

Internal models of sensorimotor integration regulate cortical dynamics

Seth W. Egger¹, Evan D. Remington^{1,2}, Chia-Jung Chang² and Mehrdad Jazayeri^{1,2*}

Sensorimotor control during overt movements is characterized in terms of three building blocks: a controller, a simulator and a state estimator. We asked whether the same framework could explain the control of internal states in the absence of movements. Recently, it was shown that the brain controls the timing of future movements by adjusting an internal speed command. We trained monkeys in a novel task in which the speed command had to be dynamically controlled based on the timing of a sequence of flashes. Recordings from the frontal cortex provided evidence that the brain updates the internal speed command after each flash based on the error between the timing of the flash and the anticipated timing of the flash derived from a simulated motor plan. These findings suggest that cognitive control of internal states may be understood in terms of the same computational principles as motor control.

Theories of embodied cognition posit an intimate link between mental states and how we move our body. Compelling theoretical arguments have advanced the hypothesis that movements are controlled through the interaction of three building blocks^{1–5} (Fig. 1a): a controller that drives the body, a simulator that simulates movements to predict future states and an estimator that updates state variables. Considering the common challenges that the dynamic control of motor and cognitive states face⁶, we set out to test whether the regulation of cortical dynamics supporting anticipation and planning can be similarly understood in terms of the interplay between a controller, a simulator and a state estimator.

To address this problem, we leveraged a recent finding in monkeys regarding how the brain controls movement-initiation time. When animals are instructed to initiate a movement after a delay, neural responses in multiple brain areas evolve toward a movement-initiation state at a speed that is inversely proportional to the delay^{7,8,9}; that is, faster for shorter intervals and slower for longer intervals. Moreover, similar to how movements are controlled by motor commands¹⁰, the speed with which neural responses evolve over time are controlled by an internally generated “speed command”⁸.

From the perspective of control, initiating a movement based on an instructed delay and in the absence of feedback is relatively straightforward; it only requires an open-loop system driven by the desired speed command. The natural next question to ask is how such speed command may be controlled in a closed-loop fashion in the presence of sensory feedback. To address this question, we developed a novel timing task in which the speed command controlling movement-initiation time had to be adjusted based on uncertain and discrete sensory cues. As noted in the literature for motor control, to integrate uncertain and delayed sensory feedback, the system would benefit from establishing a simulator and an estimator (Fig. 1a). Accordingly, we asked whether closed-loop control of an internally generated speed command could be explained in terms of augmenting the open-loop control with a simulator and an estimator.

Results

1-2-3-Go task. In the 1-2-3-Go task (Fig. 1b), animals had to attend three flashes presented around the fixation point (S1, S2 and S3), and

subsequently initiate a delayed saccade to a target (Go). For every trial, the inter-flash interval, which we refer to as the sample interval (t_s), was drawn from a fixed prior distribution (Fig. 1c). Animals had to estimate t_s from measurements of the S1-S2 and S2-S3 intervals and produce an S3-Go interval (t_p) that closely matched t_s . Animals received a reward when the relative error ($|(t_p - t_s)/t_s|$) was below an adaptively controlled threshold (see Methods). The magnitude of the reward decreased linearly with the relative error (Fig. 1d).

Animals integrate prior knowledge with multiple measurements.

Animals learned to time their saccade based on t_s (Fig. 1e). Using linear regression ($t_p = bt_s + c$), we verified that for both animals, t_p increased with t_s (monkey B: $b = 0.82 \pm 0.01$ (95% confidence interval (CI)), $t(5,806) = 8.51 \times 10^3$, $P < 0.001$; monkey G: $b = 0.81 \pm 0.02$ (95% CI), $t(2,293) = 3.84 \times 10^3$, $P < 0.001$; two-tailed t -test). The variability of t_p also increased with t_s , as evidenced by the regression slope relating the standard deviation of t_p to t_s (100 resamples of the data; monkey B: $b = 0.04 \pm 0.002$ (95% CI), $t(498) = 782.52$, $P < 0.001$; monkey G: $b = 0.11 \pm 0.003$ (95% CI), $t(498) = 1.84 \times 10^3$, $P < 0.001$). This increase in variability is consistent with scalar variability¹¹ during time-interval reproduction tasks^{12,13}.

In the presence of behavioral variability, an ideal observer would reduce variability by taking into account the prior statistics of t_s (refs. 7,14,9). This was evident in the behavior of the animals from systematic biases of t_p (“Bias”, see Methods) toward the prior mean (monkey B: Bias = 30.35 ± 1.27 ms (mean \pm s.d.), $t(99) = 23.97$, $P < 0.001$; monkey G: Bias = 29.89 ± 1.54 ms (mean \pm s.d.), $t(99) = 19.37$, $P < 0.001$; two-tailed t -test), and by the less-than-unity regression slope relating t_p to t_s (monkey B: $t(5,806) = -1,877$, $P < 0.001$; monkey G: $t(2,993) = -987.8966$, $P < 0.001$; two-tailed t -test).

An ideal observer would additionally improve accuracy by combining information from measurements of both the S1-S2 and S2-S3 intervals¹². This was also evident in the behavior of the animals based on a subset of ‘conflict trials’ for the $t_s = 800$ ms condition, in which either the S1-S2 interval, t_{S1-S2} , or S2-S3 interval, t_{S2-S3} , was made 50-ms longer or shorter. We analyzed the conflict trials to distinguish between three hypotheses: H_1 , monkeys only

¹McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. *e-mail: mjaz@mit.edu

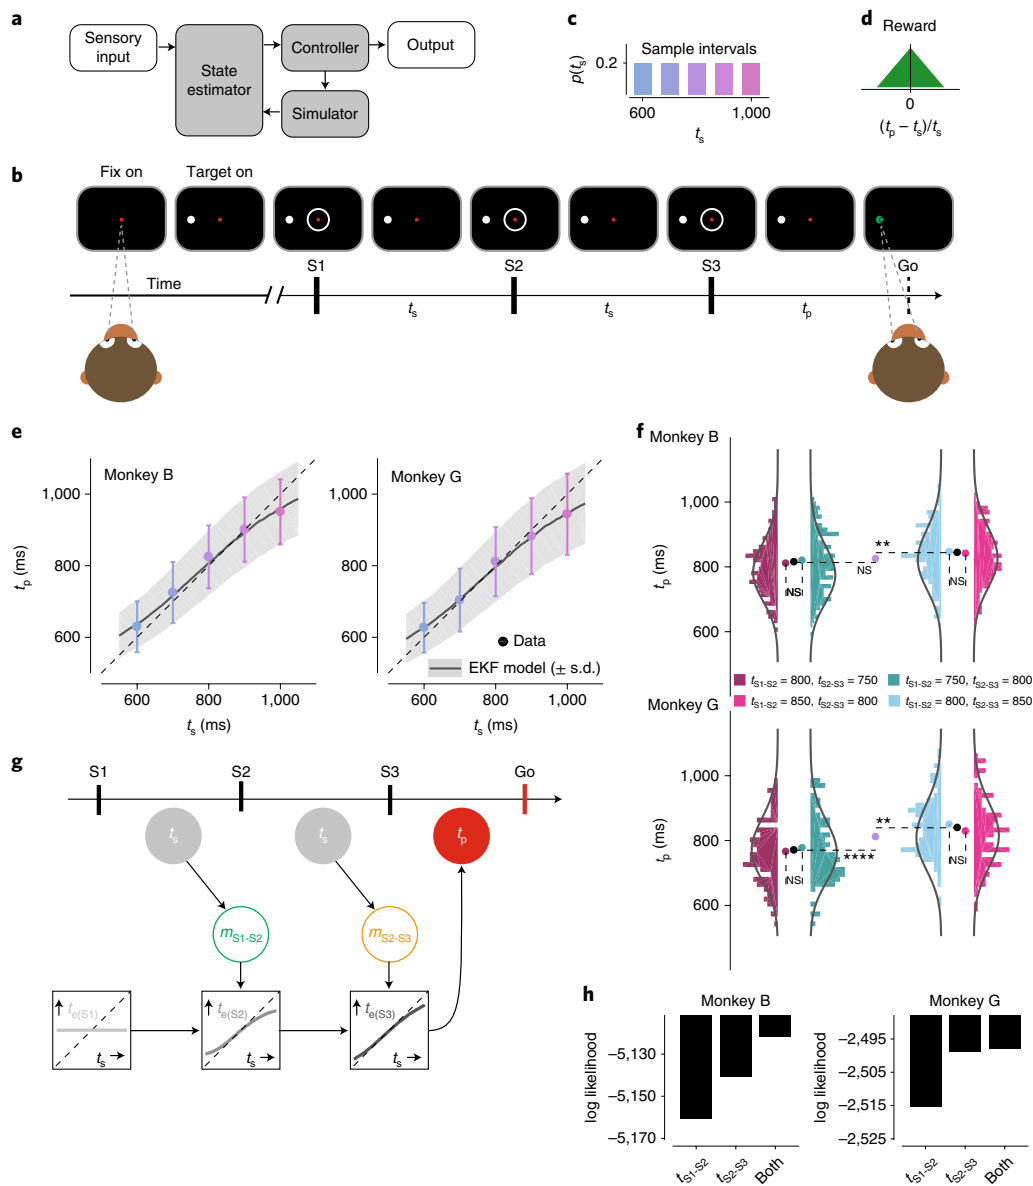


Fig. 1 | The 1-2-3-Go task, behavior and a sequential updating model. **a**, Internal model-based control. The controller drives the output. In the presence of delays and variability, the output control would benefit from a simulator and a state estimator. The simulator predicts the output and the estimator integrates the prediction with sensory inputs to update the controller. **b**, Sequence of events during a trial of 1-2-3-Go. The monkey fixates a central spot (Fix on). After the presentation of a saccadic target (Target on), three isochronous flashed annuli (S1, S2 and S3) are presented around the fixation point. The animal measures the sample interval (t_s) between consecutive flashes and aims to produce a matching interval (t_p) after S3 (Go). **c**, Sample interval distribution, $p(t_s)$. Across trials, t_s was randomly drawn from a discrete uniform prior distribution with five values ranging between 600 and 1,000 ms. **d**, Reward schedule. Maximum reward was delivered for $t_p = t_s$. The reward amount decreased linearly to zero with increasing relative error ($|(t_p - t_s)/t_s|$). **e**, Produced interval (t_p) as a function of sample interval (t_s); note that t_p increased monotonically with t_s (mean, colored circles; standard deviation, error bars; monkey B: $n = 1,412, 1,326, 407, 1,336$ and $1,326$ total trials for $t_s = 600, 700, 800, 900$ and $1,000$ ms, respectively; monkey G: $n = 699, 724, 243, 685$ and 643 trials for $t_s = 600, 700, 800, 900$ and $1,000$ ms, respectively). Responses were biased toward the median t_s and away from the unity line (broken line). Black traces and gray shadings represent the mean and standard deviation, respectively, predicted by the EKF model fit to the behavior. **f**, Analysis of behavior under the four cue conflict conditions. The histograms show the distribution of t_p in different cue conflict conditions (colors) for the two animals. Monkey B: $n = 215, 216, 224$ and 206 trials (left to right). Monkey G: $n = 117, 83, 111$ and 104 trials (left to right). Unbroken lines represent the predicted distribution under the EKF model. Pairs of nearby histograms (for example, $t_{s1-s2} = 800$ ms, $t_{s2-s3} = 750$ ms versus $t_{s1-s2} = 750$ ms, $t_{s2-s3} = 800$ ms) correspond to conflict conditions with the same mean t_s . Colored circles show the mean t_p for different cue conflict conditions. Black circles show the mean t_p across conflict conditions with the same mean t_s . Purple circles correspond to the mean t_p for $t_{s1-s2} = t_{s2-s3} = 800$ ms trials. Broken lines were added to aid comparison between mean values (NS, not significant; ** $P < 0.01$; **** $P < 0.0001$; two-sided t -test). **g**, The EKF model. At the time of S1, the model uses the mean of the prior as its initial estimate of t_s ($t_{e(S1)}$). At S2, the model derives an updated estimate ($t_{e(S2)}$) by applying a nonlinear function to the difference between $t_{e(S1)}$ and the current measurement, denoted m_{s1-s2} . At S3, the model further updates the estimate ($t_{e(S3)}$) by applying the same nonlinearity to the difference between $t_{e(S2)}$ and the second measurement, denoted m_{s2-s3} . The model uses $t_{e(S3)}$ as its final estimate, and produces t_p , which is corrupted by production noise (see Methods and Supplementary Fig. 1). Open and filled circles correspond to unobserved and observable variables, respectively. **h**, The log-likelihood of different variants of the EKF model that use t_{s1-s2} or t_{s2-s3} or both. Larger values indicate more support for a given model. $n = 861$ and 415 total trials for monkeys B and G, respectively. See also Supplementary Fig. 1.

relied on t_{S1-S2} ; H_2 , monkeys only relied on t_{S2-S3} ; or $H_{1,2}$, monkeys used both intervals. When either t_{S1-S2} or t_{S2-S3} was 850 ms, the average t_p was significantly longer than when both intervals were 800 ms (Fig. 1f, light blue for t_{S2-S3} ; one-tailed t -test; monkey B: $t(629) = 2.95$, $P = 0.002$; monkey G: $t(352) = 3.51$, $P < 0.001$; pink for t_{S1-S2} ; one-tailed t -test; monkey B: $t(611) = 2.17$, $P = 0.015$; monkey G: $t(345) = 1.51$, $P = 0.066$). This result indicated that animals used both t_{S1-S2} and t_{S2-S3} and so both H_1 and H_2 were rejected. Moreover, the change in average t_p was statistically indistinguishable between the two conflict conditions (two-tailed t -test; monkey B: $t(428) = -0.62$, $P = 0.531$; monkey G: $t(213) = -1.54$, $P = 0.125$), which suggested that animals combined the two measurements with comparable weights. Together, these results confirm that animals combined the two measurements.

