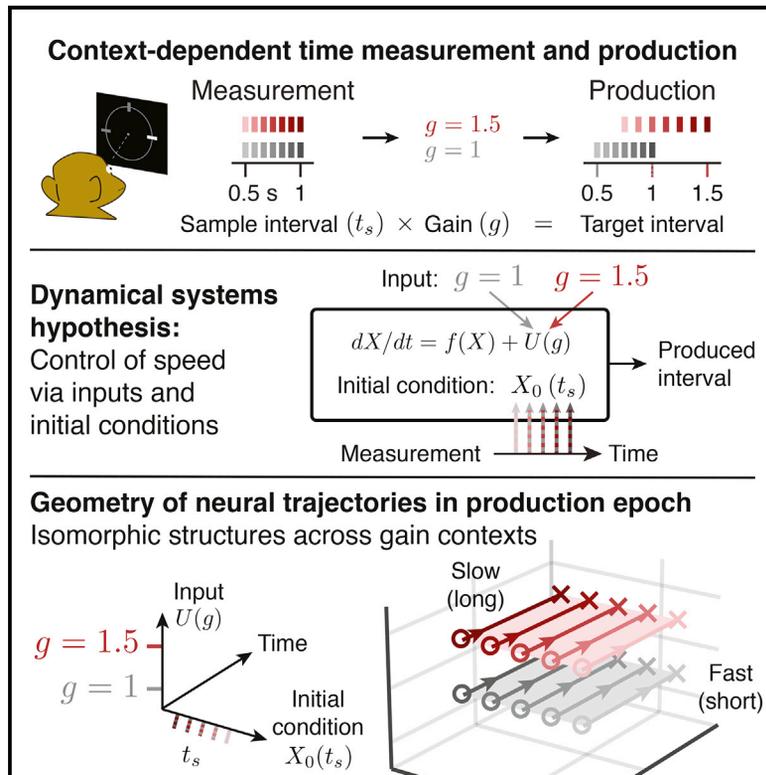


Neuron

Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics

Graphical Abstract



Authors

Evan D. Remington, Devika Narain, Eghbal A. Hosseini, Mehrdad Jazayeri

Correspondence

mjaz@mit.edu

In Brief

Remington et al. employ a dynamical systems perspective to understand how the brain flexibly controls timed movements. Results suggest that neurons in the frontal cortex form a recurrent network whose behavior is flexibly controlled by inputs and initial conditions.

Highlights

- Monkeys performed a timing task demanding flexible cognitive control
- The organization of neural trajectories in frontal cortex reflected task demands
- Flexible control was best explained in terms of inputs and initial conditions
- Recurrent neural network models validated the inferred control principles



Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics

Evan D. Remington,¹ Devika Narain,^{1,2,3,4} Eghbal A. Hosseini,² and Mehrdad Jazayeri^{1,2,5,*}

¹McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, USA

²Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

³Netherlands Institute for Neuroscience, Amsterdam, the Netherlands

⁴Erasmus Medical Center, Rotterdam, the Netherlands

⁵Lead Contact

*Correspondence: mjaz@mit.edu

<https://doi.org/10.1016/j.neuron.2018.05.020>

SUMMARY

Neural mechanisms that support flexible sensorimotor computations are not well understood. In a dynamical system whose state is determined by interactions among neurons, computations can be rapidly reconfigured by controlling the system's inputs and initial conditions. To investigate whether the brain employs such control mechanisms, we recorded from the dorsomedial frontal cortex of monkeys trained to measure and produce time intervals in two sensorimotor contexts. The geometry of neural trajectories during the production epoch was consistent with a mechanism wherein the measured interval and sensorimotor context exerted control over cortical dynamics by adjusting the system's initial condition and input, respectively. These adjustments, in turn, set the speed at which activity evolved in the production epoch, allowing the animal to flexibly produce different time intervals. These results provide evidence that the language of dynamical systems can be used to parsimoniously link brain activity to sensorimotor computations.

INTRODUCTION

Humans and nonhuman primates are capable of generating a vast array of behaviors, a feat dependent on the brain's ability to produce a vast repertoire of neural activity patterns. However, identifying the mechanisms by which the brain flexibly selects neural activity patterns across a multitude of contexts remains a fundamental and outstanding problem in systems neuroscience.

Here, we aimed to answer this question using a dynamical systems approach. Work in multiple modalities has provided support for a hypothesis that neural activity can be understood at the level of neural populations and viewed as neural trajectories of a dynamical system (Fetz, 1992; Rabinovich et al., 2008; Buonomano and Maass, 2009; Shenoy et al., 2013). For example, recent studies have provided compelling evi-

dence that low-dimensional activity in the motor cortex can be largely explained by inherent dynamical interactions (Churchland et al., 2010, 2012; Michaels et al., 2016; Seely et al., 2016). A dynamical systems view has also been used to provide explanations for neural trajectories in premotor and prefrontal cortical areas in various cognitive tasks (Rigotti et al., 2010; Mante et al., 2013; Hennequin et al., 2014; Carnevale et al., 2015; Rajan et al., 2016; Wang et al., 2018). This line of investigation has been complemented by efforts in developing, training, and analyzing recurrent neural network models that can emulate a range of motor and cognitive behaviors, leading to novel insights into the underlying latent dynamics (Rigotti et al., 2010; Laje and Buonomano, 2013; Mante et al., 2013; Hennequin et al., 2014; Sussillo et al., 2015; Chaisangmongkon et al., 2017; Wang et al., 2018). These early successes hold promise for the development of a more ambitious "computation-through-dynamics" that provides a general framework for understanding how activity patterns in the brain support flexible behaviorally relevant computations.

The behavior of a dynamical system can be described in terms of three factors: (1) the interaction between state variables, (2) the system's initial state, and (3) the external inputs to the system. Accordingly, the hope for using the mathematics of dynamical systems to understand flexible generation of neural activity patterns and behavior depends on our ability to understand the co-evolution of behavioral and neural states in terms of these three components. Assuming that synaptic couplings between neurons and other biophysical properties are approximately constant on short timescales (i.e., trial to trial), we asked whether behavioral flexibility can be understood in terms of adjustments to initial states and low-dimensional external inputs with minimal dynamics.

There is evidence that certain aspects of behavioral flexibility can be understood in terms of these factors. For example, it has been proposed that preparatory activity prior to movement initializes the system such that ensuing movement-related activity follows the appropriate trajectory (Churchland et al., 2010). Similarly, the presence of a static context input can enable a recurrent neural network to perform flexible rule- (Mante et al., 2013; Song et al., 2016; Wang et al., 2018) and category-based decisions (Chaisangmongkon et al., 2017). However, whether these initial insights would apply more



