y axis of the central plot indicates the SPAEF, a measure of the similarity of 2D filter[26] (see main text). A higher SPAEF indicates that the matrices are more similar to each other. Only significant neurons are shown; cells are sorted by spatial efficiency. **d**, Sliding window analysis for future position encoding strength. The plot shows the proportion of neurons significantly selective for a future position at several possible future delays. This curve peaks at around 700–800 ms, which corresponds to the average prediction distance for all three subjects. **e**, The distance between current prey position and future prey position at time t rises roughly linearly with time. This finding indicates that the peaks found in **d** are not likely to be an artifact of some unforeseen periodicity in the relative paths of the subject and prey. The dashed line represents the maximum distance of 1 grid, which is the diagonal distance. This provides a reference for the units on the y axis. The measure of center is the mean value 50 bootstrap, and error bars represent the standard deviation.

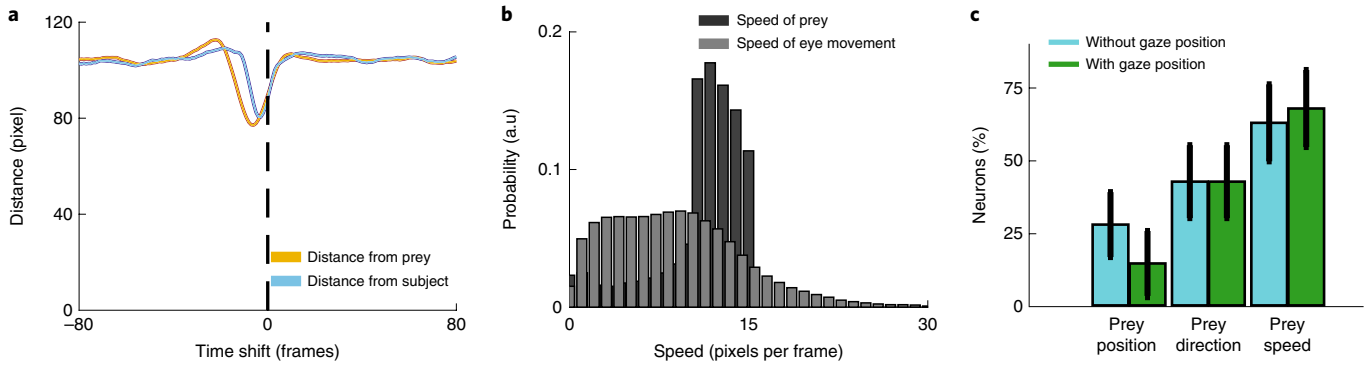

Fig. 5 | Analyses that control for potential gaze confounds. a, The Euclidean distance between the eye position at $t = 0$ and prey position (orange solid line) or self position (blue solid line). Shading represents s.e.m. Only data from subject K are shown ($n = 147{,}280$ frames). **b**, Speed distribution of prey movement and smooth eye pursuit. **c**, Proportion of neurons whose responses are selective for three key variables using our GLM procedure (cyan) and in an analysis in which all variance is assigned to eye position first (green). All three variables are still significant in the population when including gaze position (total number of neurons in the analysis is 31). The measure of center is the mean value computed by a 50× bootstrap; error bars represent the standard deviation.

We next asked how strongly dACC neurons encode the future position of the prey. We calculated the proportion of log likelihood increase (LLi) between the current position model and the current plus future model (Fig. 4a). Our neurons showed a wide range of marginal variance explained. On average, adding the future position term improved the variance explained by 6.89% (the mean of this proportion is significantly different from zero, $P < 0.001$, Wilcoxon sign-rank test, Fig. 4a, inset).

We then asked whether these newly discovered future position cells constitute a separate class of neurons from the cells that tracked the current position of the prey. To do this, we computed the explanatory variance accounted for by future position (variance explained by the combined model minus variance explained by current position) and current position, as defined by LLi in fitting. We found a positive correlation between these variables (Fig. 4b), which indicates that current and future position were multiplexed in the same population of cells ($r = 0.7394$, $P < 0.001$; see also ref. [25]).

To quantify the difference between current and future position coding, we fit separate models: one incorporated current position plus current Newtonian variables while the other was the same but used future position (assuming $t = 833$ ms) instead of current position. For the 36 neurons with significant tuning for both current and future position, we calculated the similarity between the filters, using a technique known as spatial efficiency (SPAEF)[26] (Fig. 4c). A zero SPAEF indicates orthogonal filters, whereas a positive SPAEF indicates similar filters and a negative SPAEF indicates anticorrelated filters. Although the mean of the SPAEF for our neurons was positive, it was not significantly so, and spanned a large range of values from negative to positive (mean of population SPAEF = 0.0440, Wilcoxon sign-rank test, $P = 0.3790$).

Finally, we assessed future encoding by examining the accuracy of model fitting to each of several possible future times, ranging from 0 to 1,333 ms in the future. We ran a type of sliding window analysis that involved sampling one frame (16.67 ms) every 10 frames (166.67 ms) and ignoring the intervening 9 frames. We found that the value of 833 ms fit the largest number of neurons (values around it fit many neurons too). Specifically, the plurality, 24.67% of neurons, were tuned for prey position at 833 ms (Fig. 4d). The roughly equivalent value of the neural and the behaviorally fit prospective distance (733 and 800 ms for those two subjects) suggests that these neurons encode the future position of prey on the same approximate timescale as the subject actively predicts.