A sequential updating model for the 1-2-3-Go task. Recently, we found that human performance on this task was better captured by a near-optimal sequential updating model than by an optimal Bayesian model¹². This near-optimal sequential updating model, which we refer to as the extended Kalman filter (EKF), functions as follows (Fig. 1g; Supplementary Fig. 1c). (1) At S1, the estimated interval, $t_{e(S1)}$, is set to the mean of the prior. (2) At S2, the model integrates the first measurement, denoted m_{S1-S2} , with the prior to compute an updated estimate, $t_{e(S2)}$. To do so, EKF updates $t_{e(S1)}$ by a nonlinear function of error between $t_{e(S1)}$ and m_{S1-S2} (see Methods). (3) At S3, the model repeats this process to integrate the second measurement, m_{S2-S3} , with the prior and m_{S1-S2} and computes the final estimate, $t_{e(S3)}$. After S3, the model produces t_p , which is $t_{e(S3)}$ plus signal-dependent production noise.

We found that the pattern of responses in monkeys was also more accurately captured by the EKF model (Fig. 1e, curves) compared with an optimal Bayesian model (Supplementary Fig. 1). We also used EKF to substantiate that animals integrated both measurements. We compared the log-likelihood of the data given three variants of the EKF based on H_1 , H_2 and $H_{1,2}$. The EKF that relied on both measurements was more likely than the other two variants (Fig. 1h). Furthermore, EKF fitted to trials without conflict was able to capture the distribution of t_p in conflict trials (Fig. 1f, curves). Based on these results, we concluded that monkeys integrated the prior with the two measurements using a sequential updating strategy exemplified by EKF.

Speed control in the 1-2-3-Go task. Previous work in monkeys has shown that the brain controls movement-initiation time by setting a speed command that determines how fast neural trajectories in the cortico-basal ganglia system evolve toward a movement-initiation state (that is, “threshold”⁸). Building on this finding, we considered two hypotheses regarding how the brain might control movement-initiation time in the 1-2-3-Go task. One hypothesis, which we refer to as the ‘open-loop hypothesis’, is that the brain integrates the prior with the two measurements (m_{S1-S2} and m_{S2-S3}) to compute the final estimate, $t_{e(S3)}$, which in turn determines the speed of neural trajectories in the S3-Go epoch. A key prediction of this hypothesis is that neural activity in the S1-S2 and S2-S3 epochs should be similar and reflect explicit measurements of time intervals. Moreover, activity patterns in the first two epochs should be qualitatively different from that in the third epoch, in which the responses would evolve with different speeds depending on the animal’s final estimate. The predictions of this hypothesis can be illustrated using a toy model. Let us assume that the brain tracks time by accumulating ticks from a clock. Since the first two epochs are associated with measurement of the interval, we would expect some neurons to carry signals reflecting a fixed clock and others to exhibit ramping activity with a fixed slope (Fig. 2a, left and middle). In the third epoch, however, signals would change qualitatively such that the clock speed and the slope would correspond to the animal’s estimate (Fig. 2a, right).

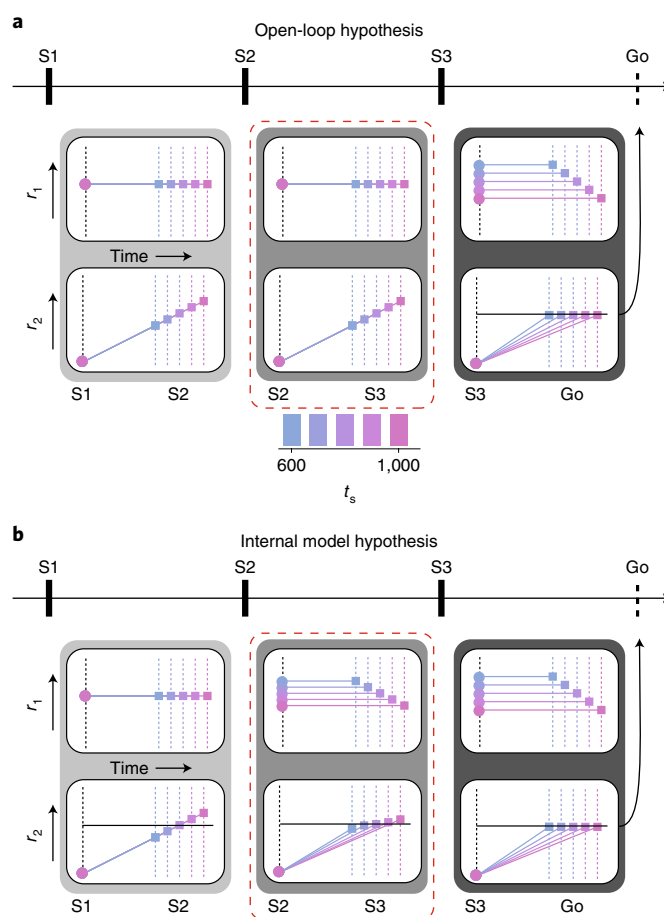


Fig. 2 | Predictions of the open-loop and internal model hypothesis in the 1-2-3-Go task. **a**, Activity profile of two hypothetical populations, r_1 and r_2 , under the open-loop hypothesis. Predictions are shown for each of the three epochs. r_1 provides the clock and r_2 integrates the clock to track elapsed time. During both the first and second epochs (light-gray and medium-gray boxes), the clock signal is independent of t_s . Therefore, r_2 at S2 or S3 represents the measured interval (different colored squares). These measurements set the clock speed (output of r_1) according to $t_{e(S3)}$, and, as a consequence, r_2 ramps to the threshold (horizontal black line) at a rate that depends on t_s from S3 to Go (dark-gray box). **b**, Activity profile under the internal model hypothesis. r_1 and r_2 represent the control signal and simulation, respectively. During the first epoch (light-gray box), the r_1 population supplies a constant speed command according to the mean of the prior. The r_2 population integrates this speed command and produces ramping activity that terminates at different points depending on the sample interval, t_s (different colored squares). The difference between the terminal point of r_2 and the predicted terminal point based on the mean of the prior (horizontal line) provides an error that allows r_1 to encode an interval-dependent speed command in the subsequent epoch (medium gray box, top). The corresponding r_2 population integrates the updated speed to generate ramping activity whose slope varies with t_s (bottom). The same process is repeated in the third epoch (dark-gray box). The broken red line highlights the key differences between the two hypotheses.

The other hypothesis, which we refer to as the ‘internal model hypothesis’, is that the brain implements a closed-loop control system similar to the EKF model, in which feedback is used to update estimates sequentially from $t_{e(S1)}$ at S1 to $t_{e(S2)}$ at S2 to $t_{e(S3)}$ at S3. To implement this strategy, the brain has to additionally establish a simulator and an estimator. The simulator would run after each flash to predict the next flash, and the estimator would

update the speed based on the prediction error at the time of the next flash.

We used the inverse relationship between interval estimate and speed of neural trajectories to formulate the following set of concrete predictions under the internal model hypothesis. (1) After S1 and before S2, the speed of neural trajectories should reflect the initial estimate, $t_{e(S1)}$. (2) After S2 and before S3, the speed command and the speed with which responses evolve should be inversely proportional to $t_{e(S2)}$. (3) After S3, the speed should be adjusted to reflect $t_{e(S3)}$. The predictions of this hypothesis can also be demonstrated using the clock-accumulator toy model (Fig. 2b). In the first epoch, the clock speed and the slope of the ramping activity would be fixed and reflect $t_{e(S1)}$ (Fig. 2b, left). In contrast, the second and third epoch would represent a simulated motor plan (S2-S3) and an actual motor plan (S3-Go). Accordingly, both the level of activity of neurons that encode the clock speed and the slope of the neurons generating ramping activity would be interval-dependent (Fig. 2b, middle and right). Specifically, in S2-S3 and S3-Go, the speed and the slope should be inversely proportional to $t_{e(S2)}$ and $t_{e(S3)}$, respectively.

Single-neuron signatures of the internal model hypothesis. We recorded neural activity ($N=63$ and 118 in monkeys B and G, respectively) in the dorsomedial frontal cortex (DMFC) (Fig. 3a), including the supplementary eye field (SEF), the dorsal region of the supplementary motor area (SMA) and the pre-SMA. Our choice of recording areas was motivated by previous work showing a central role for the DMFC in sensorimotor timing^{15–18}. We targeted the same region of DMFC where we had previously observed the inverse relationship between time and speed^{7,8,9}. DMFC responses were modulated throughout the trial and had heterogeneous profiles that were typically distinct across the three epochs (Fig. 3b). Our main interest was in evaluating the neural data in terms of the open-loop and internal model hypotheses. To do so, we analyzed the activity of 115 well-isolated neurons with respect to the aforementioned predictions of those hypotheses (Fig. 2b). Note that the key difference between the open-loop and internal model hypotheses is the nature of signals in the S2-S3 epoch. According to the open-loop hypothesis, responses in this epoch reflect an interval measurement regardless of the duration of the interval in the S1-S2 epoch. In contrast, the internal model hypothesis predicts that responses in the S2-S3 epoch evolve at different speeds depending on the animal's estimate after S2 ($t_{e(S2)}$). The predictions of the two hypotheses for the other two epochs (S1-S2 and S3-Go) are identical.

Temporal scaling across individual neurons. Temporal scaling predicts that response profiles across t_s would become more similar after undoing the temporal scaling. This was most conspicuous for the small subset of neurons with a ramp-like response profile for which the slope of the ramp had an inverse relationship to t_s (Fig. 4a, top right). However, consistent with previous findings^{8,19,20}, a majority of neurons had non-monotonic firing rate profiles, and were better fit by higher-order polynomials (see Methods; Supplementary Fig. 2). For example, cross-validated sixth-order polynomials explained 40.86, 24.42 and 20.34% more variance than cross-validated first-order polynomials in the S1-S2, S2-S3 and S3-Go epochs, respectively (Fig. 4b).

To assess the presence of temporal scaling for every neuron, we normalized response profiles in time according to the animal's estimates (that is, $t_{e(S2)}$ and $t_{e(S3)}$ for the S2-S3 and S3-Go epochs, respectively; Fig. 4c), and asked whether a polynomial fit to the normalized response profiles would explain more variance (R^2) than a polynomial of the same order fitted to original data (see Methods). We used the difference between the explanatory powers of the two polynomial fits (ΔR^2) as our metric of the degree of scaling. We also computed the same metric in the S1-S2 epoch as a

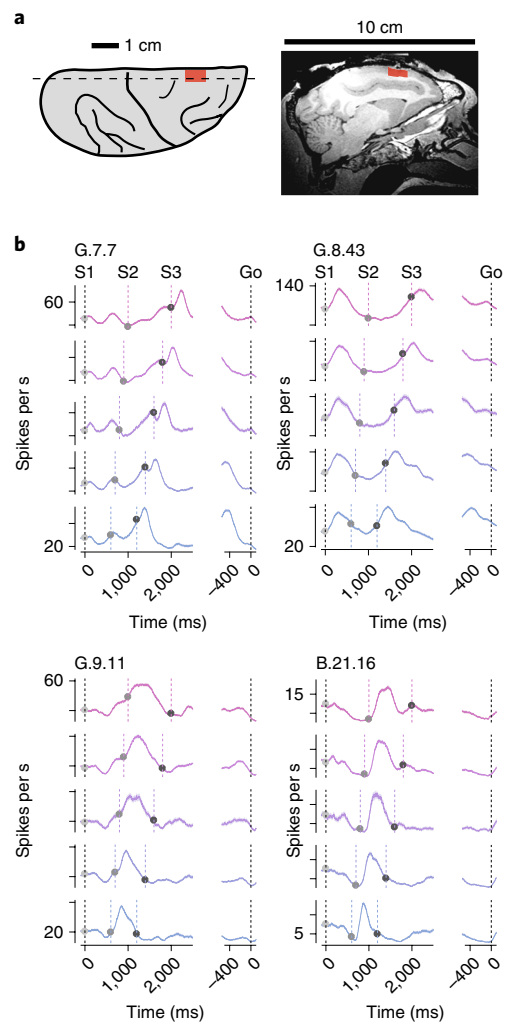


Fig. 3 | Example response profiles of individual DMFC neurons during the 1-2-3-Go task. a, Recording location. Left: illustration of the dorsal aspect of the right hemisphere, showing major sulci (black lines) and the location of recordings across monkeys (red). Right: magnetic resonance image showing a sagittal section (broken line, left) of the right hemisphere in monkey G; mean across two image acquisitions, with similar results obtained. The red region indicates the rostrocaudal extent of recordings. **b**, Average response profiles for four example neurons aligned to S1 and Go sorted according to t_s (colors). Shading indicates standard errors. Vertical broken lines indicate flash and saccade times. Trial numbers for $t_s = 600, 700, 800, 900$ and $1,000$ ms are as follows: G.7.7: 185, 206, 72, 193 and 148, respectively; G.8.43: 156, 169, 51, 138 and 133, respectively; G.9.11: 178, 172, 49, 157 and 151, respectively; B.21.16: 210, 168, 47, 176 and 173, respectively.

benchmark, since we were certain that there was no temporal scaling in this epoch (the animal did not yet know t_s). Here, we show the result for a sixth-order polynomial, which was sufficiently complex to describe the temporal responses across the DMFC population (Supplementary Fig. 2).