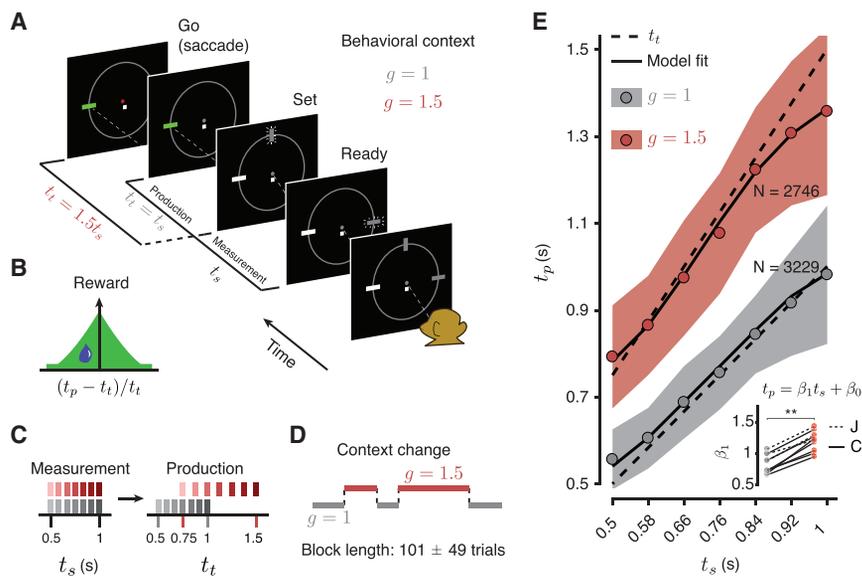


Figure 1. The RSG Task and Behavior

(A) RSG task. On each trial, three rectangular stimuli associated with “Ready,” “Set,” and “Go” events were shown on the screen arranged in a semi-circle. Following fixation, Ready and Set cues were extinguished. After a random delay, first Ready and then Set stimuli were flashed (small lines around the rectangles signify flashed stimuli). The time interval between Ready and Set demarcated a sample interval, t_s . The monkey’s task was to generate a saccade (“Go”) to a visual target such that the interval between Set and Go (produced interval, t_p) was equal to a target interval, t_t , equal to t_s multiplied by a gain factor, g ($t_t = gt_s$). The animal had to perform the task in two behavioral contexts, one in which t_t was equal to t_s ($g=1$ context), and one in which t_t was 50% longer than t_s ($g=1.5$ context). The context was cued throughout the trial by the color of fixation and the position of a context stimulus (small white square below the fixation).

(B) Reward. Animals received juice reward when the error between t_p and t_t was small, and

the reward magnitude decreased with the size of error (see STAR Methods for details). On rewarded trials, the saccadic target turned green (A). (C) Sample and target intervals. For both contexts, t_s was drawn from a discrete uniform distribution with seven values equally spaced from 0.5 to 1 s (left). The values of t_s were chosen such that the corresponding values of t_t across the two contexts were different but partially overlapping (right).

(D) Context blocks. The context changed across blocks of trials. The number of trials in a block was varied pseudorandomly (mean and SD shown).

(E) Behavior. t_p as a function of t_s for each context across both monkeys and all recording sessions. Circles indicate mean t_p across all sessions, shaded regions indicate \pm one standard deviation from the mean, dashed lines indicate t_t , and solid lines are the fits of a Bayesian observer model to behavior. Inset: slope of the regression line (β_1) relating t_p to t_s in the two contexts. Regression slopes were larger in the $g=1.5$ context, with a significant interaction between t_s and g ($p < 0.0001$) for all sessions (see text; ** $p < 0.002$ for signed-rank test). Regression intercepts (β_0) were also larger for the $g=1.5$ context (0.14 versus 0.18 s, $p = 0.04$). In all panels, different shades of gray and red are associated with $g=1$ and $g=1.5$, respectively. See Figure S1 for individual animals.

broadly when both inputs and initial conditions change is an important outstanding question.

For many behaviors, distinguishing the effects of synaptic coupling, inputs, and initial conditions in neural activity patterns is challenging. For example, neural activity during a reaching movement is likely governed by both local recurrent interactions and distal inputs from time-varying and condition-dependent reafferent signals (Todorov and Jordan, 2002; Scott, 2004; Pruszynski et al., 2011). Similarly, in many perceptual decision making tasks, it is not straightforward to disambiguate the sensory drive from recurrent activity representing the formation of a decision and the subsequent motor plan (Mante et al., 2013; Meister et al., 2013; Thura and Cisek, 2014). This makes it difficult to tease apart the contribution of recurrent dynamics governed by initial conditions from the contribution of dynamic inputs (Sussillo et al., 2016). To address this challenge, we designed a sensorimotor task for nonhuman primates in which animals had to measure and produce time intervals using internally generated patterns of neural activity in the absence of potentially confounding time-varying sensory and reafferent inputs. Using a novel analysis of the geometry and dynamics of *in vivo* activity in the dorsal medial frontal cortex (DMFC) and *in silico* activity in recurrent neural network (RNN) models trained to perform the same task, we found evidence that behavioral flexibility is mediated by the complementary action of inputs and initial conditions controlling the structural organization of neural trajectories.

RESULTS

Ready, Set, Go Task

Our aim was to ask whether flexible control of internally generated dynamics could be understood in terms of systematic adjustments made to initial conditions and external inputs of a dynamical system. We designed a “Ready, Set, Go” (RSG) timing task to directly investigate the role of these two factors. The basic sensory and motor events in the task were as follows: after fixating a central spot, monkeys viewed two peripheral visual flashes (“Ready” followed by “Set”) separated by a sample interval, t_s , and produced an interval, t_p , after Set by making a saccade to a visual target that was presented throughout the trial. In order to obtain juice reward, animals had to generate t_p as close as possible to a target interval, t_t (Figure 1B), which was equal to t_s times a “gain factor,” g ($t_t = gt_s$). The demand for flexibility was imposed in two ways (Figure 1C). First, t_s varied between 0.5 and 1 s on a trial-by-trial basis (drawn from a discrete uniform “prior” distribution). Second, g switched between 1 ($g=1$ context) and 1.5 ($g=1.5$ context) across blocks of trials (Figure 1D, mean block length = 101, SD = 49 trials).

To verify that animals learned the task (Figures 1E and S1), we used regression analyses to assess the dependence of t_p on t_s and g . First, we verified that animals learned to measure t_s by analyzing the relationship between t_s and t_p within each context ($t_p = \beta_0 + \beta_1 t_s$). Results indicated that t_p increased

monotonically with t_s ($\beta_1 > 0$, $p \ll 0.001$ for all sessions). Next, we verified that animals learned the gain by comparing regression slopes relating t_p to t_s across the two contexts. The slopes (β_1) were significantly higher in the $g = 1.5$ context compared to $g = 1$ context (mean $\beta_1 = 1.2$ versus 0.84; signed-rank test $p = 0.002$, $n = 10$ sessions; Figure 1E, inset). As a complementary analysis to show that animals displayed sensitivity to gain in individual sessions, we fit a regression model to behavior across both contexts that included additional regressors for gain and its interaction with t_s ($t_p = \beta_0 + \beta_1 t_s + \beta_2 g + \beta_3 g t_s$). Results indicated a significant positive interaction between t_s and g (mean $\beta_3 = 0.73$; $\beta_3 > 0$, $p < 0.0001$ in each session). Finally, we asked whether animals switched between contexts rapidly by fitting a regression model relating t_p to Z scored for each t_s (z_p), to the number of trials n following a context switch ($z_p = \beta_0 + \beta_1 n$, $1 \leq n \leq 25$ trials after switches). There was no evidence that z_p changed as a function of number of trials after switch (one-tailed test for β_1 , $p > 0.25$), suggesting that the switching was rapid. Together, these results confirmed that animals used an estimate of t_s to compute t_p and were able to flexibly and rapidly adjust their responses according to the gain information.

For both gains, responses were variable, and average responses exhibited a regression to the mean (mean $\beta_1 < 1$, $p = 0.005$ for $g = 1$, and mean $\beta_1 < 1.5$, $p = 0.0001$ for $g = 1.5$, one-sided signed-rank test). As with previous work (Miyazaki et al., 2005; Jazayeri and Shadlen, 2010, 2015; Acerbi et al., 2012), these behavioral characteristics were accurately captured by a Bayesian model (Figure 1E; STAR Methods) indicating that animals integrated their knowledge about the prior distribution, the sample interval, and the gain to optimize their behavior.