We considered the possibility that this peak at 833 ms was due to an unanticipated correlation between positions in the future and at the present. If this were so, then the average distance of the self and/or prey would show a local minimum at a point in the future corresponding to the peak. However, we did not see this. On the contrary, we found that the distance increases monotonically for both subjects (Fig. 4e).

State information is not confounded with gaze information. Activity in the dACC is selective for saccadic direction and may therefore correlate with gaze direction (although this has not, to our knowledge, been shown[27]). Consequently, it is possible that our spatial kernels may reflect not task state but gaze information. Specifically, what appears to be tuning for future position may instead be attributable to the fact that monkeys looked toward the predicted future prey position. We tested this possibility by calculating the Euclidean distance between eye position and prey position in a range from −80 to +80 frames (Fig. 5a). The distance between eye and prey position was the closest at −5 frames (77.09 pixels), which indicates that eye position lagged the prey position. Thus, if gaze direction was a major confound, it would show up as increased selectivity for past positions, not prediction of future positions. Likewise, the chance that prey velocity encoding is a by-product of eye velocity encoding was belied by the stark differences between gaze speed and prey speed ($P < 0.001$, Wilcoxon sign-rank test, but also clear from visual inspection of Fig. 5b). Finally, we repeated our GLM analyses (see above) but included eye position (only for the one subject from which we collected gaze data). We found that that the number of tuned neurons for the prey did not substantially change; that is, adding in gaze position as a regressor did not qualitatively change our results (Fig. 5c).

Encoding of reward and reward proximity in the dACC. Research based on conventional choice tasks indicates that dACC neurons track values of potential rewards[28]. We next tested how the dACC encodes anticipated rewards in our more complex task. We found that averaging over all other variables, the value of the pursued reward modulates activity of 8.67% of neurons ($n = 13/150$, using a simple linear regression of firing rate against value; this proportion is greater than chance, $P = 0.038$, one-way binomial test). Note that this analysis ignores the potential encoding of prey speed, which is perfectly correlated with static reward in our task design. We hypothesized that reward/speed would be encoded in a modulatory manner[29]; that is, that the pursued reward/speed would alter the shape of the tuning for other task variables, rather than be multiplexed (Fig. 6a). To test this hypothesis, we split our dataset

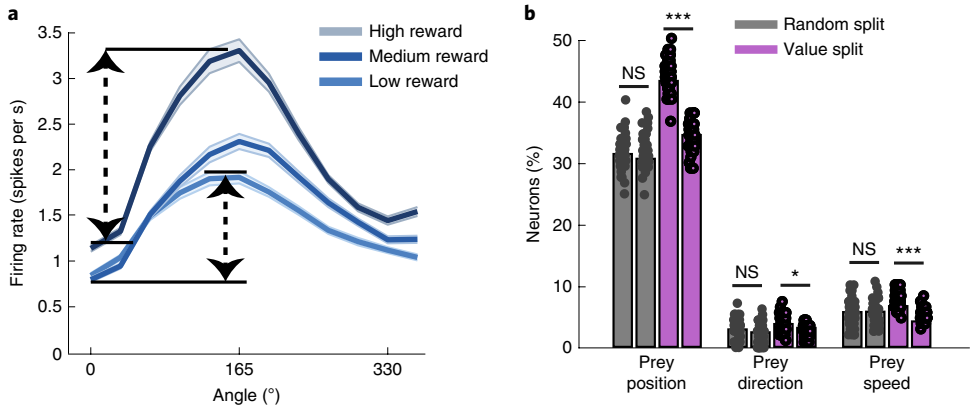

Fig. 6 | Modulatory effect of reward size on tuning for prey variables. a, Responses of an example neuron selective for the angle between self and prey; changes in the reward size of prey (divided into three bins post hoc) appear to change the gain and not the offset of the neurons; that is, reward interacts multiplicatively with angle. Shading represents the standard deviation. **b**, This pattern was also observed in the population. The proportion of neurons significantly tuned for prey variables (prey position, direction and speed) when splitting data randomly (gray bar) or according to the value of pursued prey (purple bar). The difference of value split was significant (*$P=0.0221$ for prey speed, and ***$P<0.001$ for other prey variables). Precise P values for each NS result are $P=0.1887$ (prey position), $P=0.1516$ (prey direction), $P=0.7039$ (prey speed). Significance test: two-tailed t-test. The measure of center is mean value computed by a 50× bootstrap; error bars represent the standard deviation. Each filled circle indicates a single bootstrap sample. NS, not significant.

by reward size and, as a control, split it randomly. We found that for several variables (prey position, direction and speed), value splits produced greater differences than random ones (purple bar, $P=0.0221$ for prey speed, and $P<0.001$ for other prey variables, Fig. 6b). This result indicates that the reward information encoded in dACC interacts mathematically with encoding of other variables. In other words, selectivity is mixed.