The degree of scaling in the S3-Go epoch was consistently stronger than the S1-S2 epoch (Fig. 4d, bottom; one-tailed Wilcoxon, $Z=7.08$, $P<0.001$). Importantly, the same was true in the S2-S3 epoch (Fig. 4d, top; one-tailed Wilcoxon; $Z=4.62$; $P<0.001$), and the results remained unchanged independent of the degree of the polynomial used to assess scaling. In fact, scaling was evident even when the polynomial order was chosen to maximize R^2 for unscaled data (Supplementary Fig. 3). Therefore, we concluded

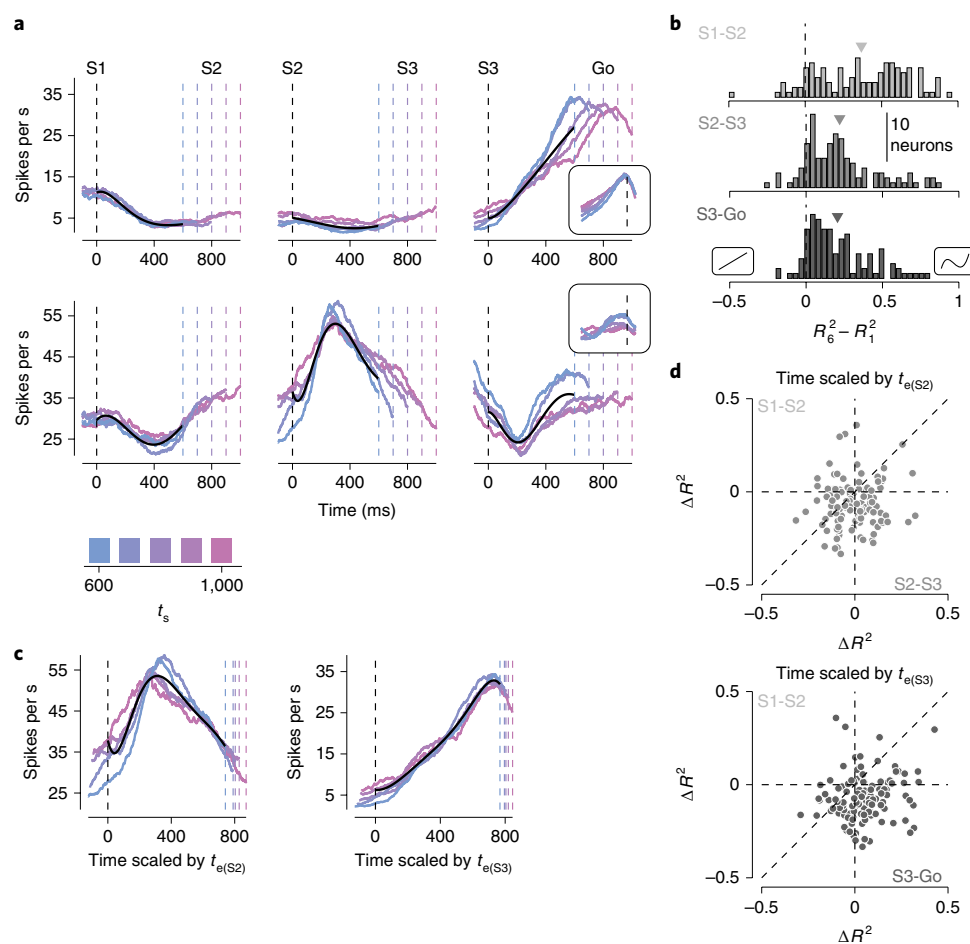


Fig. 4 | Temporal scaling of non-monotonic firing rates across individual DMFC neurons. **a**, Firing rates of two example neurons (rows) as a function of task epoch (columns) for different values of t_s (colors). Black lines represent the fit of a sixth-order polynomial to the first 600 ms of each epoch and neuron. From left to right, firing rates were aligned to S1, S2 and S3 flashes, respectively (black broken line). Colored broken lines represent the time of the subsequent flash. Insets show the firing rates aligned to Go (black broken line). **b**, Histograms of the difference in the fraction of variance explained, R^2 , between the sixth-order and first-order polynomial fit to the data. Negative and positive values correspond to neurons that favored the first-order and sixth-order polynomial, respectively. Insets on the negative and positive sides of the abscissa illustrate example first-order and sixth-order polynomials. Arrowheads indicate the mean of the distributions. Different gray levels correspond to different task epochs. **c**, Example scaled temporal responses. Left: firing rate of the second neuron in **a** during the S2-S3 epoch, temporally scaled by the animals' estimate in the S2-S3 epoch. The estimate was inferred from the model fits to the behavior. Right: firing rate profile for the S3-Go epoch of the first neuron in **a** temporally scaled by the animals' estimate in the S3-Go epoch. Black lines show the sixth-order polynomial that best fit the scaled data. Broken lines are as in **a**, but appropriately scaled. **d**, Scatter plot showing the change of explanatory power (ΔR^2) of the sixth-order polynomial fit to the scaled versus unscaled data across neurons. Top: ΔR^2 in the S2-S3 epoch compared with the S1-S2 epoch. Bottom: ΔR^2 in the S3-Go epoch compared with the S1-S2 epoch. Vertical and horizontal broken lines indicate $\Delta R^2 = 0$, and the diagonal in the unity line. See also Supplementary Figs. 2 and 3.

that responses in both S2-S3 and S3-Go were temporally stretched in accordance with t_s . These observations provide evidence that signals in the S3-Go epoch are governed by a speed-dependent motor plan, which is consistent with previous work on motor timing^{8,21}, and that responses in the S2-S3 epoch exhibit temporal scaling despite the fact that they were not associated with a motor plan. The latter observation suggests that activity in the S2-S3 epoch is predictive of the upcoming flash, and therefore provides evidence in support of the internal model hypothesis and against the open-loop hypothesis.

Speed command across individual neurons. A neuron could be a candidate for providing the hypothesized speed command if its response in the S2-S3 and S3-Go epochs were to vary systematically with t_s (Fig. 2b). In our population, the firing rate of many neurons was modulated by t_s in either the S2-S3 or S3-Go epochs (Fig. 5a). To rigorously quantify this observation, we fit a linear-nonlinear

Poisson (LNP) model to the spike count data (Fig. 5b, lines; see Methods). This yielded a parameter, β_{SX} (where $X = 1, 2$ or 3) that we used to quantify the degree to which the firing rate depended on t_s at the time of each flash. At S1, neurons cannot logically encode t_s since the animal has not yet measured t_s . Accordingly, β_{S1} across the population was distributed tightly around zero (Fig. 5c, left) and was not significant for any of the neurons. β_{S2} and β_{S3} were distributed more broadly and were significant for 27 and 20 of the neurons, respectively (Fig. 5c, middle and right, white bars; see Methods). However, the degree of sensitivity to t_s at S2 and S3 were unrelated across the population (Fig. 5d; $R^2 = 0.0013$, $P = 0.70$).

We next asked whether single neurons had an invariant representation of t_s throughout each epoch. Invariant encoding of t_s would predict high and persistent correlation between spike counts at pairs of time points within each epoch. To test this prediction, we quantified spike count correlations for each neuron as a function of t_s across all pairs of time points within the first 600 ms after the onset of each

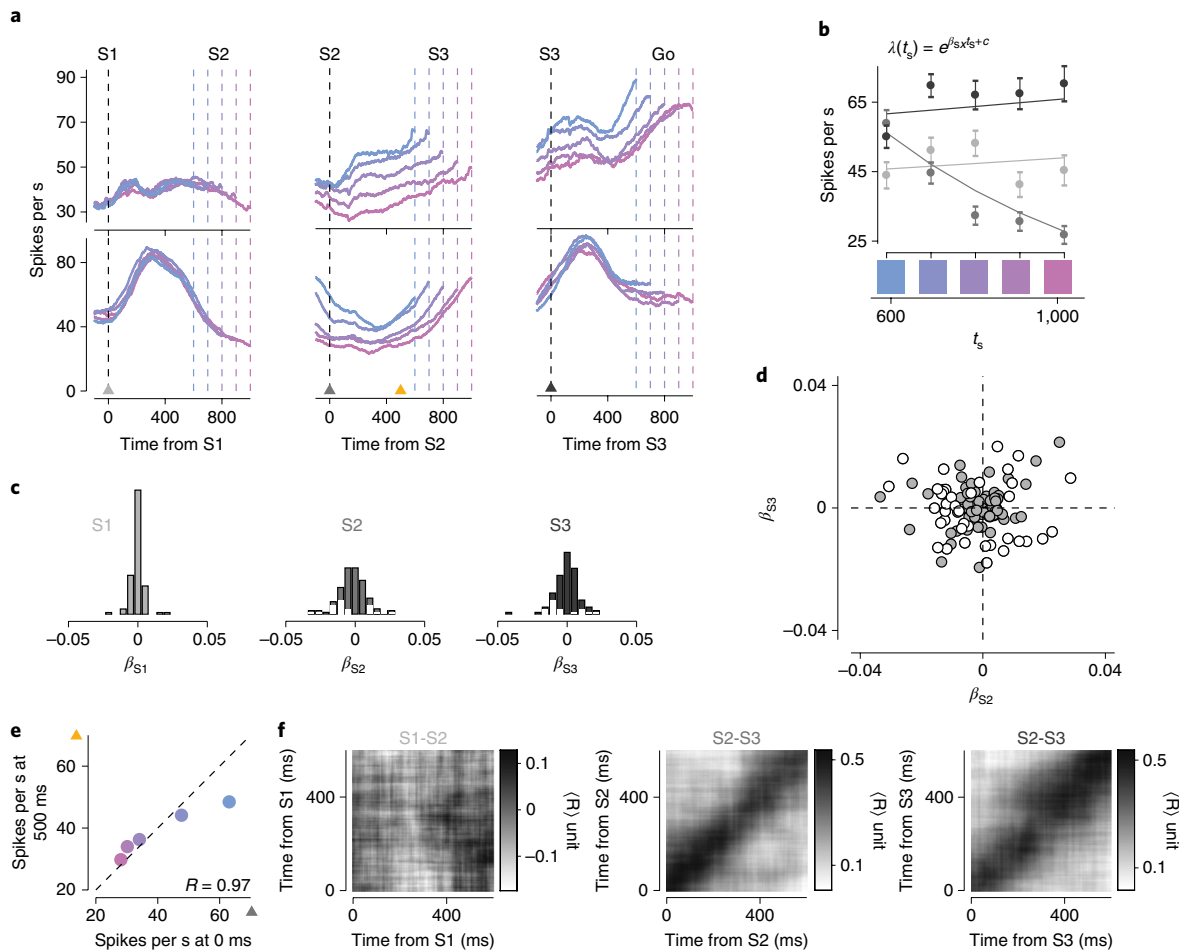


Fig. 5 | A representation of the sample interval by individual DMFC neurons. a, Firing rates of two example neurons with the same format as in Fig. 4a. Triangles indicate time points of analysis in **b** and **e**. **b**, The relationship between t_s and firing rates \pm standard error (circles and error bars; $n=30$ trials for each t_s , sampled with replacement from the test dataset) at S1, S2 and S3 for the example neuron in **a** (bottom row). Firing rates were computed from spike counts within a 150-ms window centered on S1 (light gray), S2 (medium gray) and S3 (dark gray). Lines are fits of the LNP model (functional form shown at the top). **c**, Histograms showing the sensitivity of individual neurons at the time of each flash to t_s . Sensitivity values (β_{S1} , β_{S2} and β_{S3}) were derived from the LNP model fits as shown for an example neuron in **b**. White bars include neurons for which LNP provided a superior fit to a constant firing rate model (see Methods). **d**, Scatter plot of sensitivity to t_s at S3 versus S2 across the population based on β_{S2} and β_{S3} parameters. White circles represent neurons that were significantly tuned to t_s at either S2 or S3. **e**, Pearson's correlation coefficient between firing rates, conditioned on t_s , at the time of S2 (x axis) and 500 ms afterwards (y axis) for the example neuron in **a** (middle column, bottom row). Mean values for the 0 ms time point were found from a random selection of half the total trials ($n=156, 169, 51, 138$ and 133 for $t_s=600, 700, 800, 900$ and $1,000$ ms, respectively); mean values for the 500 ms time point were found from the remaining trials. **f**, Average correlation between firing rates, conditioned on t_s , computed for different pairs of time points. The average was taken across all individual neurons ($n=115$). Left, middle and right panels correspond to the analysis performed at time points relative to S1, S2 and S3, respectively. See also Supplementary Figs. 4–7.

epoch (for example, Fig. 5e), and constructed a two-dimensional map of average correlations across t_s . In both the S2–S3 and S3–Go epochs, average correlations exhibited a diagonal structure indicating that correlations were high for nearby time points and weak for points farther apart (Fig. 5f). These results provide evidence against an invariant representation of t_s in the S2–S3 or S3–Go epochs at the level of individual neurons. This conclusion was further supported by an analysis using the LNP model (Supplementary Fig. 4). Note that although single neurons did not invariantly encode t_s , we were able to readily decode t_s from the population activity over the entire S2–S3 and S3–Go epochs (Supplementary Figs. 5–7).

Population response in the DMFC exhibits hallmarks of the internal model hypothesis. Single-neuron analyses were consistent with some of the predictions of the internal model hypothesis. However, four considerations motivated us to further analyze

the responses at the population level. First, single neurons that exhibited temporal scaling had non-monotonic response profiles (Fig. 4b), and thus could not unambiguously predict the upcoming event. We reasoned that this ambiguity might be resolved at the population level. Second, single neurons whose firing rate encoded t_s and thus could serve as a control signal did so transiently (Fig. 5f). However, at the population level, responses carried a persistent representation of t_s that was present for the entire S2–S3 or S3–Go epochs (Supplementary Figs. 4–7). Third, since our experiment did not include any dynamic stimuli or ongoing movement between consecutive flashes, the observed patterns of activity likely emerged from recurrent interactions; therefore, these patterns may be more interpretable if analyzed at the population level. Fourth, we wanted to exploit the additional statistical power afforded by the population analysis to ask whether the speed command and the simulated motor plan reflected t_s or the animal's internal estimate of t_s ($t_{e(SX)}$).