Neural Activity in the RSG Task

To assess the neural computations in RSG, we focused on the dorsal region of the medial frontal cortex (DMFC) comprising supplementary eye field, dorsal supplementary motor area (i.e., excluding the medial bank), and presupplementary motor area. DMFC is a natural candidate for our task because it plays a crucial role in timing as shown by numerous studies in humans (Halsband et al., 1993; Rao et al., 2001; Coull et al., 2004; Pfeuty et al., 2005; Macar et al., 2006; Cui et al., 2009), monkeys (Okano and Tanji, 1987; Kurata and Wise, 1988; Romo and Schultz, 1992; Isoda and Tanji, 2003; Ohmae et al., 2008; Mita et al., 2009; Merchant et al., 2011, 2013; Kunimatsu and Tanaka, 2012), and rodents (Matell et al., 2003; Kim et al., 2009, 2013; Smith et al., 2010; Murakami et al., 2014; Xu et al., 2014), and because it is involved in context-specific control of actions (Matsuzaka and Tanji, 1996; Shima et al., 1996; Brass and von Cramon, 2002; Isoda and Hikosaka, 2007; Yang and Heinen, 2014; Ray and Heinen, 2015).

We recorded from 324 units (129 from monkey C and 195 from monkey J) in DMFC (Figures S2 and S11). Between 11 and 82 units were recorded simultaneously in a given session; however, in this study, we combined data across all units irrespective of whether they were recorded simultaneously. Firing patterns were heterogeneous and varied across units, task epochs, and experimental contexts. In the Ready-Set epoch,

responses were modulated by both gain and elapsed time (e.g., units #1, 3, and 5, Figure 2A). For many units, firing rate modulations underwent a salient change at the earliest expected time of Set (0.5 s). For example, responses of some units increased monotonically in the first 0.5 s but decreased afterward (Figure 2A, units #1 and 3).

Following Set, firing rates were characterized by a mixture of (1) transient changes after Set (units #1 and 3), (2) sustained modulations during the Set-Go epoch (units #1 and 5), and (3) monotonic changes in anticipation of the saccade (units #1, 2 and 4). These characteristics were not purely sensory or motor and varied systematically with t_s and gain. For example, the amplitude of the early transient response (unit #1) depended on both t_s and gain, indicating that it was not a visually triggered response to Set. The same was true for the sustained modulations after Set and activity modulations prior to saccade initiation.

We also examined the representation of t_s and gain across the population by projecting the data on dimensions along which activity was strongly modulated by context and interval in state-space (i.e., the space spanned by the firing rates of all 324 units; see STAR Methods). Similar to individual units, population activity was modulated by both elapsed time and gain during both the Ready-Set (Figure 2B) and Set-Go (Figure 2C) epochs. We used this rich dataset to investigate whether the flexible adjustment of intrinsic dynamics across the population with respect to t_s and gain could be understood using the language of dynamical systems.

Flexible Neural Computations: A Dynamical Systems Perspective

We pursued the idea that neural computations responsible for flexible control of saccade initiation time can be understood in terms of the behavior of a dynamical system established by interactions among neurons. To formulate a rigorous hypothesis for how a dynamical system could confer such flexibility, we considered the goal of the task and worked backward logically. The goal of the animal was to flexibly control the saccade initiation time to a fixed target. Previous motor timing studies proposed that saccade initiation is triggered when the activity of a subpopulation of neurons with monotonically increasing firing rates (i.e., “ramping”) reaches a threshold (Romo and Schultz, 1987; Hanes and Schall, 1996; Roitman and Shadlen, 2002; Tanaka, 2005; Maimon and Assad, 2006; Mita et al., 2009; Kunimatsu and Tanaka, 2012). For these neurons, flexibility requires that the slope of the ramping activity be adjusted (Jazayeri and Shadlen, 2015). More recently, it was found that actions are initiated when the collective activity of neurons with both ramping and more complex activity patterns reach an action-triggering state (Churchland et al., 2006; Wang et al., 2018) and that flexible control of initiation time can be understood in terms of the speed with which neural activity evolves toward that terminal state (Shinomoto et al., 2011; Wang et al., 2018).

In a dynamical system, the speed with which activity evolves over time is determined by the derivative of the state. If we denote the state of the system by X , the derivative is usually specified by two factors, a function of the current state, $f(X)$,

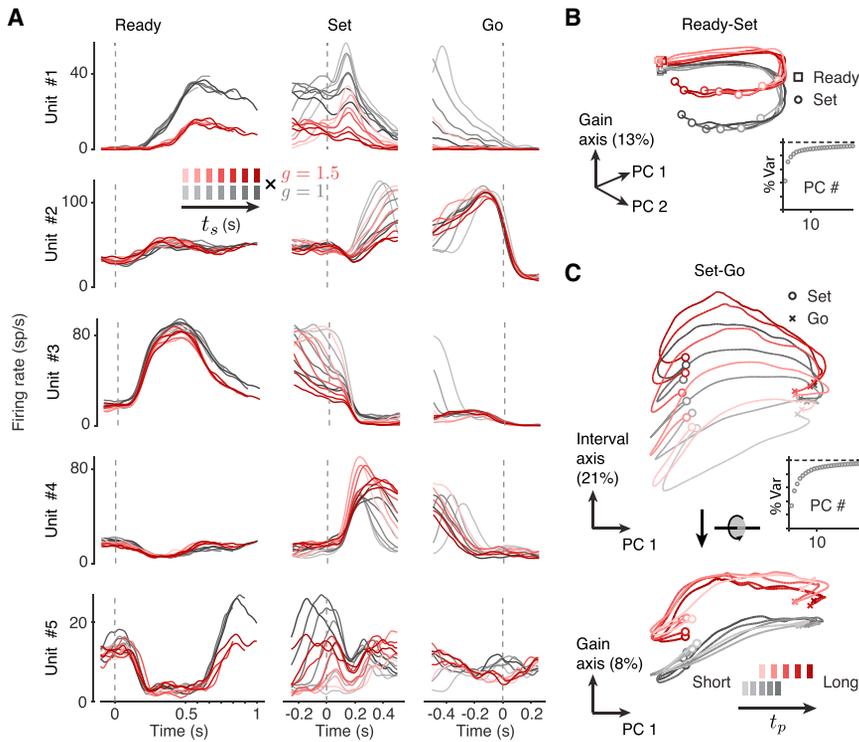


Figure 2. Neural Responses in Dorsomedial Frontal Cortex during the RSG Task

(A) Example neurons. Firing rates of 5 example units during the various phases of the task aligned to Ready (left column), Set (middle), and Go (right). Responses aligned to Ready and Set were sorted by t_s . Responses aligned to Go were sorted into 5 bins, each with the same number of trials, ordered by t_p . Gray and red lines correspond to activity during the $g=1$ and $g=1.5$ contexts, respectively, with darker lines corresponding to longer intervals.

(B) Population activity during Ready-Set. Visualization of population activity in the Ready-Set epoch sorted by t_s . The “gain axis” corresponds to the axis along which responses were maximally separated with respect to context at the time of Set. The other two dimensions (“PC1 and PC2”) correspond to the first two principal components of the data after removing the context dimension. Inset: fraction of variance explained by first 25 principal components. Dashed line indicates 100%.