A good deal of research suggests that dACC neurons also signal the approach in time of impending rewards[30–32], even in continuous tasks[33]. We therefore asked whether it does so here. We repeated our GLM, including relative (self-to-prey) distance as an explanatory variable. We found that 38.67% of neurons ($n=58/150$) were tuned for self-prey distance. Interestingly, this relationship is heterogeneous; of these 58 neurons, 31.03% ($n=18/58$) showed a positive slope and 18.97 % ($n=11/58$) showed a negative slope. This bias was not itself significant ($P=0.2649$ for rise and fall bias, $n=18/29$; $P=1.000$ for monotonic bias, $n=29/59$, binomial test in all cases). This result indicates that while dACC neurons do track the approach to reward, they do not show an overall rise or fall in activity as they do so.

Discussion

Pursuit is an important element of the behavioral repertoire of many foragers[2,5,34]. The algorithmic bases of pursuit have recently attracted the interest of scholars in ecology, engineering, psychology and other disciplines[6,33,35–39]. Nonetheless, we know very little about how pursuit decisions occur in real time, and we know even less about their neuronal underpinnings. Here, we examined how macaques pursue virtual prey in a continuous, time-varying task. We developed a generative model based on a large dataset. The result of this model suggests that our subjects follow a predictive strategy. That is, instead of pointing toward the position of the prey, they extrapolate the future positions of prey and use this prediction to adjust their heading. This strategy is more efficient (yields more reward per unit time) but may be more computationally demanding than a simpler one that would involve pointing at and tracking the current position of the prey. These results demonstrate that pursuing animals can adopt complex future-predicting strategies that improve performance.

We found that dACC neurons track the elemental physical variables our subjects use to predict the future and explicitly encode the prediction. Specifically, we found that firing rate responses of neurons in dACC encode three Newtonian variables (position, velocity and acceleration) that our subjects used to track the prey and predict future prey positions. The same neurons carry an additional representation of the future position of the prey that is multiplexed with the Newtonian variables rather than maintained in a separate pool of specialized neurons. Both representations make use of a 2D response field, akin to place fields in hippocampus, but not localized to a single position. Specifically, spatial representation in the dACC is qualitatively similar to place representations of non-grid cells in the entorhinal cortex[23]. It is notable that the dACC uses partially distinct spatial tuning functions to track the present and future positions of the prey, thus in principle allowing unambiguous decoding for a given population response.

Our work is directly inspired by important studies that identified mechanisms underlying pursuit in other animals[36,37,40]. Our work goes beyond these studies by developing a generative model; that is, a model that seeks to understand how the data are generated[41]. One benefit of the generative model is that it lets us probe how the decision is made at every time step and make guesses about the underlying mental process leading to the decision. The generative model in turn is vital for extending our understanding of mechanism to the neuronal level.

This model enabled us to generate results that provide novel insight into the role of the dACC in cognition. First, our results emphasize the core role of the dACC in prediction, a role that is central to other theories, albeit not ones that directly involve pursuit[8,14,17,18,42,43]. One recent study is particularly relevant to these results[17]. The authors examined hemodynamic activity in the human dACC during a complex decision-making task in which subjects had to track previous rewards and use a reinforcement learning-like mechanism to formulate a future prediction and make the best choice. They found that the dACC tracks multiple variables, but was particularly selective for long-term estimates of expected prediction errors. These results highlight the key role of the dACC in prediction in general and suggest its role is conserved across species (see also ref. [17]). Second, our findings highlight the importance of the dACC to navigation. While studies of navigation typically focus on the medial temporal lobe, a growing body of work has begun to explore the role of cingulate cortex, which receives direct projections from medial temporal regions[22,44].

There are several important limitations to the current work. First, and most obviously, our subjects were not performing a

truly naturalistic task; they were performing a laboratory task that required specialized training. Future studies will be needed to ascertain whether these results relate to natural pursuit contexts that are ostensibly similar, such as pursuit of insects in the peripersonal space[45,46]. Second, and relatedly, the task space we used was greatly constrained—both agents were restricted to a small rectangular space and had strict speed limits. Subjects had full information about the position of the prey at all times. To more fully understand prediction, it will be critical to extend to contexts in which some information is hidden.

Traditional laboratory tasks that study topics of interest to cognitive neuroscience—decision-making and executive control—have discrete steps and force the brain to adjust to that structure[47]. One reason we developed the prey pursuit task is that it embeds those cognitive processes in a continuous time-varying task. Doing so allows us to study one of the greatest strengths of the brain—its ability to adjust and change its mind on the fly as new evidence comes in[47–50], and to incorporate that evidence into future plans.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41593-019-0561-6.