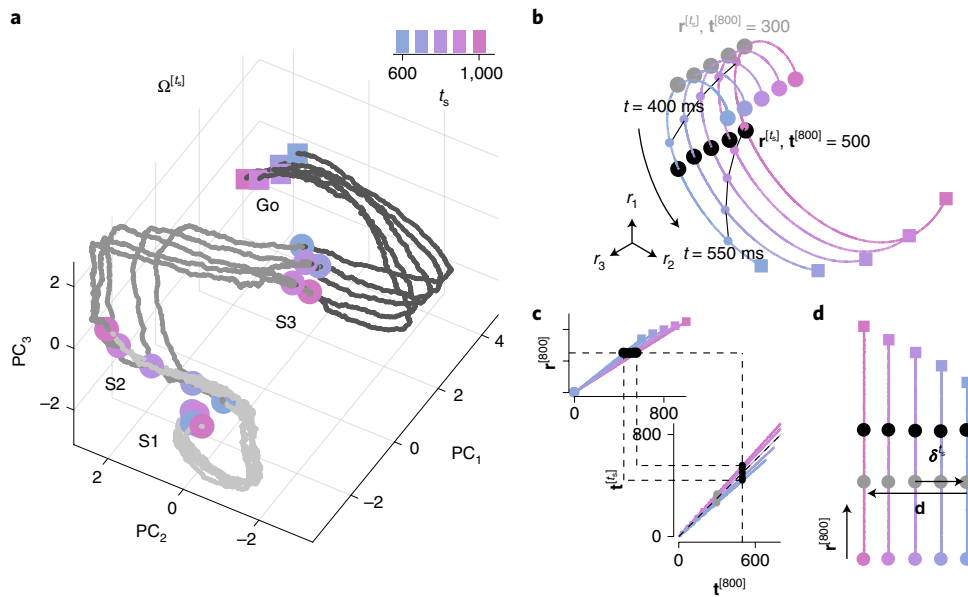


Fig. 6 | Neural trajectories and a technique for analyzing their kinematics. **a**, Neural trajectories, $\Omega^{[t_s]}$, plotted in a subspace spanned by the first three PCs computed from patterns of activity of $n=181$ units across time, from S1 to Go. Circles indicate states at S1, S2 or S3 and squares denote states 100 ms before the correct t_p for each t_s . **b-d**, Schematic illustration of KiNeT. **b**, Simulated neural trajectories generated under the internal model hypothesis and embedded in three-dimensional space by a nonlinear transformation. Gray and black markers correspond to the set of states, $\mathbf{r}^{[t_s]}$, nearest to $\mathbf{r}^{[800]}$ at $\mathbf{t}^{[800]}=300$ ms and $\mathbf{t}^{[800]}=500$ ms, respectively. Small circles connected by thin black lines are states on the trajectories at 400 ms or 550 ms after the starting point. **c**, The time required for $\Omega^{[t_s]}$ to reach $\mathbf{r}^{[t_s]}$, denoted $\mathbf{t}^{[t_s]}$, depends on the speed along each trajectory (top). Using this relationship, we can map $\mathbf{t}^{[t_s]}$ to $\Omega^{[t_s]}$ for each trajectory (bottom). Black points illustrate $\mathbf{t}^{[t_s]}$ for $\mathbf{t}^{[800]}=500$ ms. The application of KiNeT on $\Omega^{[t_s]}$ with different speeds is illustrated by the broken lines. Gray points illustrate $\mathbf{t}^{[t_s]}$ for $\mathbf{t}^{[800]}=300$ ms. **d**, Use of KiNeT to infer the relative positions of trajectories. Colored lines represent the trajectories corresponding to different values of t_s unraveled into two dimensions using KiNeT (**b** and **c**). KiNeT provides an estimate of the distance between $\Omega^{[t_s]}$ and $\Omega^{[800]}$ at different time points. δ^{t_s} is a vector that connects corresponding states on $\Omega^{[t_s]}$ and $\Omega^{[800]}$. **d** is a vector that connects corresponding states on the first and last trajectories. Distance was measured by the projection of δ^{t_s} on **d**. Black and gray circles correspond to the nearby states, as in **a**. See also Supplementary Figs. 8 and 9.

As a first step, we performed principal component analysis (PCA) to identify the dominant modes of response from S1 to Go. Projections of neural activity onto the first three principal components (PCs) as a function of t_s revealed a set of highly structured neural trajectories with two notable features. First, neural trajectories had a persistent representation of t_s after S2 (Fig. 6a). This suggests that, unlike single neurons, the population activity in the DMFC carries a t_s -dependent signal that could serve as a speed command. Second, neural trajectories were highly similar in form and terminated in nearby states at the time of S3 and Go irrespective of the duration of t_s . This observation implied that the speed at which different trajectories evolve had an inverse relationship to t_s , which is what is expected from a predictive process anticipating the next event.

Analysis of relative speed and distance of neural trajectories. When a set of neural trajectories are structured as in our experiment, it is possible to estimate their relative distance and speed using a recently developed analysis method known as KiNeT⁷. We describe KiNeT in detail in the Methods. Here, we use an analogy to describe the key idea. Imagine that we want to measure the speed of a set of runners running in parallel tracks along a winding road. If we draw a line connecting nearby points along the tracks, the time at which each runner reaches that line would be inversely proportional to their speed. We can also estimate the distance between tracks from the relative position of the runners when they reach that line. KiNeT applies the same logic to estimate the relative distances and speeds of a set of winding neural trajectories (see Methods).

Let us denote the neural trajectory associated with a specific value of t_s by $\Omega^{[t_s]}$ (Fig. 6a). We designated $\Omega^{[800]}$ as the reference

trajectory, and defined $\mathbf{t}^{[800]}$ as the vector of time points along $\Omega^{[800]}$ (for example, a vector containing 0, 1, ..., 800). Next, for every neural state along $\Omega^{[800]}$, denoted $\mathbf{r}^{[800]}$, we found the nearest neural state on the remaining $\Omega^{[t_s]}$ trajectories, denoted $\mathbf{r}^{[t_s]}$ (Fig. 6b), and defined $\mathbf{t}^{[t_s]}$ as the vector of time points when activity reached $\mathbf{r}^{[t_s]}$ (Fig. 6c). Finally, we quantified the distance between $\Omega^{[t_s]}$ and $\Omega^{[800]}$, denoted δ^{t_s} , using the vector connecting $\mathbf{r}^{[t_s]}$ and $\mathbf{r}^{[800]}$, and the speed of $\Omega^{[t_s]}$ relative to $\Omega^{[800]}$ ($\mathbf{V}^{[t_s]}$) based on the values of $\mathbf{t}^{[t_s]}$ relative to $\mathbf{t}^{[800]}$ (see Methods).

DMFC population activity is governed by speed-dependent predictive dynamics. We first formulated the predictions of KiNeT for the open-loop and internal model hypotheses. For both hypotheses, the speed in the S1-S2 epoch should not depend on t_s (Fig. 7a, top left). Accordingly, in the first epoch, $\mathbf{t}^{[t_s]}$ should be identical to $\mathbf{t}^{[800]}$ irrespective of t_s . However, the two hypotheses make different predictions for the S2-S3 epoch. According to the open-loop hypothesis, the computation in the S2-S3 epoch is similar to the S1-S2 epoch. Therefore, for the open-loop hypothesis, $\mathbf{t}^{[t_s]}$ should remain identical to $\mathbf{t}^{[800]}$ irrespective of t_s (Fig. 7a, top). In contrast, the internal model hypothesis predicts that activity in the S2-S3 epoch is predictive and, therefore, the speed exhibits an inverse relationship to t_s (Fig. 7a, bottom). Therefore, $\mathbf{t}^{[t_s]}$ should increase faster than $\mathbf{t}^{[800]}$ for slower trajectories and vice versa. In the S3-Go epoch, both hypotheses predict that the speed would depend on t_s ; therefore, the relationship between $\mathbf{t}^{[t_s]}$ and $\mathbf{t}^{[800]}$ should depend on t_s (Fig. 7a, bottom). Consequently, the key prediction that differentiates between the two hypotheses is whether or not activity in the S2-S3 epoch is predictive or, equivalently, whether the relationship between $\mathbf{t}^{[t_s]}$ and $\mathbf{t}^{[800]}$ depends on t_s .

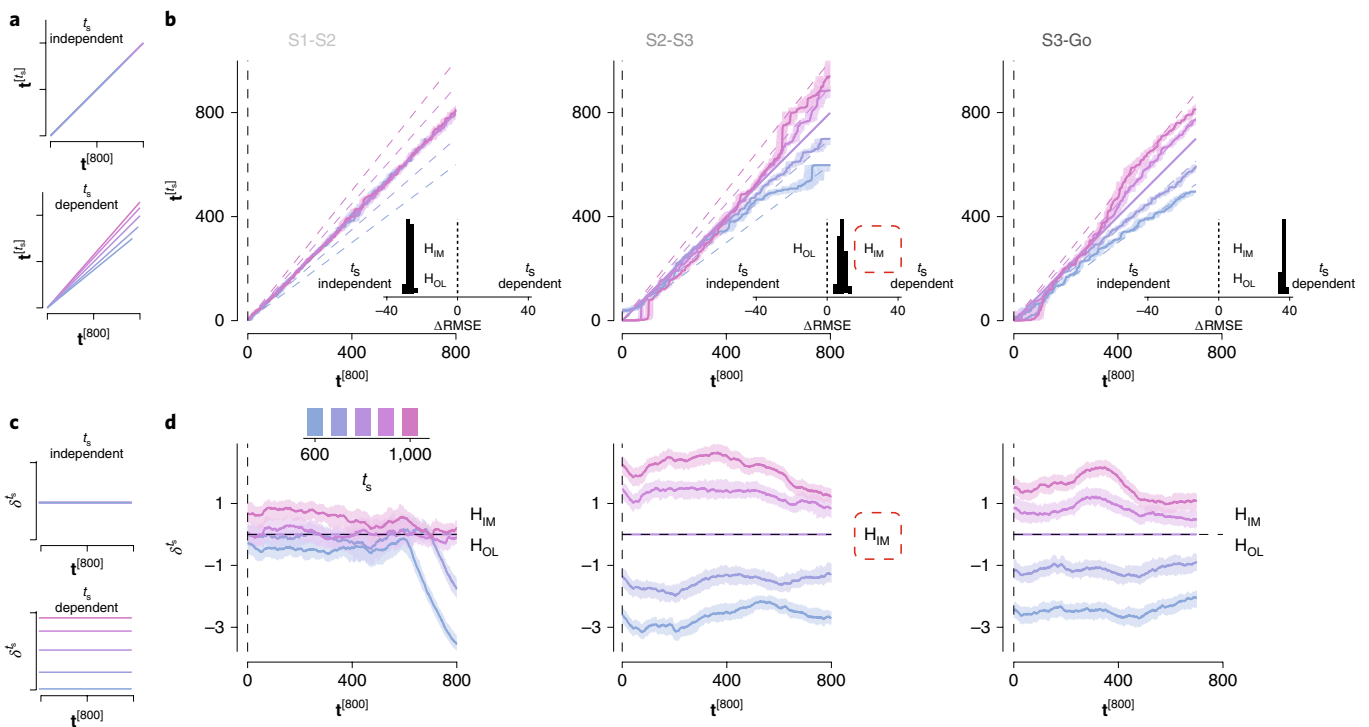


Fig. 7 | Relative speed and distance between neural trajectories during 1-2-3-Go. **a**, t_s -independent and t_s -dependent representations of speed. **b**, Speed of each neural trajectory, $\Omega^{[t_s]}$, compared with the speed of the reference trajectory, $\Omega^{[800]}$. Colored lines show the progression of elapsed time on $\Omega^{[t_s]}$ ($t^{[t_s]}$) as a function of elapsed time on $\Omega^{[800]}$ ($t^{[800]}$) for different t_s values. Shadings indicate the median \pm 95% CIs computed from $n=100$ bootstrap resamples. The unity line corresponds to no difference in speed. The broken lines represent the expected relationship between $t^{[t_s]}$ and $t^{[800]}$ under the internal model hypothesis for an observer with perfect knowledge of t_s . Insets show the difference in root mean square error (RMSE) under the assumption of t_s -independent speed (left of zero) versus t_s -dependent speed (right of zero). In the S1-S2 epoch, both the open-loop and internal hypotheses (H_{OL} and H_{IM} , respectively) predict t_s -independent speed; t_s -dependent predictions were based on fits of EKF to behavior. In the S2-S3 epoch, the H_{OL} predicts t_s -independent speed, and the H_{IM} predicts t_s -dependent speed; the data supports the H_{IM} (red broken line). In the S2-S3 epoch, both hypotheses predict t_s -dependent speed. **c**, t_s -independent and t_s -dependent representations of the speed command. **d**, Distance (δ^s) between nearby states on $\Omega^{[t_s]}$ and $\Omega^{[800]}$ as a function of time. The horizontal broken line ($\delta^s = 0$) corresponds to overlapping trajectories. The data support the H_{IM} (red broken line). Shadings indicate the median \pm 95% CIs computed from $n=100$ bootstrap resamples. See also Supplementary Figs. 10–12.