(C) Population activity during Set-Go. Visualization of population activity in the Set-Go epoch sorted into 5 bins, each with the same number of trials, ordered by t_p . Top: activity plotted in 2 dimensions spanned by PC1 and the dimension of maximum variance with respect to t_p within each context (“Interval axis”). Bottom: same as top rotated 90°

(circular arrow) to visualize activity in the plane spanned by the context axis (“Gain axis”) and PC1. For plotting purposes, PCs were orthogonalized with respect to the interval and gain axes. Squares, circles, and crosses in the state space plots represent Ready, Set, and Go, respectively. Percentage variance explained by interval and gain are shown numerically near the corresponding axes. See STAR Methods for the calculation of percent variance explained by interval and gain. See Figure S2 for individual animals and Figure S11 for different recording sites.

and an external input, U , that may be constant or context and/or time dependent:

$$\frac{dX}{dt} = f(X) + U.$$

When analyzing the collective activity of a specific population of neurons, this formulation has a straightforward interpretation. The state represents the collective firing rate of neurons under investigation, $f(X)$ accounts for the interactions among those neurons, and U corresponds to external inputs coming from other populations of neurons, possibly driven by external stimuli. The only additional information needed to determine the behavior of this system is its initial condition, X_0 , which specifies the initial neural state prior to generating a desired dynamic pattern of activity.

To assess the utility of the dynamical systems perspective for understanding behavioral flexibility, we assumed that $f(X)$ (i.e., synaptic coupling in DMFC) is fixed across trials. Furthermore, we only considered external inputs U of low dimensionality and low temporal complexity (i.e., tonic and transient inputs), as a complex input could trivially account for the observed dynamics without providing any meaningful insight regarding the observed neural activity. This leaves initial conditions and low-dimensional inputs as the only “dials” for achieving flexibility with respect to t_s and gain (Figure 3).

First, we focused on behavioral flexibility with respect to t_s for each gain context. How can a dynamical system adjust the

speed at which activity during Set-Go evolves in a t_s -dependent manner? In RSG, within each context, there are no sensory inputs (exafferent or reafferent) that could serve as a t_s -dependent input drive. Therefore, we hypothesized that the t_s -dependent adjustment of speed in the Set-Go epoch results from a parametric control of initial conditions at the time of Set. Because the speed of the dynamics at any given point in time is governed by the current state of the system, if the initial conditions at Set determine the path that neural activity takes through the state space, then it follows that initial conditions control speed. The corollary to this hypothesis is that the time-varying activity during the Ready-Set epoch is responsible for adjusting this initial condition based on the desired speed during the ensuing Set-Go epoch.

Second, we asked how speed might be controlled across the two gain contexts. One possibility is that the system uses initial conditions to encode $t_t = gt_s$, which has implicit information about both gain and t_s (i.e., $X_0(gt_s)$). This would predict that neural trajectories form a single organized structure with respect to the target time ($t_t = gt_s$). In the extreme case, neural trajectories associated with the same value of gt_s across the two contexts (e.g., 1.5×0.5 and 1.0×0.75) should terminate in the same state at the time of Set and should evolve along identical trajectories during the Set-Go epoch. We refer to this solution as A_1 (Figure 3A).

Alternatively, a tonic gain-dependent input could be responsible for rapid reconfiguration of dynamics across contexts

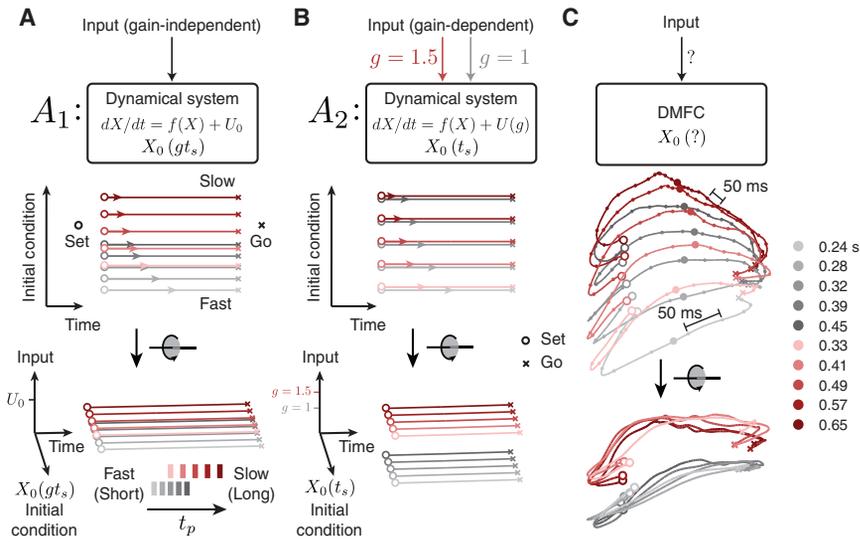


Figure 3. Dynamical Systems Predictions for the RSG Task

(A and B) Schematic illustrations of two hypothetical dynamical systems solutions to RSG across contexts through manipulation of initial conditions or external inputs. (A) The first solution (A_1). Top: in A_1 , the initial condition (X_0) depends on the target interval $t_t = gt_s$ (g , gain, t_s , sample interval), and the system is driven by a gain-independent input (U_0). Middle: state trajectory between Set and Go in the plane spanned by initial condition and time axes. After Set (open circles), activity evolves toward an action-triggering state (crosses) with a speed (colored arrows) fully determined by position along the initial condition axis (ordinate). Activity across contexts is organized according to $t_t = gt_s$. Bottom: same trajectories, rotated to show an oblique view. Trajectories are separated only along the initial condition axis across both contexts such that trajectory structure reflects t_t explicitly. There is no separation along the Input axis. (B) The second solution (A_2). Top: in A_2 , X_0 depends on t_s within contexts,

but the system is driven by a tonic gain-dependent input ($U(g)$; red and gray arrows for the two gains). Middle: state trajectory between Set and Go in the plane spanned by initial condition and time axes. Because initial condition encodes t_s and not t_t , in this plane, trajectories associated with the same t_s but different gains appear overlapping. Bottom: oblique view. A context-dependent external input creates two sets of neural trajectories in the state space for the two contexts in the Set-Go epoch. This input controls speed in conjunction with initial conditions, generating a structure which reflects t_s and g explicitly, but not t_t . In both A_1 and A_2 , responses would be initiated when activity projected onto the time axis reaches a threshold.

(C) DMFC data. Top: unknown mechanism of RSG control in DMFC. Middle, bottom: three-dimensional projection of DMFC activity in the Set-Go epoch (from Figure 2C). Middle: qualitative assessment indicated that neural trajectories within each context for different t_p bins were associated with different initial conditions and remained separate and ordered through the response. Bottom: across the two contexts, neural trajectories formed two separated sets of neural trajectories without altering their relative organization as a function of t_p . Both of these features were consistent with A_2 . Filled circles depict states along each trajectory at a constant fraction of the trajectory length, illustrating speed differences across trajectories (circles are closer for the longer intervals associated with slower speed).

(Figure 3B). As exemplified by previous work, a tonic input inevitably pushes activity in recurrent neural network models to a different region of the state space (Mante et al., 2013; Hennequin et al., 2014; Sussillo et al., 2015; Song et al., 2016; Chaisangmongkon et al., 2017; Wang et al., 2018). This solution, which we refer to as A_2 , makes the following predictions. First, it predicts that the two gains should lead to two sets of neural trajectories in two different regions of the state space. Second, it predicts that neural trajectories should be organized with respect to t_s and t_p (i.e., within each context) but not necessarily with respect to t_t (i.e., across contexts). Because the context information in RSG was provided as an external visual input (fixation cue) and was available throughout the trial, we predict that this solution offers the more plausible prediction for how the brain might solve the task.