Received: 26 December 2018; Accepted: 20 November 2019;
Published online: 06 January 2020

References

1. Stephens, D. W. & Krebs, J. R. *Foraging Theory* (Princeton Univ. Press, 1986).
2. Ydenberg, R. C. & Dill, L. M. The economics of fleeing from predator. *Adv. Study Behav.* **16**, 229–249 (1986).
3. Broom, M. & Ruxton, G. D. Evolutionarily stable stealing: game theory applied to kleptoparasitism. *Behav. Ecol.* **9**, 397–403 (1998).
4. Cooper, W. E. & Frederick, W. G. Optimal flight initiation distance. *J. Theor. Biol.* **244**, 59–67 (2007).
5. Helfman, G. S. Threat-sensitive predator avoidance in damselfish–trumpetfish interactions. *Behav. Ecol.* **24**, 47–58 (2010).
6. Fujioka, E., Aihara, I., Sumiya, M., Aihara, K. & Hiryu, S. Echolocating bats use future-target information for optimal foraging. *Proc. Natl Acad. Sci. USA* **113**, 4848–4852 (2016).
7. Clayton, N. S., Bussey, T. J. & Dickinson, A. Opinion: can animals recall the past and plan for the future? *Nat. Rev. Neurosci.* **4**, 685–691 (2003).
8. Kolling, N., Scholl, J., Chekroud, A., Trier, H. A. & Rushworth, M. F. S. Prospection, perseverance, and insight in sequential behavior. *Neuron* **99**, 1069–1082.e7 (2018).
9. Merchant, H. & Georgopoulos, A. P. Neurophysiology of perceptual and motor aspects of interception. *J. Neurophysiol.* **95**, 1–13 (2006).
10. Seligman, M. E. P., Railton, P., Baumeister, R. F. & Sripada, C. Navigating into the future or driven by the past. *Perspect. Psychol. Sci.* **8**, 119–141 (2013).
11. Suddendorf, T. & Corballis, M. C. The evolution of foresight: what is mental time travel and is it unique to humans? *Behav. Brain Sci.* **30**, 299–351 (2007).
12. Raby, C. R., Alexis, D. M., Dickinson, A. & Clayton, N. S. Planning for the future by western scrub-jays. *Nature* **445**, 919–921 (2007).
13. Suddendorf, T. & Busby, J. Making decisions with the future in mind: developmental and comparative identification of mental time travel. *Learn. Motiv.* **36**, 110–125 (2005).
14. Alexander, W. H. & Brown, J. W. Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* **14**, 1338–1344 (2011).
15. Heilbronner, S. R. & Hayden, B. Y. Dorsal anterior cingulate cortex: a bottom-up view. *Annu. Rev. Neurosci.* **39**, 149–170 (2016).
16. Rushworth, M. F. S., Noonan, M. A. P., Boorman, E. D., Walton, M. E. & Behrens, T. E. Frontal cortex and reward-guided learning and decision-making. *Neuron* **70**, 1054–1069 (2011).
17. Wittmann, M. K. et al. Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nat. Commun.* **7**, 12327 (2016).
18. Brown, J. W. & Braver, T. Risk prediction and aversion by anterior cingulate cortex. *Cogn. Affect. Behav. Neurosci.* **7**, 266–277 (2007).
19. Bush, G., Luu, P. & Posner, M. I. Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn. Sci.* **4**, 215–222 (2000).
20. Paus, T. Primate anterior cinculate cortex: wher motor control, drive and cognition interface. *Nat. Rev. Neurosci.* **2**, 417–424 (2001).
21. Shenhav, A., Botvinick, M. M. & Cohen, J. D. Review the expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217–240 (2013).
22. Kaplan, R. et al. The neural representation of prospective choice during spatial planning and decisions. *PLoS Biol.* **15**, e1002588 (2017).
23. Hardcastle, K., Maheswaranathan, N., Ganguli, S. & Giocomo, L. M. A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron* **94**, 375–387.e7 (2017).
24. Pillow, J. W. et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
25. Blanchard, T. C., Piantadosi, S. & Hayden, B. Y. Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J. Neurophysiol.* **119**, 1305–1318 (2017).
26. Koch, J., Demirel, M. C. & Stisen, S. The SPAtial EFficiency metric (SPAEF): multiple-component evaluation of spatial patterns for optimization of hydrological models. *Geosci. Model Dev.* **11**, 1873–1886 (2018).
27. Hayden, B. Y. & Platt, M. L. Neurons in anterior cingulate cortex multiplex information about reward and action. *J. Neurosci.* **30**, 3339–3346 (2010).
28. Azab, H. & Hayden, B. Y. Correlates of economic decisions in the dorsal and subgenual anterior cingulate cortices. *Eur. J. Neurosci.* **47**, 979–993 (2018).
29. Yoo, S. B. M. & Hayden, B. Y. Economic choice as an untangling of options into actions. *Neuron* **99**, 434–447 (2018).
30. Shima, K. & Tanji, J. Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* **282**, 1335–1338 (1998).
31. Hayden, B. Y., Pearson, J. M. & Platt, M. L. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat. Neurosci.* **14**, 933–939 (2011).
32. Shidara, M. & Richmond, B. Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* **296**, 1709–1711 (2002).
33. Qi, S. et al. How cognitive and reactive fear circuits optimize escape decisions in humans. *Proc. Natl Acad. Sci. USA* **115**, 3186–3191 (2018).
34. Stephens, D. W., Brown, J. & Ydenberg, R. (eds) *Foraging: Behavior and Ecology* (Univ. Chicago Press, 2007).
35. Cliff, D. & Miller, G. Tracking the Red Queen: methods for measuring co-evolutionary progress in open-ended simulations. In *Proc. European Conference on Artificial Life* (eds Morán, F., Moreno, A., Merelo, J.J. & Chacón, P.) **929**, 200–218 (Springer, 1995).
36. Catania, K. C. Tentacled snakes turn C-starts to their advantage and predict future prey behavior. *Proc. Natl Acad. Sci. USA.* **106**, 11183–11187 (2009).
37. Mischiati, M. et al. Internal models direct dragonfly interception steering. *Nature* **517**, 333–338 (2015).
38. MacIver, M. A., Schmitz, L., Mugan, U., Murphey, T. D. & Mobley, C. D. Massive increase in visual range preceded the origin of terrestrial vertebrates. *Proc. Natl Acad. Sci. USA* **114**, E2375–E2384 (2017).
39. Mugan, U. & MacIver, M. A. The shift to life on land selected for planning. Preprint at *bioRxiv* https://www.biorxiv.org/content/10.1101/585760v1 (2019).
40. Lin, H., Leonardo, A., Lin, H. & Leonardo, A. Heuristic rules underlying dragonfly prey selection and interception article heuristic rules underlying dragonfly prey selection and interception. *Curr. Biol.* **27**, 1124–1137 (2017).
41. Iqbal, S. N. et al. Latent goal models for dynamic strategic interaction. *PLoS Comput. Biol.* **15**, e1006895 (2019).
42. Holroyd, C. B. & Coles, M. G. H. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* **109**, 679–709 (2002).
43. Matsumoto, M., Matsumoto, K., Abe, H. & Tanaka, K. Medial prefrontal cell activity signaling prediction errors of action values. *Nat. Neurosci.* **5**, 647–656 (2007).
44. Mashhoori, A., Hashemnia, S., McNaughton, B. L., Euston, D. R. & Gruber, A. J. Rat anterior cingulate cortex recalls features of remote reward locations after disfavoured reinforcements. *eLife* **7**, e29793 (2018).
45. Rothman, J. M., Raubenheimer, D., Bryer, M. A. H., Takahashi, M. & Gilbert, C. C. Nutritional contributions of insects to primate diets: implications for primate evolution. *J. Hum. Evol.* **71**, 59–69 (2014).
46. Sussman, R. W., Tab Rasmussen, D. & Raven, P. H. Rethinking primate origins again. *Am. J. Primatol.* **75**, 95–106 (2013).
47. Cisek, P. Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* **22**, 927–936 (2012).
48. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
49. Pezzulo, G. & Cisek, P. Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends Cogn. Sci.* **20**, 414–424 (2016).
50. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).