KiNeT provided clear evidence to support the internal model hypothesis. In the S1-S2 epoch, $\mathbf{V}^{[t_s]}$ did not exhibit any systematic relationship to t_s ; that is, the slope of $t^{[t_s]}$ as a function of $t^{[800]}$ was not significantly different from one (Fig. 7b, left; two-tailed t -test; $t_s=600$: $t(99)=0.1225$, $P=0.9028$; $t_s=700$: $t(99)=-0.3136$, $P=0.7545$; $t_s=900$: $t(99)=-0.5884$, $P=0.5576$; $t_s=1,000$: $t(99)=0.7642$, $P=0.4466$; Fig. 7a, left). Accordingly, $\mathbf{V}^{[t_s]}$ in the S1-S2 epoch was better explained by a t_s -independent model than a t_s -dependent model (Fig. 7b, left inset; one-tailed t -test, $t(99)=-22.4067$, $P<0.001$). In the S2-S3 epoch, values of $t^{[t_s]}$ increased faster than values of $t^{[800]}$ for longer t_s values (Fig. 7b, middle; one-tailed t -test; $t_s=900$: $t(99)=5.3862$, $P<0.001$; $t_s=1,000$: $t(99)=9.1030$, $P<0.001$) and slower for shorter t_s values (one-tailed t -test; $t_s=600$: $t(99)=-14.4431$, $P<0.001$; $t_s=700$: $t(99)=-9.6668$, $P<0.001$; Fig. 7b, middle). This indicates that $\mathbf{V}^{[t_s]}$ in the S2-S3 epoch was ordered according to t_s , which is consistent with the internal model hypothesis (Fig. 7b, middle inset; one-tailed t -test, $t(99)=4.1760$, $P<0.001$). The organization of $t^{[t_s]}$ with respect to $t^{[800]}$ in the S3-Go epoch was also consistent with adjustments of speed depending on t_s (Fig. 7b, right; one-tailed t -test, $t(99)=35.2322$, $P<0.001$). These results are consistent with DMFC responses reflecting a simulated and an actual motor plan in the S2-S3 and S3-Go epochs, respectively.

DMFC population activity reflects an interval-dependent speed command. The open-loop and the internal model hypotheses also

make distinct predictions regarding the representation of a t_s -dependent speed command. The open-loop hypothesis predicts a t_s -independent speed command in the first two epochs (Fig. 7c, top) and a t_s -dependent signal in the third epoch (Fig. 7c, bottom). In contrast, the internal model hypothesis predicts a t_s -dependent signal in both the second and third epochs (Fig. 7c, bottom). Qualitative examination of $\Omega^{[t_s]}$ supported the internal model hypothesis; that is, neural states evolved along parallel trajectories in both the S2-S3 and S3-Go epochs and were systematically organized in neural state space in accordance with t_s (Fig. 6a). To quantitatively examine this organization, we measured the distance (δ^s) between each $\Omega^{[t_s]}$ and $\Omega^{[800]}$ as a function of time after using KiNeT (see Methods; Fig. 6d).

In the first epoch, trajectories were highly overlapping and there was no significant distance between $\Omega^{[700]}$ and $\Omega^{[900]}$ (Fig. 7d, left; one-tailed t -test; $t(99)=0.4796$, $P=0.6325$), the two test trajectories. In contrast, in the S2-S3 (Fig. 7d, middle) and S3-Go epochs (Fig. 7d, right), test trajectories were systematically separated according to t_s (one-tailed t -test; S2-S3: $t(99)=8.2495$, $P<0.001$; S3-Go: $t(99)=6.2927$, $P<0.001$; see Methods). Moreover, the magnitude of δ^s was nearly constant between consecutive flashes, as evidenced by the slope of the regression line relating δ^s to $t^{[800]}$, which was less than 0.1% of the offset in both epochs. These observations indicate that population responses in the DMFC harbor a representation of t_s throughout the S2-S3 and S3-Go epochs, which is consistent with the internal model hypothesis.

We performed a number of additional analyses to verify the robustness of our results. For example, the preceding results were based on the application of KiNeT to the first ten PCs derived from activity within each epoch, which captured more than 90% of variance (Supplementary Figs. 8 and 9). However, the same results were found when we applied KiNeT to PCs of activity across the three epochs (Supplementary Fig. 10). Moreover, all results were consistent across the two animals (Supplementary Fig. 11). We also tested KiNeT on responses aligned to the time of saccade to ensure that results in the S3-Go epoch were not due to an averaging of firing rates across trials of different durations (due to different saccade times; Supplementary Fig. 12). Finally, we note that the systematic modulations of the relative distance and speed with t_s in the S2-S3 and S3-Go epochs cannot be a simple artifact of applying KiNeT to complex spatiotemporal responses, as we found no evidence for t_s -dependent speed in the S1-S2 epoch despite the presence of complex non-monotonic response profiles.

Linking behaviorally relevant computations to internal models.

Results of KiNeT provided evidence for the presence of a speed command driving a simulated motor plan during the S2-S3 epoch, and an actual motor plan in the S3-Go epoch. However, according to the internal model hypothesis, the speed command and the simulation have to be guided by an estimator. Therefore, we predicted that δ^t and $\mathbf{V}^{[t_s]}$ should reflect the animal's internal estimate of t_s , and not the experimentally controlled t_s . To test this prediction, we used the fits of EKF to the animal's behavior to infer its internal estimates immediately after S1, S2 and S3 – $t_{e(S1)}$, $t_{e(S2)}$ and $t_{e(S3)}$. According to EKF, $t_{e(S1)}$ was equal to the mean of the prior (Fig. 8a, left), $t_{e(S2)}$ follows t_s but was biased toward the prior mean (Fig. 8a, center), and $t_{e(S3)}$ follows t_s with less bias than $t_{e(S2)}$ (Fig. 8a, right).

We leveraged these prior dependent biases to ask whether the speed of different neural trajectories was ordered according to t_s (no bias) or $t_{e(SX)}$ (biased) as prescribed by EKF. Since speed is inversely proportional to time, we used KiNeT to restate this question as follows: does $\mathbf{t}^{[t_s]}$ more accurately correspond to t_s or t_e ? We converted $\mathbf{t}^{[t_s]}$ to an inferred estimate, \hat{t}_e , by multiplying the slope of the regression line relating $\mathbf{t}^{[t_s]}$ to $\mathbf{t}^{[800]}$ ($\Delta \mathbf{t}^{[t_s]} / \Delta \mathbf{t}^{[800]}$; Fig. 8b, left) by the reference interval of 800 ms (see Methods). Clearly, \hat{t}_e exhibited the biases observed in $t_{e(S2)}$ and $t_{e(S3)}$ during S2-S3 and S3-Go epochs, respectively (Fig. 8b, center and right panels). This relationship was not explained by the monotonic relationship between \hat{t}_e with t_s , as the organization of \hat{t}_e derived from neural activity was far better predicted by the animal's estimates than t_s (Fig. 8b, insets, see Methods).

We also asked whether the t_s -dependent position of neural trajectories that provide a correlate of the speed command reflected \hat{t}_e more accurately than t_s . To test this, we asked whether average distances between neural trajectories within each epoch were more accurately captured by the experimentally controlled t_s or by the behaviorally inferred \hat{t}_e in that epoch (see Methods). Again, we found that the distances between trajectories were more similar to $t_{e(S2)}$ and $t_{e(S3)}$ than t_s (Fig. 8c, center and right). These observations suggest that the neural correlates of both the speed command and the predictive dynamics associated with the simulated motor plan correspond to the animal's estimates (see Supplementary Fig. 13 for individual animals). In other words, changes in neural activity after each flash reflect the estimation process prescribed by EKF.

A final prediction of the internal model hypothesis is that the speed and distance of neural trajectories ought to reflect the integration of each new measurement. This integration was clearly evident during the transition from the S1-S2 to the S2-S3 epoch. That is, speeds and distances were independent of t_s before S2, and were systematically organized according to t_s after S2. To assess whether a similar integration was evident during the transition from the S2-S3 to S3-Go epoch, we returned to the predictions of the EKF model

(Fig. 8a). According to the model, both $t_{e(S2)}$ and $t_{e(S3)}$ exhibit a systematic bias toward the mean of the prior distribution, but the overall bias was smaller in the S3-Go epoch¹² (Fig. 8a, right). To examine whether this pattern was present in neural trajectories, we compared the bias in both the speeds and distances of neural trajectories in the S2-S3 and S3-Go epochs (see Methods). The analysis of speeds provided clear evidence that biases were larger during the S2-S3 epoch than the S3-Go epoch (Fig. 8d, top). We also measured the degree of bias in the distance between neural trajectories. Because neural trajectories in the S2-S3 and S3-Go epochs were analyzed using PCs derived from activity in their respective epochs, a direct comparison of the distances was not warranted. Therefore, we computed the bias based on relative distances within each epoch (see Methods). Again, results indicated that the degree of bias was larger during the S2-S3 epoch than the S3-Go epoch (Fig. 8d, bottom). These observations validated our hypothesis that DMFC responses reflected the animal's estimate in accordance with sequential updating after both S2 and S3 flashes.

Discussion

Our results demonstrate that internal models can be used to understand the algorithm the brain uses to control internally generated cortical activity patterns in the absence of overt movements. This is an important finding as it indicates that the internal model hypothesis, advanced by decades of research on sensorimotor function^{5,22–26} and motor control^{4,10,27–29}, may also be used to understand how the brain controls mental states that are not observable.

A prominent result from our analysis of individual neurons was the heterogeneity of responses. A few neurons exhibited characteristics consistent with a speed command (that is, control signal) and predictive dynamics (that is, simulation), but the responses of most neurons reflected a mixture of both. This motivated a complementary analysis across the population, which led to three important findings. First, the distance between neural trajectories provided evidence for the presence of a persistent speed command consistent with the output of a controller. Second, the speed with which trajectories evolved through time revealed a predictive process consistent with the output of a simulated motor plan in the S2-S3 epoch and an ongoing motor plan in the S3-Go epoch. Third, the speed information inferred from neural activity reflected the sequential updating consistent with the output of an underlying estimator.

Although the activity in the DMFC was consistent with the computations associated with the controller, simulator and estimator, we do not know whether and how the DMFC might implement these computations. We have previously established that inactivation of the DMFC interferes with time interval production⁸. Given that active and simulated motor plans rely on similar dynamics, it is possible that the DMFC contributes to both³⁰. However, the tight interaction between the DMFC and other cortical and subcortical areas^{31–35} suggests that other brain areas are also involved.

Our previous work on a simple time production task suggests that the speed command in the DMFC may originate in the thalamus⁸. It is therefore possible that the speed command in the 1-2-3-Go task is also driven by a tonic thalamic input. The source of the speed command in the thalamus has not been identified. One possibility is the dentate nucleus of the cerebellum that trans-thalamically drives the DMFC^{36–38}, and is thought to support interval timing in non-motor tasks^{39–42}. This interpretation is consistent with experimental and modeling work suggesting a role for ascending cerebellar signals in driving cortical dynamics⁴³. However, if the speed command is indeed provided by the cerebellum, then performing sequential updating as in the 1-2-3-Go task will have to engage the entire cortico-cerebellar loop so that the command can be updated in the cerebellum depending on the state of the simulated motor plan in cortical circuits.