We thus set out to characterize the geometry of the observed neural trajectories quantitatively to determine the most likely explanation for their structure. The dynamical systems perspective in RSG leads us to the following specific hypotheses: (1) the evolution of activity in the Ready-Set epoch parametrizes the initial conditions needed to control the speed of dynamics in the production epoch within each context, and (2) the context cue acts as a tonic external input leading the system to establish structurally similar yet distinct sets of neural trajectories associated with the two gains, consistent with A_2 .

Visualization of neural trajectories from Set to Go in state space (Figure 3C, same as in Figure 2C) provided qualitative support for these hypotheses. First, within each context, neural tra-

jectories for different t_p bins were clearly associated with different initial conditions and remained separate and ordered throughout the Set-Go epoch. Second, context information seemed to displace the entire group of neural trajectories to a different region of neural state space without altering their relative organization as a function of t_p . Third, indexing time along nearby trajectories suggested that the speed with which responses evolved along each trajectory was systematically related to the desired t_t ; i.e., slower for longer t_t . To validate these observations quantitatively, we developed an analysis technique that we termed “kinematic analysis of neural trajectories” (KiNeT) that helped us measure the relative speed and position of multiple, possibly curved (Figure S3), neural trajectories.

Control of Neural Trajectories by Initial Condition within Contexts

We first employed KiNeT to validate that animals’ behavior was predicted by the speed with which neural trajectories evolved over time. We reasoned that neural states evolving faster will reach the same destination on the trajectory in a shorter amount of time. Therefore, we estimated relative speed across the trajectories by performing a time alignment to identify the times when neural activity reached nearby points on each trajectory (Figure 4A). We then used this approach to analyze the geometrical structure of trajectories through the Set-Go epoch.

To perform KiNeT, we binned trials from each gain and recording session into five groups according to t_p . Neural responses from these trials were averaged, and then PCA was

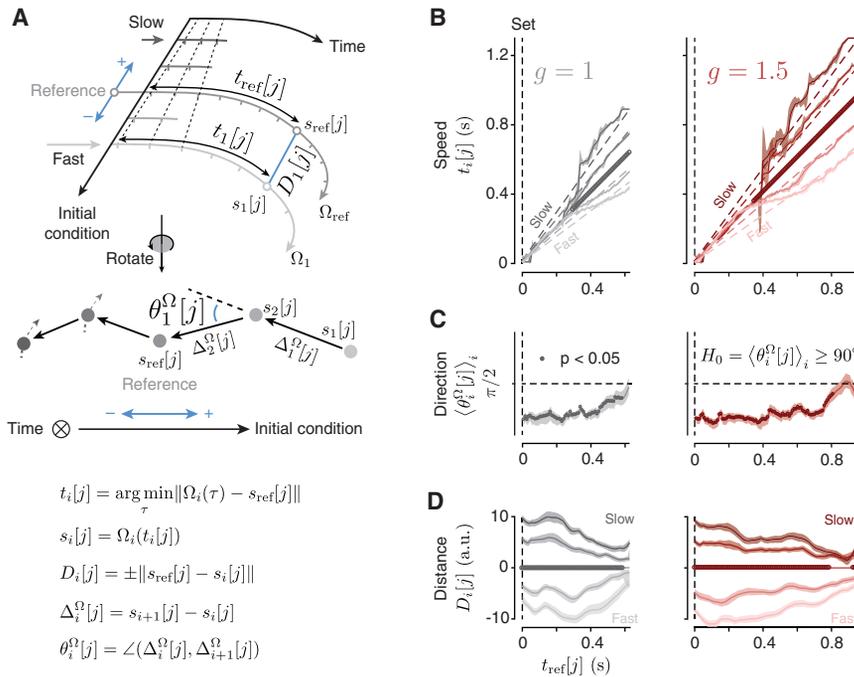


Figure 4. Kinematic Analysis of Neural Trajectories

(A) Illustration of kinematic analysis of neural trajectories KiNeT. Top: a collection of trajectories originate from Set, organized by initial condition. Tick marks on the trajectories indicate unit time. Darker trajectories evolve at a lower speed as demonstrated by the distance between tick marks and the dashed line connecting tick marks. KiNeT quantifies the position of trajectories and the speed with which states evolve along them relative to a reference trajectory (middle trajectory, Ω_{ref}). To do so, it finds a collection of states $\{s_i[j]\}_j$ on each Ω_i that are closest to Ω_{ref} through time (for visual clarity, only the first and reference trajectories are illustrated). Trajectories which evolve at a slower speed require more time to reach those states leading to larger values of $t_i[j]$. KiNeT quantifies relative position by a distance measure, $D_i[j]$ (distance between Ω_i and Ω_{ref} at $t_i[j]$) that is signed (blue arrows) and is considered positive for the largest value of t_p (see STAR Methods). Middle: trajectories rotated such that the time axis is normal to the plane of illustration, denoted by a circle with an inscribed cross. Filled circles represent the states $\{s_i[j]\}_j$ aligned to $s_{ref}[j]$ for a particular j . Vectors $\Delta_i^\Omega[j]$ connect states on trajectories of shorter to longer t_p . Angles $\theta_i^\Omega[j]$

between successive $\Delta_i^\Omega[j]$ provide a measure of t_p -related structure. Bottom: equations defining the relevant variables.

(B) Speed of neural trajectories compared to Ω_{ref} computed for each context separately. Shortly after Set, all trajectories evolved with similar speed (unity slope). Afterward, Ω_i associated with shorter t_s evolved faster than Ω_{ref} as indicated by a slope of less than unity (i.e., $\{t_i[j]\}_j$ smaller than $\{t_{ref}[j]\}_j$), and Ω_i associated with longer t_s evolved slower than Ω_{ref} . Circles on the unity line indicate j values for which $\{t_i[j]\}_j$ was significantly correlated with i (bootstrap test, $r > 0$, $p < 0.05$, $n = 100$).

(C) Relative position of adjacent neural trajectories computed for each context separately. $\langle \theta_i^\Omega[j] \rangle_i$ (angle brackets with subscript i signify average across trajectories) were significantly smaller than 90° (circles) for the majority of the Set-Go epoch (bootstrap test, $\langle \theta_i^\Omega[j] \rangle_i < 90^\circ$, $p < 0.05$, $n = 100$) indicating that $\Delta_i^\Omega[j]$ were similar across Ω_i .

(D) Distance of neural trajectories to Ω_{ref} computed for each context separately. Distance measures ($D_i[j]$) indicated that $\{\Omega_i\}_j$ had the same ordering as the corresponding t_p values. The magnitude of $D_i[j]$ decreased with time, indicating that trajectories coalesce as they get closer to the time of Go. When trajectories coalesce, small deviations in their relative position due to variability in firing rate estimates may cause trajectories to appear disorganized. This is consistent with the observation that $\langle \theta_i^\Omega[j] \rangle_i$ were closer to 90° near the time of Go (C). Shaded regions represent 90% bootstrap confidence intervals. Significance was tested using bootstrap samples for each j ($p < 0.05$, $n = 100$). See Figure S5 for individual animals and Figure S11 for different recording sites.

applied to generate five neural trajectories within the state space spanned by the first 10 PCs that explained 89% of variance. We denote each trajectory after the time of Set by Ω_i (see Table S1 for an abridged list of notations and symbols) where the subscript i indexes the trajectory. We use curly brackets to refer to a collection of trajectories ($\{\Omega_i\}_j$). We estimated speed and position along each Ω_i relative to the trajectory associated with the middle (third) bin, which we refer to as the reference trajectory, Ω_{ref} . We denoted neural states on the reference trajectory by $s_{ref}[j]$, where j indexes states through time along Ω_{ref} . We also use curly brackets to refer to a collection indices over times. For example, $\{s_{ref}[j]\}_j$ refers to all states on Ω_{ref} , and $\{t_{ref}[j]\}_j$ corresponds to time points on Ω_{ref} associated with those states.