Methods

Subjects. All animal procedures were approved by the University Committee on Animal Resources at the University of Rochester and/or the University of Minnesota and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Three male rhesus macaques (*M. mulatta*) served as subjects for the behavior experiments; two of them also served as subjects for the physiology experiments. The subjects were aged 9 years (subject K), 10 years (subject C) and 10 years (subject H). The subjects had never previously been exposed to decision-making tasks in which they could use a joystick to pursue a moving prey. The previous training history for these subjects included two types of foraging tasks[51,52], intertemporal choice tasks[53], several types of gambling tasks[54–56], attentional tasks (similar to those in ref. [57]) and two types of reward-based decision tasks[58,59].

Experimental apparatus. The joystick was a modified version of commercially available joysticks with a built-in potentiometer (Logitech Extreme Pro 3D). The control bar was removed and replaced with a control stick (a 15-cm plastic dowel) topped with a 3-cm diameter plastic sphere designed to be easy for macaques to manipulate. The joystick position was read out by a custom-coded program in Matlab running on the stimulus-control computer. The joystick was controlled by an algorithm that detected the positional change of the joystick and limited the maximum pixel movement to within 23 pixels in 16.67 ms.

Task design. At the beginning of each trial, two shapes appeared on a gray computer monitor placed directly in front of the subject. The yellow (subject K) or purple (subjects H and C) circle (15-pixels in diameter) represented the subject. Subject position was determined by the joystick and was limited by the screen boundaries. A square shape (30 pixels in length) represented the prey. The movement of the prey was determined by a simple artificial intelligence algorithm (see below). Each trial ended with either the successful capture of the prey or after 20 s, whichever came first. Successful capture was defined as any spatial overlap between the avatar circle and the prey square. Capture resulted in an immediate juice reward, whereby the juice amount corresponded to prey color as follows: 0.3 ml for orange; 0.4 ml for blue; 0.5 ml for green; 0.6 ml for violet; and 0.7 ml for cyan.

The path of the prey was generated interactively using A-star pathfinding methods, which are commonly used in video gaming[60]. For every frame (16.67 ms), we computed the cost of 15 possible future positions the prey could move to in the next time-step. These 15 positions were equally spaced on the circumference of a circle centered on the current position of the prey, with a radius equal to the maximum distance the prey could travel within one time-step. The cost in turn was based on the following two factors: the position in the field and the position of the avatar of the subject. The field that the prey moved in had a built-in bias for cost, which made the prey more likely to move toward the center (Fig. 1b). The cost due to distance from the avatar of the subject was transformed using a sigmoidal function: the cost became zero beyond a certain distance so that the prey did not move, and it became greater as distance from the avatar of the subject decreased. Eventually, the costs from these 15 positions were calculated and the position with the lowest cost was selected for the next movement. If the next movement was beyond the screen range (1,920 × 1,080 resolution), then the position with the second lowest cost was selected, and so on.