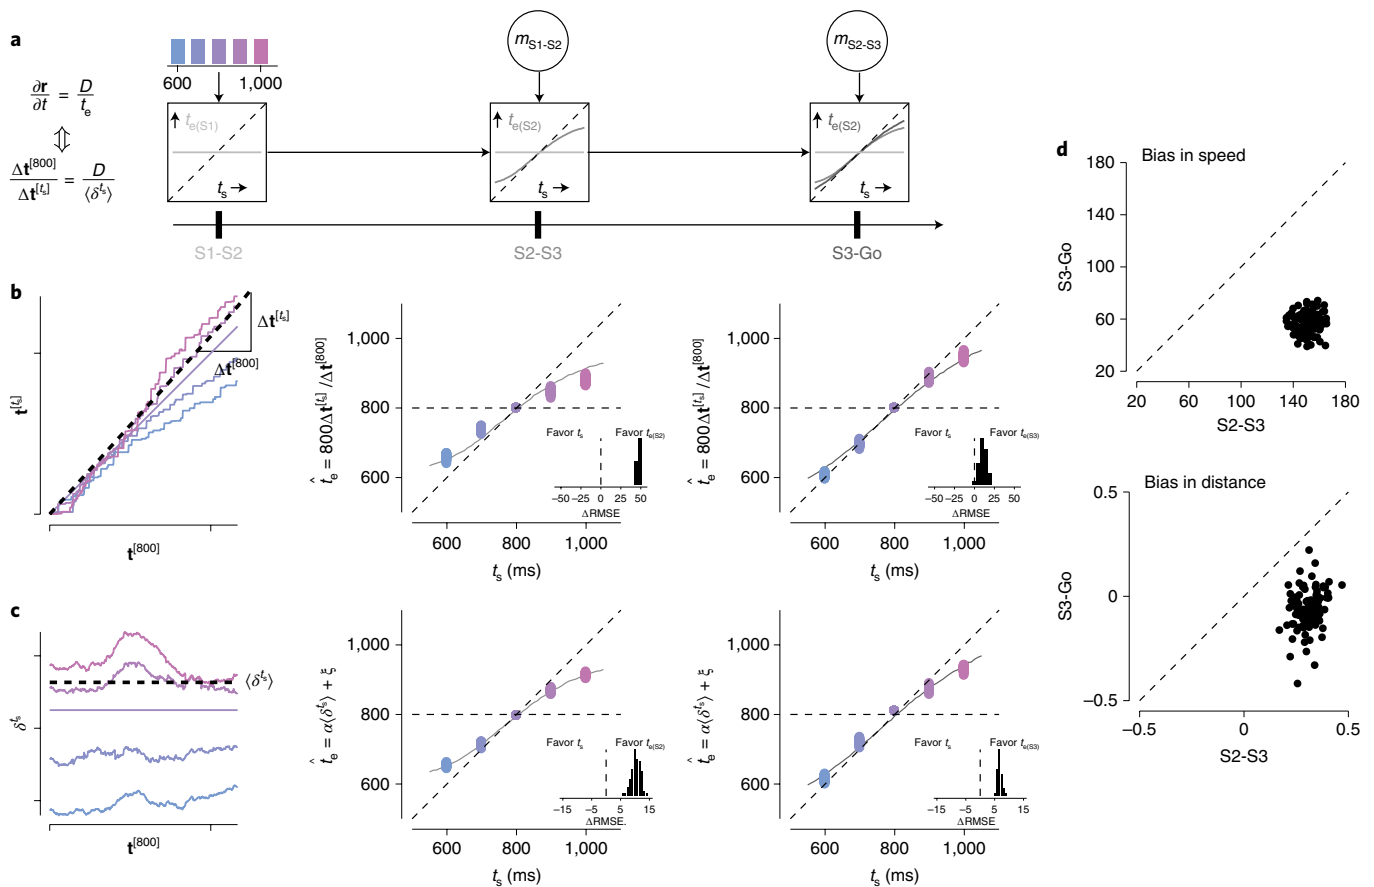


Fig. 8 | Speed and distance between neural trajectories reflect animals' internal estimates. **a**, Interval estimates derived from EKF model (same format as in Fig. 1g). **b**, Interval estimates inferred from the speed of neural trajectories (\hat{t}_e) as a function of t_s . Left: the procedure for inferring \hat{t}_e for each t_s . The colored traces show $\mathbf{t}^{[t_s]}$ as a function of $\mathbf{t}^{[800]}$ for a random draw from the bootstrap distribution (see Methods). To infer \hat{t}_e for each t_s we measured the slope of the regression line relating $\mathbf{t}^{[t_s]}$ to $\mathbf{t}^{[800]}$ ($\Delta \mathbf{t}^{[t_s]} / \Delta \mathbf{t}^{[800]}$) and then scaled the slope by the duration of the reference interval (800 ms). The broken line represents the regression slope for $t_s = 900$ ms. Middle and right: \hat{t}_e as a function of t_s inferred from the speed of neural trajectories in the S2-S3 and S3-Go epochs. The colored circles show multiple estimates of \hat{t}_e derived from bootstrapping ($n = 100$). The unbroken curves show interval estimates derived from EKF model fits ($t_{e(S2)}$ in the middle panel and $t_{e(S3)}$ in the right panel). Unity indicates perfect estimates of t_s and the horizontal line represents the mean of the prior. Insets show the difference in RMSE between models that assume that the speed of neural trajectories reflects t_s versus $t_{e(S2)}$ (middle) or $t_{e(S3)}$ (right). Positive values indicate neural data that were more accurately captured by EKF. **c**, The value of \hat{t}_e inferred from the distance between neural trajectories as a function of t_s . Left: the procedure for inferring \hat{t}_e for each t_s . The colored traces show δ^{t_s} as a function of time for a random draw from the bootstrap distribution (see Methods). We inferred \hat{t}_e for each t_s using a linear transformation of the average distance ($\alpha \langle \delta^{t_s} \rangle + \xi$). The broken line shows the average distance, $\langle \delta^{t_s} \rangle$, for $t_s = 900$ ms. Middle and right: \hat{t}_e as a function of t_s inferred from the distances between neural trajectories in the S2-S3 and S3-Go epochs (same format as **b**). Insets show the differences in RMSE between models assuming the distances of neural trajectories reflect t_s versus $t_{e(S2)}$ (middle) or $t_{e(S3)}$ (right). **d**, Comparison of the degree of Bias in speed and distance during the S2-S3 and S3-Go epochs. Broken lines represent the unity. See also Supplementary Fig. 13.

More generally, since the generation of simulated and actual motor plans involves recurrent interactions among neurons, making definitive statements about the nature of the control signal is difficult. A dynamical system can generate task-dependent dynamics either through adjustment of initial conditions or via direct inputs^{7,44}. In our work, we interpreted the presence of the persistent t_s -dependent signals across the population as a persistent control signal. However, it is also possible that this persistent signal results from setting the initial conditions of the system. Consistent with this interpretation, trajectories associated with the S1-S2 and S2-S3 epochs terminated at different states, thus providing a potential substrate for using initial conditions as the control signal. If so, we would need to reinterpret the action of the control signal as being mediated through initial conditions. Exerting control through initial conditions instead of a tonic input does not change our conclusions with respect to understanding the control

of dynamics in terms of internal models, but points to a different underlying implementation.

We also do not know how the DMFC alone, or through interactions with other brain areas, implements an estimator. The function of the estimator is to use the state of the simulated motor plan at the time of feedback to update the speed command. Since the establishment of a suitable estimator demands learning and optimization^{5,10,45}, a likely substrate for establishing an estimator could be synaptic learning of the recurrent connections in various cortical areas, including the DMFC. The key requirement for estimation is to have a network that can adjust its responses to inputs (for example, flashes) in a state-dependent manner (that is, the state of the simulated plan). Recent theoretical work suggests that this computation is well within the capacity of recurrent neural networks^{7,46,47,9} and can be mediated through suitable connectivity⁴⁸. If so, alternations of neural states immediately before the

flash using precise perturbation techniques should lead to specific errors in updating.

Finally, our findings regarding the implementation of a simulated motor plan at the level of cortical population activity highlights the possibility of a novel mechanism for performing predictive computations. Theoretical considerations have long noted the computational advantages of predictive information processing⁴⁹, and computational models have proposed a number of possible mechanisms for its implementation in cortical microcircuits⁵⁰. Our work provides evidence for a novel mechanism for generating top-down predictions afforded by population-level latent dynamics of cortical networks.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, statements of code and data availability and associated accession codes are available at <https://doi.org/10.1038/s41593-019-0500-6>.

Received: 4 April 2018; Accepted: 16 August 2019;
Published online: 7 October 2019

References

- Wolpert, D. M. & Ghahramani, Z. Computational principles of movement neuroscience. *Nat. Neurosci.* **3**, 1212–1217 (2000).
- Todorov, E. Optimality principles in sensorimotor control. *Nat. Neurosci.* **7**, 907–915 (2004).
- Scott, S. H. Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* **5**, 532–546 (2004).
- Shadmehr, R. & Krakauer, J. W. A computational neuroanatomy for motor control. *Exp. Brain Res.* **185**, 359–381 (2008).
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).
- Ito, M. Control of mental activities by internal models in the cerebellum. *Nat. Rev. Neurosci.* **9**, 304–313 (2008).
- Remington, E. D., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019.e5 (2018).
- Wang, J., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
- Sohn, H., Narain, D., Meirhaeghe, N. & Jazayeri, M. Bayesian computation through cortical latent dynamics. *Neuron* **103**, 934–947.e5 (2019).
- Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
- Gallistel, C. R. & Gibbon, J. Time, rate, and conditioning. *Psychol. Rev.* **107**, 289–344 (2000).
- Egger, S. W. & Jazayeri, M. A nonlinear updating algorithm captures suboptimal inference in the presence of signal-dependent noise. *Sci. Rep.* **8**, 12597 (2018).
- Acerbi, L., Wolpert, D. M. & Vijayakumar, S. Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Comput. Biol.* **8**, e1002771 (2012).
- Jazayeri, M. & Shadlen, M. N. A neural mechanism for sensing and reproducing a time interval. *Curr. Biol.* **25**, 2599–2609 (2015).
- Kurata, K. & Wise, S. P. Premotor and supplementary motor cortex in rhesus monkeys: neuronal activity during externally- and internally-instructed motor tasks. *Exp. Brain Res.* **72**, 237–248 (1988).
- Mita, A., Mushiaki, H., Shima, K., Matsuzaka, Y. & Tanji, J. Interval time coding by neurons in the presupplementary and supplementary motor areas. *Nat. Neurosci.* **12**, 502–507 (2009).
- Kunimatsu, J. & Tanaka, M. Alteration of the timing of self-initiated but not reactive saccades by electrical stimulation in the supplementary eye field. *Eur. J. Neurosci.* **36**, 3258–3268 (2012).
- Merchant, H., Pérez, O., Zarco, W. & Gámez, J. Interval tuning in the primate medial premotor cortex as a general timing mechanism. *J. Neurosci.* **33**, 9082–9096 (2013).
- Merchant, H., Zarco, W., Pérez, O., Prado, L. & Bartolo, R. Measuring time with different neural chronometers during a synchronization-continuation task. *Proc. Natl Acad. Sci. USA* **108**, 19784–19789 (2011).
- Murakami, M., Vicente, M. I., Costa, G. M. & Mainen, Z. F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* **17**, 1574–1582 (2014).
- Merchant, H. & Averbeck, B. B. The computational and neural basis of rhythmic timing in medial premotor cortex. *J. Neurosci.* <https://doi.org/10.1523/JNEUROSCI.0367-17.2017> (2017).
- Franklin, D. W. & Wolpert, D. M. Computational mechanisms of sensorimotor control. *Neuron* **72**, 425–442 (2011).
- Kawato, M., Furukawa, K. & Suzuki, R. A hierarchical neural-network model for control and learning of voluntary movement. *Biol. Cybern.* **57**, 169–185 (1987).
- Angelaki, D. E., Shaikh, A. G., Green, A. M. & Dickman, J. D. Neurons compute internal models of the physical laws of motion. *Nature* **430**, 560–564 (2004).
- Sommer, M. A. & Wurtz, R. H. Influence of the thalamus on spatial visual processing in frontal cortex. *Nature* **444**, 374–377 (2006).
- Desmurget, M. & Grafton, S. Forward modeling allows feedback control for fast reaching movements. *Trends Cogn. Sci.* **4**, 423–431 (2000).
- Sabes, P. N. The planning and control of reaching movements. *Curr. Opin. Neurobiol.* **10**, 740–746 (2000).
- Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
- Sommer, M. A. & Wurtz, R. H. A pathway in primate brain for internal monitoring of movements. *Science* **296**, 1480–1482 (2002).
- Cadena-Valencia, V., García-Garibay, O., Merchant, H., Jazayeri, M. & de Lafuente, V. Entrainment and maintenance of an internal metronome in supplementary motor area. *eLife* **7**, e38983 (2018).
- Shook, B. L., Schlag-Rey, M. & Schlag, J. Primate supplementary eye field. II. Comparative aspects of connections with the thalamus, corpus striatum, and related forebrain nuclei. *J. Comp. Neurol.* **307**, 562–583 (1991).
- Shook, B. L., Schlag-Rey, M. & Schlag, J. Primate supplementary eye field: I. Comparative aspects of mesencephalic and pontine connections. *J. Comp. Neurol.* **301**, 618–642 (1990).
- Picard, N. & Strick, P. L. Motor areas of the medial wall: a review of their location and functional activation. *Cereb. Cortex* **6**, 342–353 (1996).
- Schall, J. D. in *Extrastriate Cortex in Primates* (eds Rockland, K. S., Kaas, J. H. & Peters, A.) 527–638 (Springer, 1997).
- Matelli, M., Luppino, G. & Rizzolatti, G. Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. *J. Comp. Neurol.* **311**, 445–462 (1991).
- Asanuma, C., Thach, W. R. & Jones, E. G. Anatomical evidence for segregated focal groupings of efferent cells and their terminal ramifications in the cerebellothalamic pathway of the monkey. *Brain Res.* **286**, 267–297 (1983).
- Lynch, J. C., Hoover, J. E. & Strick, P. L. Input to the primate frontal eye field from the substantia nigra, superior colliculus, and dentate nucleus demonstrated by transneuronal transport. *Exp. Brain Res.* **100**, 181–186 (1994).
- Prevosto, V., Graf, W. & Ugolini, G. Cerebellar inputs to intraparietal cortex areas LIP and MIP: functional frameworks for adaptive control of eye movements, reaching, and arm/eye/head movement coordination. *Cereb. Cortex* **20**, 214–228 (2010).
- Kunimatsu, J., Suzuki, T. W., Ohmae, S. & Tanaka, M. Different contributions of preparatory activity in the basal ganglia and cerebellum for self-timing. *eLife* **7**, e35676 (2018).
- Ashmore, R. C. & Sommer, M. A. Delay activity of saccade-related neurons in the caudal dentate nucleus of the macaque cerebellum. *J. Neurophysiol.* **109**, 2129–2144 (2013).
- Ohmae, S., Uematsu, A. & Tanaka, M. Temporally specific sensory signals for the detection of stimulus omission in the primate deep cerebellar nuclei. *J. Neurosci.* **33**, 15432–15441 (2013).
- Narain, D., Remington, E. D., Zeeuw, C. I. D. & Jazayeri, M. A cerebellar mechanism for learning prior distributions of time intervals. *Nat. Commun.* **9**, 469 (2018).
- Gao, Z. et al. A cortico-cerebellar loop for motor planning. *Nature* <https://doi.org/10.1038/s41586-018-0633-x> (2018).
- Remington, E. D., Egger, S. W., Narain, D., Wang, J. & Jazayeri, M. A dynamical systems perspective on flexible motor timing. *Trends Cogn. Sci.* **22**, 938–952 (2018).
- Berniker, M. & Kording, K. Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci.* **11**, 1454–1461 (2008).
- Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
- Chaisangmongkon, W., Swaminathan, S. K., Freedman, D. J. & Wang, X.-J. Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron* **93**, 1504–1517.e4 (2017).
- Mastrogiuseppe, F. & Ostojic, S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* **99**, 609–623.e29 (2018).
- Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
- Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).

Acknowledgements

M.J. is supported by the NIH (NINDS-NS078127), the Sloan Foundation, the Klingenstein Foundation, the Simons Foundation (grants 325542 and 542993SPI),

the McKnight Foundation, the Center for Sensorimotor Neural Engineering (grant UWSC6200 (BPO4405)) and the McGovern Institute.

Author contributions

S.W.E. and M.J. designed the task. S.W.E. collected behavioral and neural data from monkey B. C.-J.C. collected behavioral and neural data from monkey G. E.D.R. designed the KiNeT analysis. S.W.E. performed all analyses. S.W.E. and M.J. interpreted the results and wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41593-019-0500-6>.

Correspondence and requests for materials should be addressed to M.J.