For each $s_{ref}[j]$, we found the nearest point on all non-reference trajectories ($i \neq ref$) as measured by Euclidean distance. We denoted the collection of the nearest states on Ω_i to Ω_{ref} by $\{s_i[j]\}_j$, and the corresponding time points by $\{t_i[j]\}_j$. The corresponding time points along different trajectories provided the means for comparing speed: if $\{t_i[j]\}_j$ were systematically greater than $\{t_{ref}[j]\}_j$, we could conclude that Ω_i evolves at a slower speed compared to Ω_{ref} (Figure 4A). This relationship can be readily inferred from the slope of the line that relates

$\{t_i[j]\}_j$ to $\{t_{ref}[j]\}_j$. While a unity slope indicates that the speeds are the same, higher and lower values would indicate slower and faster speeds of Ω_i compared to Ω_{ref} , respectively.

Applying KiNeT to neural trajectories in the Set-Go epoch indicated that Ω_i evolved at similar speeds immediately following the Set cue (unity slope). Later, speed profiles diverged such that neural trajectories associated with longer intervals slowed down and trajectories associated with shorter intervals sped up for both gain contexts (Figure 4B). This is consistent with previous work that the key variable predicting t_p is the speed with which neural trajectories evolve (Wang et al., 2018). One common concern in this type of analysis is that averaging firing rates across trials of slightly different duration could lead to a biased estimate of neural trajectory. To ensure that our estimates of average speed were robust, we applied KiNeT to neural trajectories while aligning trials to Go instead of Set. Results remained unchanged and confirmed that the speed of neural trajectories predicted t_p across trials (Figure S4).

Having validated speed as the key variable for predicting t_p , we focused on our first hypothesis that the evolution of activity in the Ready-Set epoch parametrizes the initial conditions to control the speed of dynamics in the production epoch for

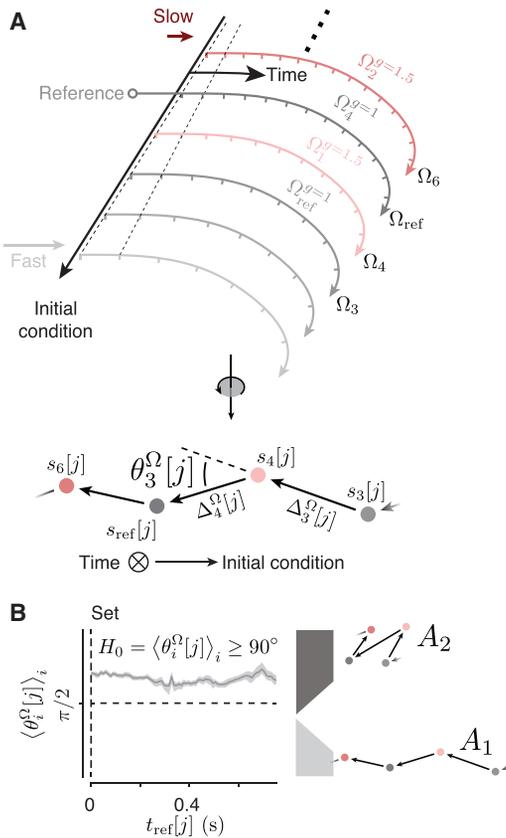


Figure 5. Neural Trajectories across Contexts Do Not Form a Single Structure Reflecting t_p

(A) A schematic illustrating neural trajectories across the two contexts after Set. Top: the expected geometrical structure under A_1 . Neural trajectories for the gain of 1 (gray) and 1.5 (red) are organized along a single initial condition axis and ordered with respect to t_p . Tick marks indicate unit time. Bottom: a rotation of the top showing neural trajectories with the time axis normal to the plane of illustration. If the neural trajectories were organized as such, then the angle between vectors connecting nearby points (e.g., $\theta_3^{\Omega}[j]$) would be less than 90° (A_1 , Figure 3A).

(B) Left: orientation of vectors connecting adjacent neural trajectories combined across the two contexts. Right: possible geometrical structures A_1 (bottom), and A_2 (top). $\langle \theta_i^{\Omega}[j] \rangle_i$ was larger than 90° for all j in the Set-Go interval, consistent with A_2 . Shaded regions represent 90% bootstrap confidence intervals. See Figure S6 for individual animals and Figure S11 for different recording sites.

each context. Because speed is a scalar variable and has an orderly relationship to t_p , this hypothesis predicts that the neural trajectories (and their initial conditions) should also have an orderly organizational structure with respect to t_p . In other words, there should be a systematic relationship between the vectors connecting nearest points across neural trajectories and the t_p to which they correspond. We tested this prediction in two complementary ways. First, we performed an *analysis of direction* testing whether the vectors connecting nearby trajectories were more aligned than expected by chance. Second, we performed an *analysis of distance* asking whether the distance between the reference trajectory and the other trajectories respected the distance between the corresponding speeds.

Analysis of Direction

We used KiNeT to measure the angle between vectors connecting nearest points (Euclidean distance) across consecutive trajectories ordered by t_p . We used $\Delta_i^{\Omega}[j]$ to denote the vector connecting nearest points across trajectories (superscript Ω) between $s_i[j]$ and $s_{i+1}[j]$. According to our hypothesis, the direction of $\Delta_i^{\Omega}[j]$ should be similar to $\Delta_{i+1}^{\Omega}[j]$ connecting $s_{i+1}[j]$ to $s_{i+2}[j]$. To test this, we measured the angle between these two vectors, denoted by $\theta_i^{\Omega}[j]$. The null hypothesis of unordered trajectories predicts that $\Delta_i^{\Omega}[j]$ and $\Delta_{i+1}^{\Omega}[j]$ should be unaligned on average ($\langle \theta_i^{\Omega}[j] \rangle_i = 90^\circ$; angle brackets with the subscript i signify average over index i). Results indicated that $\theta_i^{\Omega}[j]$ was substantially smaller than 90° for both contexts (Figures 4C and S5). This provides the first line of quantitative evidence for an orderly organization of neural trajectories with respect to t_p .

Analysis of Distance

We used KiNeT to measure the length of the vectors connecting nearest points on Ω_i and Ω_{ref} , denoted by $D_i[j]$, at different time points. This analysis revealed that trajectories evolving faster than Ω_{ref} and those evolving slower than Ω_{ref} were located on the opposite sides of Ω_{ref} and that the magnitude of $D_i[j]$ increased progressively for larger speed differences (Figures 4D and S5). This analysis provided evidence that, for each context, the relative position of neural trajectories and their initial conditions in the state space were orderly and predictive of t_p .

Control of Neural Trajectories across Contexts by External Input

To identify the mechanism by which speed might be controlled across contexts, we first tested whether $t_t = gt_s$ was encoded by initial conditions (A_1). According to this alternative, neural trajectories should follow the organization of t_p across both contexts (Figure 3A), in addition to within each context (Figure 4C). To test A_1 , we sorted neural trajectories across the two contexts according to t_p (Figure 5A, top), and asked whether the angle between vectors connecting nearest points ($\theta_i^{\Omega}[j]$) was significantly less than 90° (Figure 5A, bottom). Unlike the within-context results (Figure 4C), when neural trajectories from both contexts were combined, the angle between nearby neural trajectories was not less than 90° (Figures 5B and S6). This indicates that trajectories across contexts do not have an orderly relationship to t_t (A_1 : less than 90°).