The maximum speed of the subject was 23 pixels per frame (and each frame was 16.67 ms). The maximum and minimum speeds of the prey varied across subjects and were set by the experimenter to obtain a large number of trials (Fig. 1). Specifically, speeds were selected so that subjects could capture prey on <85% of trials; these values were modified using a staircase method. If subjects missed the prey three times consecutively, then the speed of the prey was reduced. Once the subject intercepted the prey in a trial where the staircase method was used, then the selection of prey speed was randomized again. To ensure sufficient time of pursuit, the minimum distance between the initial position of each subject avatar and prey was 400 pixels.

Training level estimation. Three subjects were trained for the same amount of time (8 weeks). As training progressed, each subject was exposed to a progressively more difficult (faster) suite of prey, up to a fixed maximum. Subjects K and H reached a similar range for the maximum speed of prey during the training period (K:15 pixels per frame; H: 14 pixels per frame). However, subject C only attained a maximum speed of 8 pixels per frame (Supplementary Fig. 6). It is for this reason we refer to him as the less well-trained subject.

Behavioral model. To fit each the movement of each subject, each trial was divided into 1-s-long segments. Each segment included 61 data points (because we used 16.67-ms resolution). We modeled these trajectories using a single prediction and a single force parameter for the entire trial as a simplifying assumption. Nonetheless, it is reasonable to assume that throughout a long 20-s period there would be an active adjustment of prediction and force. An actual comparison using AIC supported our intuition, and we used segment as the unit of analysis throughout (values of AIC of segment/AIC of trial was 0.9328, 0.9214 and 0.9227 for subjects K, H and C (or whatever), respectively.

Overall, the position of the subject was generated according to the following:

$$P_{\text{subject}}(t+1) = \boldsymbol{\kappa} f(P_{\text{subject}}(t)) + m$$

where $P_{\text{subject}}(t)$ is position of the subject at time t, m is the inertia of the subject as calculated from the joystick and κ is the force parameter. The vector $\boldsymbol{\kappa} f(P_{\text{subject}}(t))$ was then summed with the inertia m that was defined as follows:

$$m = P_{\text{subject}}(t) - P_{\text{subject}}(t-1)$$

$P_{\text{prey}}(t)$ indicates the position of the prey at time t. The function with respect to subject position at time t was defined as follows:

$$f(P_{\text{subject}}(t)) = P_{\text{prey}}(t+1) - P_{\text{subject}}(t)$$

Then the position of the prey at time $t+1$ was as follows:

$$P_{\text{prey}}(t+1) = P_{\text{prey}}(t) + \tau \sum_{n=0}^{k} \frac{\mathrm{d}^n P_{\text{prey}}(t)}{\mathrm{d}t^n}$$

Where the n indicates the order of derivation with respect to the time. Thus, $n=1$ indicates the velocity and $n=2$ indicates the acceleration.

The PVBP model incorporates one previous time step to predict the next position of the prey. This approach is similar to a Kalman filter[61]. The other two models we tested do not utilize any past information. The model assuming prediction using the CCMP model considers only the lowest cost location at the next time step. The model assuming VP automatically finds the exact position of the prey at the next time step. Once the position of the prey on the next time step is predicted, the model computes how far this predicted position is from the current position of the agent. A prediction value of 1 indicates that the future position will be as far as from the current position of the agent as the current position of the prey. The best-fitting parameter pairs were determined by performing a grid search across the ranges of both parameters.

During this search, we tested the range of the prediction parameter between −400 and 400 for subjects H and C, and between −200 and 200 for subject K. Units for this range correspond to the distance the prey moved in the previous time step. Subjects H and C had a larger range because over 5% of their trajectories resulted in either −200 or 200 in the prediction parameter value. Representative parameters for explaining each segment were selected based on the value of the sum of squared error between the actual segment and the segment generated by the model.

Significance testing for estimating parameters of the behavioral model. To determine whether the positive prediction parameter was significantly greater than zero, we performed a bootstrap analysis of heatmap slices from each segment. This resampling was performed 500 times, and selected heatmaps were added. Then, the parameter set resulting in the lowest cost was selected in each resampling.

Model evaluation. To evaluate model performance and compare among models, we computed the AIC using the likelihood of each model (Fig. 2; Supplementary Figs. 4 and 5). We first calculated the mean and variance of all the sum of squared errors across trajectories. Then we estimated the likelihood assuming a normal distribution centered on the mean of the sum of squared errors with a variance equivalent to the variance of the sum of squared errors across all trajectories. To validate whether subjects used a single prediction and force across all the trials or adaptively changed their prediction method, we compared the AIC value between cases for which the parameter pair varied across all trajectories using only the single best parameter pair.