Peer review information *Nature Neuroscience* thanks Matthew Kaufman, Reza Shadmehr and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

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

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2019

Methods

Experiments involved the in vivo extracellular recording of neural activity from the DMFC of two awake, behaving male monkeys (*Macaca mulatta*) trained to perform the 1-2-3-Go behavioral task. During experimental sessions, monkeys were seated comfortably in a dark and quiet room. Stimuli and behavioral contingencies were controlled using MWorks (<https://mworks.github.io/>). Visual stimuli were presented on a frontoparallel 23-inch Acer H236HL monitor at a resolution of 1,920 × 1,080 and a refresh rate of 60 Hz. Eye positions were tracked with an infrared camera (Eyelink 1000; SR Research) and sampled at 1 kHz. All experimental procedures conformed to the guidelines of the National Institutes of Health and were approved by the Committee of Animal Care at Massachusetts Institute of Technology. Randomization and blinding were not applicable at the level of subjects, as the animals were not assigned to experimental groups. However, trials and conditions for each animal were randomized. When warranted, data analyses were subjected to blinding using cross-validation.

Task and stimuli. Each trial began with the presentation of a central fixation point (FP; circular, red, 0.5° in diameter). Following fixation (with delay selected from a uniform hazard, minimum = 100 ms, mean = 400 ms), a saccadic target (circular, white, 1.5° in diameter) was presented 10° to the left or right of the FP. Afterwards (uniform hazard, minimum = 500 ms, mean = 1,000 ms), three flashed visual stimuli (S1, S2 and S3) were presented. Each flash was presented for 100 ms as an annulus around the FP (white, inside diameter = 2.5°, outer diameter = 3°). The time between consecutive flashes was fixed within a trial and was drawn at random from a discrete uniform distribution across trials (Fig. 1c; 5 values, minimum = 600 ms, maximum = 1,000 ms). We refer to the time between flashes as the sample interval, t_s . Monkeys had to maintain fixation until after S3 within an electronic window around the FP (7° in diameter). After S3, monkeys had to proactively initiate a saccade to the saccadic target to produce an interval equal to t_s . The produced interval, t_p , was measured as the time between S3 and when the eye of the animal entered an electronic window (7° in diameter) around the saccadic target. Early fixation breaks (before 200 ms post S3) or extremely long production times ($t_p > t_s + 1,000$ ms) were considered aborted trials and were discarded. We additionally removed outlier trials whereby the value of t_p was more than three standard deviations from the mean t_p for a given t_s . Animals received trial-by-trial visual feedback on the magnitude and sign of error by a visual stimulus (circular, 0.5° in diameter) presented immediately after the saccade along the line connecting the FP to the saccadic target. The magnitude of the relative error ($|t_p - t_s|/t_s$) was represented by the distance of the feedback stimulus to the saccadic target, and positive (negative) errors were associated with locations farther away (closer to) the FP. When the magnitude of the relative error was smaller than a threshold, both the saccadic target and the feedback stimulus turned green and the animal received a juice reward. The magnitude of the reward decreased linearly with the magnitude of the relative error (Fig. 1d). If the relative error was larger than the threshold, the target stimulus and analog feedback remained white and no reward was delivered. Performance was evaluated according to the relative error (as opposed to absolute error) to accommodate the scalar variability of time interval production¹¹. The reward threshold was initialized at 0.15 at the start of every session and was adjusted adaptively and on a trial-by-trial basis in a one-up one-down fashion with a fixed increment/decrement of 0.001 for unrewarded/rewarded trials, respectively. With this scheme, animals received reward in approximately 50% of the trials.

We quantified the overall bias of the animal as the root-mean square of bias for each t_s :

$$\text{Bias} = \sqrt{\frac{1}{K} \sum_{i=1}^K (\bar{t}_p - t_{s_i})^2} \quad (1)$$

where \bar{t}_p is the mean t_p across trials of the i -th sample interval t_{s_i} .

Bayesian integration model of behavior. We developed a Bayesian observer model for the 1-2-3-Go task based on a model that was previously used to account for behavior in a timing task with a single measurement interval^{12,14,51}. The model computed the posterior of t_s from the prior, $p(t_s)$, and the likelihood function, $\lambda(m|t_s)$, from two independent measurements m_{S1-S2} and m_{S2-S3} and used the mean of the posterior (shown with angle brackets) as the final estimate, e_{S3-Go} , as follows:

$$p(t_s | m_{S1-S2}, m_{S2-S3}) = \frac{\lambda(m_{S1-S2}|t_s)\lambda(m_{S2-S3}|t_s)p(t_s)}{p(m_{S1-S2}, m_{S2-S3})} \quad (2)$$

$$e_{S3-Go} = f_{\text{BLS}}(m_{S1-S2}, m_{S2-S3}) = \langle p(t_s | m_{S1-S2}, m_{S2-S3}) \rangle \quad (3)$$

Where f_{BLS} represents the Bayes-least-squares estimator of t_s . The likelihood function for each measurement was modeled based on the assumption that the measurement is subject to zero-mean Gaussian noise whose standard deviation scales with t_s according to the Weber fraction w_m (Supplementary Fig. 1a).

$$\lambda(m|t_s) = \frac{1}{\sqrt{2\pi w_m^2 t_s^2}} \exp\left(-\frac{(m - t_s)^2}{2w_m^2 t_s^2}\right) \quad (4)$$

Finally, we assumed that the production interval, t_p , is subject to zero-mean Gaussian noise whose standard deviation scales with e_{S3-Go} according to the Weber fraction w_p (Supplementary Fig. 1a). To fit the model to animals' behavior, we augmented the model by an offset term, b , to account for stimulus- and prior-independent biases observed in responses as follows:

$$p(t_p | e_{S3-Go}) = \frac{1}{\sqrt{2\pi w_p^2 e_{S3-Go}^2}} \exp\left(-\frac{(t_p - e_{S3-Go} - b)^2}{2w_p^2 e_{S3-Go}^2}\right) \quad (5)$$

EKF model of interval estimation. Following our previous efforts to model human behavior in a similar task, we developed an updating model of interval estimation that approximates the fully Bayesian estimate¹². The model implements a sequential updating scheme akin to the Kalman filter⁵². According to EKF, at S1, the estimated duration is the mean of the prior distribution as follows:

$$e_{S1-S2} = \int t_s p(t_s) dt_s \quad (6)$$

At the time of S2, the observer combines the measurement of the S1-S2 interval, m_{S1-S2} , with e_{S1-S2} to generate a new estimate, e_{S2-S3} , as follows:

$$e_{S2-S3} = e_{S1-S2} + k_{S2} f^*(m_{S1-S2} - e_{S1-S2}) \quad (7)$$

with $k_{S2} = 1 \cdot f^*$ is a nonlinear function based on the Bayesian model after one measurement, $f^*(m_{S1-S2}) = \langle p(t_s | m_{S1-S2}) \rangle$. At the time of S3, the model combines the new measurement m_{S2-S3} with e_{S2-S3} to generate a final estimate of t_s as follows:

$$e_{S3-Go} = e_{S2-S3} + k_{S3} f^*(m_{S2-S3} - e_{S2-S3}) \quad (8)$$

with $k_{S3} = 0.5$. This final estimate is used to guide production, with noise in the production phase described by equation (5). See also Supplementary Fig. 1.

Model fitting. We used a maximum likelihood procedure to fit the Bayesian and EKF models to the data. Assuming t_p values are conditionally independent across trials, the log-likelihood of model parameters can be formulated as follows:

$$\log p(t_p^1, t_p^2, \dots, t_p^N | t_s, w_m, w_p, b) = \sum_{i=1}^N \log p(t_p^i | t_s, w_m, w_p, b) \quad (9)$$

where the superscripts denote the trial number. We found the parameters that maximize equation (9) with a constrained optimization function.

We used the fitted model to additionally infer the expected value of the animal's estimate after each of the S1, S2 and S3 flashes, which we denote as $t_{e(S1)}$, $t_{e(S2)}$ and $t_{e(S3)}$, respectively. Because the only information available to the animal after S1 is the prior distribution, $p(t_s)$, $t_{e(S1)}$ was set to the prior mean. We derived $t_{e(S2)}$ and $t_{e(S3)}$ by averaging the EKF estimates (e_{S2-S3} and e_{S3-Go} , respectively) across measurements for each t_s .

Cue conflict trials. To test whether animals integrated both measurements, 33% of the trials with $t_s = 800$ ms were modified such that either S1-S2 or S2-S3 were made either 50-ms longer or shorter. We verified that animals used both measurements in two ways. First, we tested whether t_p was biased in a direction expected under the integration hypothesis, and whether the magnitude of bias depended on which epoch was jittered. Second, we asked which of the following three models best explained the behavior in the cue conflict trials: (1) monkeys used both m_{S1-S2} and m_{S2-S3} to guide production; (2) monkeys used m_{S1-S2} only; or (3) monkeys used m_{S2-S3} only. All three models had the same parameters (w_m , w_p , b). We used a maximum likelihood procedure to fit the models to the data without conflict and compared the likelihood of the conflict data under each model.

Analysis of DMFC data. Electrophysiological recordings were made using one or two 16- or 24-channel laminar V-probes (Plexon) through craniotomies located above the DMFC. In total, we analyzed neural activity from 16 recording sessions (12 from monkey B and 4 from monkey G). We targeted the region of interest using the following procedure. We started by gathering anatomical information about the regions of DMFC associated with the SEF^{31,53-55} and the pre-SMA^{55,56} using previous studies of rhesus macaques. We compiled information about reported stereotactic coordinates, histological plots and MRI scans to identify the general region of interest. We then sampled the DMFC starting from the center of the identified region and moving outward in a quasi-systematic fashion, looking for regions with strong modulation during any epoch of the task. Neural activity in monkey B was recorded from between 0.5 mm to 2.9 mm lateral of the midline and 2.2 mm posterior to 8.1 mm anterior of the genu of the arcuate sulcus. Neural activity in monkey G was recorded from between 0.9 mm to 4.6 mm lateral of the midline and 1.8 mm to 9.3 mm anterior of the genu of the arcuate sulcus. Because

the targeted recording sites were defined anatomically, we conservatively refer to the recorded region as the DMFC. This anatomically defined region potentially comprised the functionally defined areas referred to as the SEF, the pre-SMA and the dorsal SMA (that is, outside the medial bank). Data were stored using a Blackrock Cerebus Neural Signal Processor (Blackrock Microsystems) and analyzed offline using custom code in Matlab (R2017a, The MathWorks, 2000).

Single-unit and multi-unit activity were extracted from band-pass-filtered voltage traces by an initial threshold crossing followed by spike sorting. We used the following procedure for spike sorting. (1) A Gaussian mixture model was used to cluster waveforms on each channel that crossed a threshold. Let us assume that the data were best fit using a Gaussian mixture model with N clusters. (2) We used the fitted model to generate surrogate samples from each of the N clusters. (3) We used maximum likelihood to classify the surrogate samples (that is, which cluster were they sampled from). (4) For each cluster, we used the classification results to quantify the probability of assigning surrogate samples from other clusters as originating to the current cluster, which we refer to as the probability of false alarm, P_{FA} . Note that the ground truth for the surrogate samples was known. (5) We only accepted units as well-isolated if surrogate samples from the corresponding cluster were associated with $P_{FA} < 0.05$.

The spiking data for each neuron were first smoothed over time using a 150-ms boxcar after alignment to different events in the task (that is, S1, S2, S3 and Go). Trial-by-trial data were then conditioned on t_s to measure peristimulus time histograms and to perform PCA. We assigned all conflict trials to the $t_s = 800$ ms condition for these analyses. Population analyses were based on all 181 units collected across the 16 individual recording sessions, after excluding units with poor signal-to-noise or that were unmodulated by the task. Single-unit analysis was based on 115 well-isolated units. No statistical methods were used to predetermine sample sizes, but our sample sizes were similar to those reported in previous publications^{7,15,16}. Supplementary analysis was separately performed on units from each monkey.

To characterize the firing rate profiles over time, we fit a polynomial of order n to the mean firing rates across t_s calculated using a random half of the trials without replacement. We assessed the goodness of fit by measuring the correlation coefficient between the expected firing rate, based on the polynomial fit, and the mean firing rate calculated from the left-out trials. We also assessed polynomial fits to temporally scaled firing rates using the same procedure.

To assess the degree of stationarity in the encoding of t_s by individual neurons, we first assigned the trial-by-trial spike counts into two groups at random (without replacement). For each group, we calculated the firing rate as a function of time and t_s , $r_1(t, t_s)$ and $r_2(t, t_s)$. We calculated the correlation coefficient between the firing rates at two different time points, $r_1(t', t_s)$ and $r_2(t'', t_s)$, across t_s . We performed this analysis for each combination of times, t' and t'' , with t' and t'' between 0 and 600 ms relative to the onset of the epoch of interest. We used the resulting matrix of correlation coefficients to assess the degree of similarity in firing rates as a function of t_s at different time points. This procedure enabled us to measure similarity without making assumptions about the exact nature of encoding of t_s .

LNP model. Neural encoding of t_s by individual neurons was also assessed by fitting a LNP model to spike count data at $t = 0, 100, 200, 300, 400$ or 500 ms following each flash (S1, S2 and S3). We modeled spike count data as a nonhomogeneous Poisson process with a t_s -dependent rate function as follows:

$$\lambda_n(t) = e^{\beta(t)t_s + c} \quad (10)$$

With $\beta(t)$ parameterizing the sensitivity of the neuron to t_s at time t . The probability of a given spike count on a particular trial is as follows:

$$p[r_n(t)|\lambda_n(t)] = \frac{\lambda_n(t)^{r_n(t)} \exp[-\lambda_n(t)]}{r_n(t)!} \quad (11)$$

To fit the parameters of the model and assess goodness of fit, we split the data into training and validation sets, designating every other trial as training data. We fit the model by maximizing the log-likelihood of model parameters given t_s and the observed spike count, $r_n(t)$, of the n th neuron across $i = 1, \dots, M$ training trials at time t as follows:

$$\mathcal{L}(\beta(t), c) = \sum_{i=1}^M [r_n^i(t) \log \lambda_n(t) - \lambda_n(t)] \quad (12)$$

We assessed the significance of sensitivity to t_s by comparing the LNP model with the t_s -dependent rate function (above) to another model with a t_s -independent rate function as follows:

$$\lambda_n(t) = c \quad (13)$$

We compared the models using Bayesian information criterion (BIC) and took $\text{BIC} > 10$ in favor of the t_s -dependent model as evidence for significant sensitivity to t_s .