Next, we investigated the hypothesis that the context cue acts as a tonic external input (A_2 ; Figure 3B), leading the system to establish structurally similar but distinct collections of neural trajectories across contexts (Figures 6A and 6B). This hypothesis imposes a set of specific geometrical constraints on neural trajectories in the Set-Go epoch. We determined whether the data met these constraints by testing whether the converse of each could be rejected, as illustrated in Figures 6C–6F. If we denote the collection of neural trajectories in the two contexts by $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$, these constraints and tests can be formalized as follows.

(1) $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ should evolve in the same direction as a function of time with different average speeds (i.e., slower for $\{\Omega_i^{g=1.5}\}_i$). If the converse were true (i.e., trajectories evolving in different directions, Figure 6C, left), we would expect no systematic relationship between time points across the two contexts. Results from KiNeT across contexts (see STAR Methods)

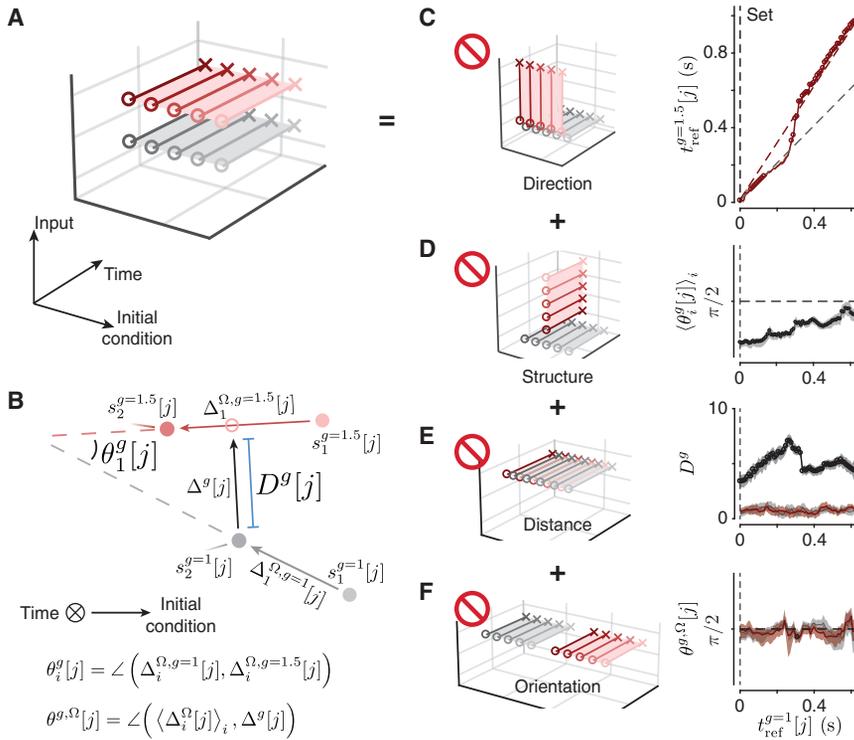


Figure 6. Neural Trajectories Comprise Distinct but Similar Structures across Gains

(A) A schematic showing the organization of neural trajectories in a subspace spanned by Input, Initial condition, and Time if context were controlled by tonic external input. If DMFC were to receive a gain-dependent input, we would expect neural trajectories from Set to Go to be separated along an input subspace, generating two similar but separated t_p -related structures for each context (A₂, Figure 3B). We verified this geometrical structure by excluding alternative structures (interdictory circles indicate rejected alternatives). (B) An illustration of neural trajectories for $g = 1$ (gray filled circle) and $g = 1.5$ (red filled circle) with the time axis normal to the plane of illustration. Gray and red arrows show vectors connecting nearby points in each context independently ($\Delta^{\Omega, g=1.5}$ and $\Delta^{\Omega, g=1}$). When the neural trajectories associated with the two gains are structured similarly, these vectors are aligned and the angle between them (θ^g) is less than 90° . We used KiNeT to test this possibility (see STAR Methods).

(C) Left: schematic illustrating a condition in which the time axes for trajectories in the two contexts (gray and red) are not aligned. Right: $\{t_{ref}^{g=1}[j]\}_j$ increased monotonically with $\{t_{ref}^{g=1.5}[j]\}_j$, indicating that the time axes across contexts were aligned. Values of $\{t_{ref}^{g=1.5}[j]\}_j$ above the unity line indicate that activity evolved at a slower speed in

the $g = 1.5$ context. The dashed gray line represents unity and the dashed red line represent expected values for $\{t_{ref}^{g=1.5}[j]\}_j$ if speeds were scaled perfectly by a factor of 1.5.

(D) Left: schematic illustrating an example configuration in which $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ establish dissimilar t_p -related structures. Right: $\langle \theta_i^g[j] \rangle_i$ was significantly less than 90° for all j indicating that the t_p -structure was similar across the two contexts.

(E) Left: schematic illustrating a condition in which $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ are overlapping. Right: the minimum distance D^g across contexts (black line) was substantially larger than that found between subsets of trajectories within contexts (red and gray lines, see STAR Methods), indicating that the two sets of trajectories were not overlapping.

(F) Left: schematic illustrating a condition in which $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ are separated along the same direction that neural trajectories within each context were separated. Right: the vector associated with the minimum distance between the two manifolds ($\Delta^g[j]$) was orthogonal to the vector connecting nearby states for both $g = 1$ (gray, $\langle \Delta_i^{\Omega, g=1}[j] \rangle_i$) and $g = 1.5$ (red, $\langle \Delta_i^{\Omega, g=1.5}[j] \rangle_i$). In (C)–(E), shaded regions represent 90% bootstrap confidence intervals, and circles represent statistical significance ($p < 0.05$, bootstrap test, $n = 100$). See Figure S7 for individual animals and Figure S11 for different recording sites.

revealed a monotonically increasing relationship between $\{t_{ref}^{g=1}[j]\}_j$ and $\{t_{ref}^{g=1.5}[j]\}_j$, confirming that Set-Go trajectories across contexts evolved in the same direction (Figures 6C, right, S7A, and S7E). Moreover, $\{t_{ref}^{g=1.5}[j]\}_j$ had a higher rate of change than $\{t_{ref}^{g=1}[j]\}_j$ indicating that average speeds were slower in the $g = 1.5$ condition. This suggests that speed control played a consistent role across contexts (Figure 6A). Consistent with the within-gain analyses (Figure 4B), KiNeT indicated that Ω_i evolved at similar speeds across context immediately following the Set cue (unity slope).

(2) $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ should be organized similarly with respect to t_p . In other words, the vector that connects nearby points in $\{\Omega_i^{g=1}\}_i$ should be aligned to its counterpart that connect nearby points in $\{\Omega_i^{g=1.5}\}_i$. To evaluate this constraint, we used the angle between pairs of vectors that connect nearby points within each context. We use an example to illustrate the procedure (Figure 6B). Consider one vector connecting nearby points in two successive neural trajectories in the gain of 1 (e.g., $\Omega_1^{g=1}$ and $\Omega_2^{g=1}$), and another vector connecting the corresponding points in the gain of 1.5 (e.g., $\Omega_1^{g=1.5}$ and $\Omega_2^{g=1.5}$). A similar orientation between the two groups of trajectories (Fig-

ure 6A) would cause the angle between these vectors (θ_i^g) to be significantly smaller than 90° . If instead, $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ were oriented differently (Figure 6D, left) or had no consistent relationship, these vectors would be on average orthogonal. Using KiNeT, we found that this angle ($\theta_i^g[j]$) was consistently smaller than 90° throughout the Set-Go epoch (Figures 6D, right, S7B, and S7F), providing quantitative evidence that the collection of neural trajectories associated with the two gains were structurally similar (Figure 6A).