Electrophysiological recording. One subject (H) was implanted with multiple floating microelectrode arrays (Microprobes for Life Sciences) in the dACC. This is the region that we define as Area 24 (ref. [15]) and corresponds to the dACC in most other primate studies, including those from our laboratory[52,54,62]. Each floating microelectrode array had 32 electrodes (impedance of 0.5 MOhm, 70% Pt, 30% Ir) of various lengths to reach multiple layers within the dACC. Neurons from subject K were recorded with laminar V probes (Plexon) that had 24 contact points with 150 µm of inter-contact distance. Continuous, wideband neural signals were amplified, digitized at 40 kHz and stored using the Grapevine Data Acquisition System (Ripple). Spike sorting was done manually offline (Plexon Offline Sorter). Spike sorting was performed blind to any experimental conditions to avoid bias.

Details of the linear–nonlinear model. To test the selectivity of neurons for various experimental variables, we constructed GLMs with navigational variables[23,24]. The GLMs estimated the spike rate (r_i) of one neuron during time bin t as an exponential function of the weighted sum of the relevant value of each variable at time t, for which the weights are determined by a set of coefficients (w_i). The estimated firing rates from the GLMs can be expressed as follows:

$$\mathbf{r} = \exp\left(\sum_i X_i^T w_i\right)/\mathrm{d}t$$

Where $\mathbf{r}$ denotes a vector of firing rates for one neuron over T time points across the session, and i indexes the variables of interest, for example, the position of

the avatar on the screen. The vector of firing rates over T time points provide the benefit for modeling the neural activity without the need of specifically time-locking to a behavioral event. X_i is a matrix in which each column represents a set of 'state variables' of the animal (for example, 1 of 12 speeds, determined by post hoc binning) obtained from binning the continuous variable so that all the columns for a particular row are 0, except for one column. Unlike conventional tuning curve analysis, GLM analysis does not assume the parametric shape of the tuning curve a priori. Instead, the weights, which define the shape of tuning for each neuron, were optimized by maximizing the Poisson log-likelihood of the observed spike train given the model-expected spike number, with additional regularization for the smoothness of parameters in a continuous variable, and a lasso regularization for parameters in a discrete variable. Position parameters were separately smoothed across rows and columns. The regularization hyperparameter was chosen by maximizing the cross-validation log-likelihood based on several randomly selected neurons. The unconstrained optimization with gradient and Hessian was performed (Matlab fminunc function). The model performance of each neuron was quantified by the log-likelihood of held out data under the model. This cross-validation procedure was repeated ten times (tenfold cross-validation), and overfitting was penalized. Through multiple levels of penalties, we compared the performance of models with varying complexity.

Forward model selection. Model selection was based on the cross-validated log-likelihood value for each model. We first fit n models with a single variable, where n is the total number of variables. The best single model was determined by the largest increase in spike-normalized log-likelihood from the null model (that is, the model with a single parameter representing the mean firing rate). Then, additional variables ($n-1$ in total) were added to the best single variable model. The best two-variable model was preferred over the single variable model only if it significantly improved the cross-validation log-likelihood value (Wilcoxon signed-rank test, $\alpha = 0.05$). Likewise, the procedure was continued for the three-variable model and beyond if adding more variables significantly improved the model performance, and the best, simplest model was selected. The cell was categorized as not tuned to any of the variables considered if the log-likelihood increase was not significantly higher than baseline, which was the mean firing rate of fitted neurons across the session.

Future position models. We examined the effect of future position by fitting a GLM having 'future position' and 'current position' together as the input variable. Then we compared this model to the GLM model with only the current position as the input. The difference between the two models provided evidence that the additional variance was explained by including future position.

Comparison between current and future position filters. For this purpose, we constructed two GLMs: one with current position and current Newtonian variables (velocity and acceleration), and another with future position and current Newtonian variables. Then we selected the neurons that showed significant tuning for both models. To compare the similarity between two positional filters, we used the SPAEF that prior literature suggests to be more robust than the 2D spatial correlation[26]. It quantifies the similarity between two maps as follows:

$$\mathrm{SPAEF} = 1 - \sqrt{(A-1)^2 + (B-1)^2 - (C-1)^2}$$

where A is the Pearson correlation between two maps, B is the ratio between the coefficients of variation for each map and C is the activity similarity measured by histogram profiles. Values near -1 indicate anticorrelated maps (one tends to be high when the other is low), 0 indicates uncorrelated maps and 1 indicates perfect matching between the two.

Velocity-dependent PVMP prediction bias. We examined whether PVBP is preferred when the velocity of prey is high (Supplementary Fig. 6). We first obtained the average velocity of the prey at each segment, and then categorized each segment as belonging to either the physics or non-physics variable-based prediction based on which fit result was best. With the prey velocity and segment category, we performed logistic regression with velocity as a predictor and category as the dependent variable (glmfit in Matlab).

Statistics. Error bars of log-likelihood fit increase in the neural analysis were obtained using a bootstrapping cross-validation procedure (50 times; Fig. 4a). Error bars in the percentage of tuned neurons were obtained by calculating the mean for 50-time bootstrapping processes (Figs. 4d, 5c and 6b). Other error bars represent the standard error of the mean (s.e.m.) (Figs. 5e and 6a; Supplementary Fig. 1). No statistical methods were used to predetermine sample sizes, but our sample sizes were similar to those reported in previous publications both from our laboratory and other laboratories (ref. [41] for the behavioral model, refs. [51,52] for the neural data). Data were assumed to be normally distributed, but this was not formally tested. Trial conditions were randomly generated. The selection of subjects was arbitrary. Data collection and analyses were not performed blinded to the conditions of the experiments. No data points were excluded.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

A portion of the data is available on Github (https://github.com/sbyoo/prospectpursuit/). Full data are available from the corresponding author upon reasonable request.