Decoding t_s using the LNP model. We used the LNP model fits to the training data to compute the maximum likelihood estimate of t_s , \hat{t}_s^i , on validation trial i , assuming that the firing of the N neurons is conditionally independent:

$$\mathcal{L}(t_s)^i = \sum_{n=1}^N [r_n^i(t) \log \lambda_n(t) - \lambda_n(t)] \quad (14)$$

$$\hat{t}_s^i = \text{argmax}_{t_s} \mathcal{L}(t_s)^i \quad (15)$$

We assessed decoding performance by calculating the correlation coefficient between the actual and decoded values of t_s . We also considered the performance of a decoder that assumes that the representation of t_s is stationary following each flash. To do so, we held $\beta(t)$ at the value fit at the time of S1, S2 or S3 for decoding t_s during the S1-S2, S2-S3 and S3-Go epochs, respectively.

PCA. To perform PCA, we averaged spike times by using a 150-ms boxcar filter and used the square root of spike counts to reduce the effect of Poisson noise on analyses⁵⁷. Data from each unit were then trial-averaged and z-scored by removing the mean and normalizing by the standard deviation calculated across time and conditions. Finally, z-scored data over time were arranged in a data matrix to perform PCA, with rows of the data matrix concatenating data across time and t_s condition and columns corresponding to units collected across the 16 individual recording sessions. To visualize activity patterns from S1 to Go (Fig. 6a) within a single coordinate system, we applied PCA to neural activity from S1 to Go. We also performed PCA on activity within each epoch.

KiNeT. We developed an analysis technique to infer the relative position and speed of multiple low-dimensional neural trajectories based on the organization of sets of points along those trajectories. KiNeT can be broken down to the following six steps. (1) Choose a subspace to perform the analysis. In our case, we focused on the projection of the neural activity onto the first ten PCs, which captured more than 93% of variance in the data. (2) Designate one of the trajectories as a reference trajectory. In our case, we designated the median t_s (800 ms) as the reference. (3) Sample the reference trajectory at a given time resolution and compute a vector of associated time points. In our case, we used a time resolution of 1 ms and denoted the resulting vector of time points by $\mathbf{t}^{[800]}$. (4) For each element of $\mathbf{t}^{[800]}$, find the corresponding neural state vector on $\Omega^{[800]}$, $\mathbf{r}^{[800]}$. The set of state vectors was arranged in a matrix, $R^{[800]}$, with each column associated with a time point in $\mathbf{t}^{[800]}$. (5) For each neural state in $R^{[800]}$, find the nearest point (minimum Euclidean distance) on all the other (non-reference) trajectories. This provides a corresponding matrix of states on each non-reference trajectory, which we denoted by $R^{[t_s]}$ (the superscript is a reference to the trajectory associated with t_s). (6) Find the time when the neural state reaches each state in $R^{[t_s]}$. This furnishes a corresponding set of times, which we denoted by $\mathbf{t}^{[t_s]}$. The four arrays, $\mathbf{t}^{[800]}$, $R^{[800]}$, $\mathbf{t}^{[t_s]}$ and $R^{[t_s]}$ can then be used to evaluate the relative geometry and speed of the neural trajectories.

We estimated the relative position of neural trajectories by measuring the distance between each $\mathbf{r}^{[t_s]}$ and the corresponding reference state $\mathbf{r}^{[800]}$ along the direction defined by the local vector connecting $\mathbf{r}^{[1,000]}$ to $\mathbf{r}^{[600]}$. Mathematically, this is equivalent to finding the projection of the difference vector $\delta^{t_s} = \mathbf{r}^{[t_s]} - \mathbf{r}^{[800]}$ onto the base vector, $\mathbf{d} = \mathbf{r}^{[1,000]} - \mathbf{r}^{[600]}$, according to $\delta^{t_s} = \mathbf{d}^T \delta^{t_s} / \|\mathbf{d}\|$. Repeating the process for each state in $R^{[t_s]}$ results in a scalar quantity for each point in $\mathbf{t}^{[800]}$. The result of this analysis was robust against the choice of the base vector as long as it captured the dimension along which trajectories differed (for example, the base vector cannot be orthogonal to the difference vectors).

Confidence intervals for the various hypotheses tested via the KiNeT analysis were computed by resampling the data 100 times. To assess the significance of our distance metric, we compared the bootstrap distribution of δ^{t_s} for $t_s = 700$ and 900 ms, the two values of t_s that were held out of the distance metric. Significance was computed using a t -test.

Comparison of KiNeT results to behavior. The time it takes for the state variable to reach a point in $\mathbf{r}^{[t_s]}$ is inversely proportional to the speed with which it changes with time. Using this relationship and denoting the speed of the reference and non-reference trajectories by $\mathbf{V}^{[800]}$ and $\mathbf{V}^{[t_s]}$, we approximated the relative speed, $\mathbf{V}^{[t_s]} / \mathbf{V}^{[800]}$ by $\mathbf{t}^{[800]} / \mathbf{t}^{[t_s]}$; that is, the slope of the regression line relating $\mathbf{t}^{[t_s]}$ to $\mathbf{t}^{[800]}$ that passes through the origin.

We also asked whether $\mathbf{V}^{[t_s]}$ used by the simulator was governed by the experimentally controlled t_s or the EKF estimate, t_s , inferred from the behavior of the animal (Fig. 8b). These two possibilities can be written in terms of the following two hypotheses, respectively:

$$H_1 : \mathbf{V}^{[t_s]} t_s \approx \mathbf{V}^{[800]} 800 \quad (16)$$

$$H_2 : \mathbf{V}^{[t_s]} t_s \approx \mathbf{V}^{[800]} 800 \quad (17)$$

Rearranging and exploiting the inverse relationship between speed and time, we can write:

$$H_1 : \hat{t}_e = 800 t^{[t_s]} / t^{[800]} \quad (18)$$

$$H_1 : t_s = 800 t^{[t_s]} / t^{[800]} \quad (19)$$

We evaluated these hypotheses by comparing the root mean squared error (RMSE) under the two hypotheses.

We asked whether the relative position of neural trajectories was organized in neural state space according to the experimentally controlled t_s or the EKF estimate, t_e , inferred from the animal's behavior (Fig. 8c). Using a linear model, we asked whether the relative position was more directly associated with t_e or t_s as follows:

$$H_1 : \hat{t}_e = \alpha \langle \delta^t \rangle t^{[800]} + \xi \quad (20)$$

$$H_2 : t_s = \alpha \langle \delta^t \rangle t^{[800]} + \xi \quad (21)$$

where $\langle \delta^t \rangle t^{[800]}$ is the average of δ^t over all $t^{[800]}$. We evaluated these hypotheses by fitting the parameters α and ξ under the two hypotheses and compared the goodness of fit using the RMSE.

To determine whether the second interval was integrated into the speed of the population response, we compared the overall bias in \hat{t}_e estimated from speeds during S2-S3 and S3-Go. We measured the overall bias as follows:

$$\text{Bias} = \sqrt{\frac{1}{K} \sum_{i=1}^K (\hat{t}_{e_i} - t_{s_i})^2} \quad (22)$$

with i indexing the sample interval.

We also used the average distance between trajectories to test whether the animals integrated the second measurement. The sigmoidal bias observed in behavior predicts that the inter-trajectory distances should exhibit a compressive nonlinearity; that is, distances between consecutive trajectories should become smaller for t_e values that are farther away from the middle values of 800 ms. We inferred the bias in each epoch from the average inter-trajectory distances using the following formulation:

$$\text{Bias} = 1 - \frac{(\langle \delta^{600} \rangle - \langle \delta^{700} \rangle) + (\langle \delta^{900} \rangle - \langle \delta^{1,000} \rangle)}{\langle \delta^{700} \rangle - \langle \delta^{900} \rangle} \quad (23)$$

In this formulation, positive values reflect the presence of the aforementioned compressive nonlinearity. Since the integration of the second measurement should reduce this nonlinearity, we asked whether the Bias computed using equation (23) was smaller in the S3-Go epoch relative to the S2-S3 epoch.

Statistics. The significance of the differences in slope from zero or unity of the linear model fit to animal behavior was assessed using a two-tailed t -test (Fig. 4e).

We tested the increase in variance with t_p by resampling the data 100 times with replacement to generate an estimate of the variance in each associated with each t_s . We then fit a linear model to the data and tested the significance of the difference from zero with a two-tailed t -test. The significance of the Bias in behavior was determined by resampling to generate 100 estimates of the statistic followed by a two-tailed t -test. Determination of the significance of differences in response to different conflict conditions was determined using one-tailed t -tests (Fig. 2f). Significance of the difference in the degree of scaling of individual neuron responses between task epochs was assessed using one-tailed Mann–Whitney–Wilcoxon test for equal medians (Fig. 4d). For population analyses (Figs. 7 and 8), we tested statistical significance by resampling trials to generate 100 bootstrap estimates of the statistic of interest. A t -statistic, estimated from the mean and standard deviation of the bootstrap distribution, was then used to perform either one- or two-tailed sample tests depending the assumptions associated with the test of interest. For all t -tests, assumptions of normality were not formally tested.

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

Data availability

The data that support the findings of this study are available at: https://drive.google.com/drive/folders/1T-U4hHW8iIEEea8ngBHryr9qfL_GNu_y?usp=sharing.

Code availability

Standalone code, including custom analyses used for generating the plots in the paper, are provided at: https://drive.google.com/drive/folders/1T-U4hHW8iIEEea8ngBHryr9qfL_GNu_y?usp=sharing.

References

- Jazayeri, M. & Shadlen, M. N. Temporal context calibrates interval timing. *Nat. Neurosci.* **13**, 1020–1026 (2010).
- Stengel, R. F. *Optimal Control and Estimation* (Dover Publications, 1994).
- Schlag, J. & Schlag-Rey, M. Evidence for a supplementary eye field. *J. Neurophysiol.* **57**, 179–200 (1987).
- Huerta, M. F. & Kaas, J. H. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J. Comp. Neurol.* **293**, 299–330 (1990).
- Fujii, N., Mushiaki, H. & Tanji, J. Distribution of eye- and arm-movement-related neuronal activity in the SEF and in the SMA and Pre-SMA of monkeys. *J. Neurophysiol.* **87**, 2158–2166 (2002).
- Matsuzaka, Y., Aizawa, H. & Tanji, J. A motor area rostral to the supplementary motor area (presupplementary motor area) in the monkey: neuronal activity during a learned motor task. *J. Neurophysiol.* **68**, 653–662 (1992).
- Yu, B. M. et al. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).

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

Spike waveforms were collected and visualized on-line using Blackrock Microsystems Neural Signal Processor (Firmware v6.5.4) and Central Suite v6.5.4. Behavioral data was collected using MWorks (<https://mworks.github.io/>).

Data analysis

Spiking data was extracted in MATLAB (R2014b, R2016b, R2017a, R2017b) using Blackrock Microsystems Neural Processing MATLAB Kit. Additional analyses used custom MATLAB (R2014b, R2016b, R2017a, R2017b) code.

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 reasonable 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

Life sciences study design

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

Sample size	We analyzed 181 neurons (118 in one animal and 63 in the second animal). No statistical methods were used to predetermine sample size.
Data exclusions	We excluded neurons with poor signal to noise and those that were not modulated by the task using pre-established criteria.
Replication	Both animals were trained over a period of 2 years and performed thousands of trials. Results were highly consistent across behavioral sessions. Neural results were robust to resampling and relaxation of exclusion criteria.
Randomization	Not applicable since no experimental groups were assigned. Trials and conditions per animal were randomized.
Blinding	Not applicable since no experimental groups were assigned. When warranted, data analyses were subjected to blinding using cross-validation.

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	Included in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input type="checkbox"/>	<input checked="" type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

Methods

n/a	Included in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input type="checkbox"/>	<input checked="" type="checkbox"/> MRI-based neuroimaging

Animals and other organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research

Laboratory animals	Two adult (8-9 years old) male rhesus macaques (<i>Macaca mulatta</i>), weighing 8.0 and 10 kg.
Wild animals	Not applicable. Study did not involve the use of wild animals.
Field-collected samples	Not applicable. Study did not involve the use samples collected from the field.
Ethics oversight	All experimental procedures were approved by the Committee of Animal Care at Massachusetts Institute of Technology.

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

Magnetic resonance imaging

Experimental design

Design type	N/A
Design specifications	NA
Behavioral performance measures	NA

Acquisition

Imaging type(s)	Structural
Field strength	3T
Sequence & imaging parameters	Structural MRIs were acquired with a 0.5 mm, isotropic MP-RAGE (Magnetization Prepared RApid Gradient Echo; sagittal slices; 448 mm field of view; 224x224 matrix size; TE: 2.69ms; TR: 2100ms; 8 deg flip angle).
Area of acquisition	Whole brain scan.
Diffusion MRI	<input type="checkbox"/> Used <input checked="" type="checkbox"/> Not used

Preprocessing

Preprocessing software	OsiriX Lite 10.0.1
Normalization	N/A
Normalization template	N/A
Noise and artifact removal	N/A
Volume censoring	N/A

Statistical modeling & inference

Model type and settings	N/A
Effect(s) tested	N/A
Specify type of analysis:	<input type="checkbox"/> Whole brain <input type="checkbox"/> ROI-based <input type="checkbox"/> Both
Statistic type for inference (See Eklund et al. 2016)	N/A
Correction	N/A

Models & analysis

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Functional and/or effective connectivity
<input checked="" type="checkbox"/>	<input type="checkbox"/> Graph analysis
<input checked="" type="checkbox"/>	<input type="checkbox"/> Multivariate modeling or predictive analysis