Code availability

Code is available at https://github.com/sbyoo/prospectpursuit/.

References

51. Blanchard, T. C. & Hayden, B. Y. Monkeys are more patient in a foraging task than in a standard intertemporal choice task. *PLoS One* **10**, e0117057 (2015).
52. Blanchard, T. C., Strait, C. E. & Hayden, B. Y. Ramping ensemble activity in dorsal anterior cingulate neurons during persistent commitment to a decision. *J. Neurophysiol.* **114**, 2439–2449 (2015).
53. Blanchard, T. C., Pearson, J. M. & Hayden, B. Y. Postreward delays and systematic biases in measures of animal temporal discounting. *Proc. Natl Acad. Sci. USA* **110**, 15491–15496 (2013).
54. Azab, H. & Hayden, B. Y. Correlates of decisional dynamics in the dorsal anterior cingulate cortex. *PLoS Biol.* **15**, e2003091 (2017).
55. Strait, C. E. et al. Neuronal selectivity for spatial position of offers and choices in five reward regions. *J. Neurophysiol.* **115**, 1098–1111 (2016).
56. Heilbronner, S. R. & Hayden, B. Y. The description-experience gap in risky choice in nonhuman primates. *Psychon. Bull. Rev.* **23**, 593–600 (2016).
57. Hayden, B. Y. & Gallant, J. L. Working memory and decision processes in visual area V4. *Front. Neurosci.* **7**, 18 (2013).
58. Sleezer, B. J. & Hayden, B. Y. Differential contributions of ventral and dorsal striatum to early and late phases of cognitive set reconfiguration. *J. Cogn. Neurosci.* **28**, 1849–1864 (2016).
59. Wang, M. Z. & Hayden, B. Y. Reactivation of associative structure specific outcome responses during prospective evaluation in reward-based choices. *Nat. Commun.* **8**, 15821 (2017).
60. Hart, P. E. & Nils, J. Formal basis for the heuristic determination of minumum cost path. *IEEE Trans. Syst. Sci. Cyber.* **4**, 100–107 (1968).
61. Kalman, R. E. A new approach to linear filtering and prediction problems. *J. Basic Eng.* **82**, 35 (1960).
62. Blanchard, T. C. & Hayden, B. Y. Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J. Neurosci.* **34**, 646–655 (2014).

Acknowledgements

The authors thank A. Thomé for his critical role in designing the task, for devising the training protocols and for developing our joysticks. They thank M. Mancarella for his critical help with joystick training, and appreciate invaluable help from M. Schieber, A. Rouse and S. Heilbronner. This work was supported by an award from the Templeton Foundation to B.Y.H. and by an R01 from NIDA (DA038615).

Author contributions

S.B.M.Y. and B.Y.H. conceptualized and designed the experiment. S.B.M.Y. collected the data. S.B.M.Y. and S.T.P. developed the behavioral model, S.B.M.Y., J.C.T. and B.Y.H. developed the physiological model and analyzed the data. S.B.M.Y. and B.Y.H. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41593-019-0561-6.

Correspondence and requests for materials should be addressed to S.B.M.Y.

Reprints and permissions information is available at www.nature.com/reprints.

nature research

Corresponding author(s): Seng Bum Michael Yoo

Last updated by author(s): Nov 7, 2019

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Matlab(MATLAB 2016a,MATLAB 2017b)and python(ver 3.5.2), Plexon Offline sorter 3 and 4, Grapevine software 1.10 is used.
Data analysis	Matlab(MATLAB 2016a,MATLAB 2017b)and python(ver 3.5.2) is used.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The data that support the findings of this study are available from the corresponding author upon request

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	The data is quantitative data collected by computerized experiment.
Research sample	Three male rhesus macaque was tested repetitively.
Sampling strategy	For the macaques, two different types of sample sizes were used. First, total number of monkeys: n = 2. This number was picked based on standard practices in electrophysiology literature. Second, total number of trials or repetitions for the behavioral and neural data. The number of trials were predefined, but were compared with the references cited.
Data collection	The data was collected from chaired rhesus macaque with eyetracker (Eyelink), customized joystick (Logitech), neural recording probes (Plexon, Microprobe), and neural recording device (Plexon, Ripple). The computerized task was operated by MATLAB 2016a and MATLAB 2017b.
Timing	Collection from initial subject started in 2016-Jan and the collection from last subject ended in 2017-Dec.
Data exclusions	No data was excluded.
Non-participation	No subjects were excluded.
Randomization	N/A for experimental group, but trial generation was randomized.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☐	☒ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research

Laboratory animals	Three rhesus macaque were used (age 8 for one subject, 10 for two subjects at the time of project).
Wild animals	No wild animals were used in this study.
Field-collected samples	There was no field-collected samples.
Ethics oversight	Ethical oversight was taken by IACUC of University of Rochester and University of Minnesota.

Note that full information on the approval of the study protocol must also be provided in the manuscript.