(3) If context information is provided as a tonic input, $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ should be separated in state space along a context axis throughout the Set-Go epoch. To verify this constraint, we assumed that neural trajectories for each context were embedded in distinct manifolds and compared the minimum distance between the two manifolds (D^g) to an analogous distance metric within each manifold (Figure 6B; see STAR Methods). These distance measures should be the same if the groups of trajectories associated with the two contexts overlap in state space (Figure 6E, left). However, we found distances to be substantially larger across contexts compared to within contexts (Figures 6E, right, S7C, and S7G). This confirms that the groups

of trajectories associated with the two contexts were separated in state space (Figure 6A).

(4) The results so far reject a number of alternative hypotheses (Figures 6C–6E) and leave two possibilities: either $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ are separated along the same dimension that separates trajectories within each context (Figure 6F, left), or they are separated along a distinct input axis in accordance with A_2 (Figure 6A). To distinguish between these two, we asked whether the vector associated with the minimum distance $D^g[j]$ ($\Delta^g[j]$) between the two manifolds was aligned to vectors connecting nearby states within each context ($\{\Delta_i^g[j]\}_i$). Analysis of the angle between these vectors ($\langle\theta_i^{g,\Omega}[j]\rangle$) indicated that the two were orthogonal for almost all j (Figures 6F, right, S7D, and S7H). This ruled out the remaining possibility that trajectories across contexts were separated along the same dimension as within context (Figure 6F, left).

Having validated these constraints quantitatively, we concluded that population activity across gains formed two groups of isomorphic speed-dependent neural trajectories (Figure 6A). These results support our primary hypothesis that flexible control of speed based on gain context was established by a context-dependent tonic external input (Figure 3B).

Variability in Neural Trajectories Systematically Predicted Behavioral Variability

To further substantiate the link between the geometry of neural trajectories and behavior, we asked whether variability in t_p for each t_s and gain could be explained in terms of systematic fluctuations of speed and location of neural trajectories in the space. Considering the orderly structure of nearby neural trajectories, we reasoned that two types of deviations along neural trajectories should lead to systematic biases in t_p . First, on trials in which the trajectory deviates toward the trajectory associated with a shorter (or longer) t_s , the trial should result in a shorter (or longer) t_p . Second, when a trajectory evolves too slowly (closer to its positions at Set) or too fast (closer to its positions at Go), the trial should result in a longer (or shorter) t_p . We tested this by using KiNeT to examine the relative geometrical organization of neural trajectories associated with larger and smaller values of t_p for the same t_s and gain. If we denote deviations across trajectories by ε^Ω and deviations along trajectories by ε^t (Figure 7A), these predictions can be formalized as follows.

(P_1) Deviations ε^Ω off of one trajectory toward a trajectory associated with longer t_s should lead to longer t_p , and vice versa (Figure 7B). To test P_1 , we divided trials for each t_s into two bins. One bin contained all trials in which t_p was shorter than median t_p , and the other contained all trials in which t_p was longer than median t_p . We computed neural trajectories for the short and long t_p bins and denoted the corresponding states by $s_i^{\text{short}}[j]$ and $s_i^{\text{long}}[j]$, respectively. If P_1 is correct, then the geometric relationship between $s_i^{\text{short}}[j]$ and $s_i^{\text{long}}[j]$ should be similar to that between $s_{i-1}[j]$ (shorter t_s) and $s_{i+1}[j]$ (longer t_s). Therefore, the vector pointing from $s_i^{\text{short}}[j]$ to $s_i^{\text{long}}[j]$ ($\Delta_i^p[j]$) and the vector pointing from $s_{i-1}[j]$ to $s_{i+1}[j]$ ($\Delta_i^\Omega[j]$) should be aligned, and the angle between them, denoted by $\langle\theta_i^{p,\Omega}[j]\rangle$, should be acute. See STAR Methods description for calculation of $\Delta_i^\Omega[j]$ for t_s .

Consistent with P_1 , we found that average $\langle\theta_i^{p,\Omega}[j]\rangle$ ($\langle\langle\theta_i^{p,\Omega}[j]\rangle\rangle_i$) was less than 90° throughout the Set-Go epoch (Figures 7D,

7E, S8A, and S8C blue trace). This result indicates that neural trajectories that correspond to larger values of t_p were shifted in state space toward larger values of t_s . Notably, the systematic relationship between t_p and neural activity was already present at the time of the Set, indicating that t_p was influenced by variability during the Ready-Set measurement epoch and thus the initial condition in the production epoch.

(P_2) Deviations ε^t along trajectories should influence the time it takes for activity to reach the Go state and should therefore influence t_p (Afshar et al., 2011; Michaels et al., 2015). If P_2 is correct, then $s_i^{\text{short}}[j]$ should be ahead of $s_i^{\text{long}}[j]$. Therefore, $\Delta_i^p[j]$ should point backward in time, and the angle between $\Delta_i^p[j]$ and $\Delta_i^t[j]$ that connects $s_i[j-1]$ to $s_i[j+1]$, denoted by $\langle\theta_i^{p,t}[j]\rangle$, should be obtuse. See STAR Methods description for calculation of $\Delta_i^t[j]$.

Consistent with P_2 , we found that $\langle\theta_i^{p,t}[j]\rangle_i$ was greater than 90° . This result indicates that t_p was larger (smaller) when speed along the neural trajectory was slower (faster) (Figures 7D, S8B, and S8D, green trace). The angle $\langle\theta_i^{p,t}[j]\rangle_i$ was initially close to 90° consistent with the observation that trajectories evolved at similar speeds early in the Set-Go epoch (Figure 4B). We also measured the angle between $\Delta_i^\Omega[j]$ and $\Delta_i^t[j]$, denoted by $\langle\theta_i^{\Omega,t}[j]\rangle$. $\langle\theta_i^{\Omega,t}[j]\rangle_i$ was not significantly different than expected by chance (90°) for most time points.

We also considered whether deviations of neural trajectories along the gain axis that presumably reflect variability in the context-dependent input might be related to systematic changes in t_p . We did not find evidence that this was the case. While there was statistically significant alignment between $\Delta_i^p[j]$ and vectors connecting trajectories of identical t_s across gains ($\Delta_i^g[j]$), this alignment was explained by the component of $\Delta_i^g[j]$ along the direction of $\Delta_i^\Omega[j]$ (Figure S9).

This analysis extends the correspondence between behavior and the organization of neural trajectories to include animals' within-condition variability. Together, these results provide evidence for our hypothesis that activity during Ready-Set epoch parametrically adjusts the system's initial condition (i.e., neural state at the time of Set) within contexts, which in turn controls the speed of neural trajectory in the Set-Go epoch and the consequent t_p .

RNN Models Recapitulate the Predictions of Inputs and Initial Conditions

The geometry and dynamics of DMFC responses were consistent with the hypothesis that behavioral flexibility in the RSG task relies on systematic adjustments of initial conditions and external inputs of a dynamical system. Motivated by recent advances in the use of recurrent neural networks (RNNs) as a tool for testing hypotheses about cortical dynamics (Mante et al., 2013; Hennequin et al., 2014; Sussillo et al., 2015; Chaisangmongkon et al., 2017; Wang et al., 2018), we investigated whether RNNs trained to perform the RSG task would establish similar geometrical structures and dynamics.

We created RNNs with two different architectures, one in which the gain information was provided by the magnitude of a tonic input, and another in which the gain information was provided by the magnitude of a transient pulse before the Ready cue. We refer to these networks as tonic-input RNNs and transient-input RNNs, respectively (Figure 8A). We used the

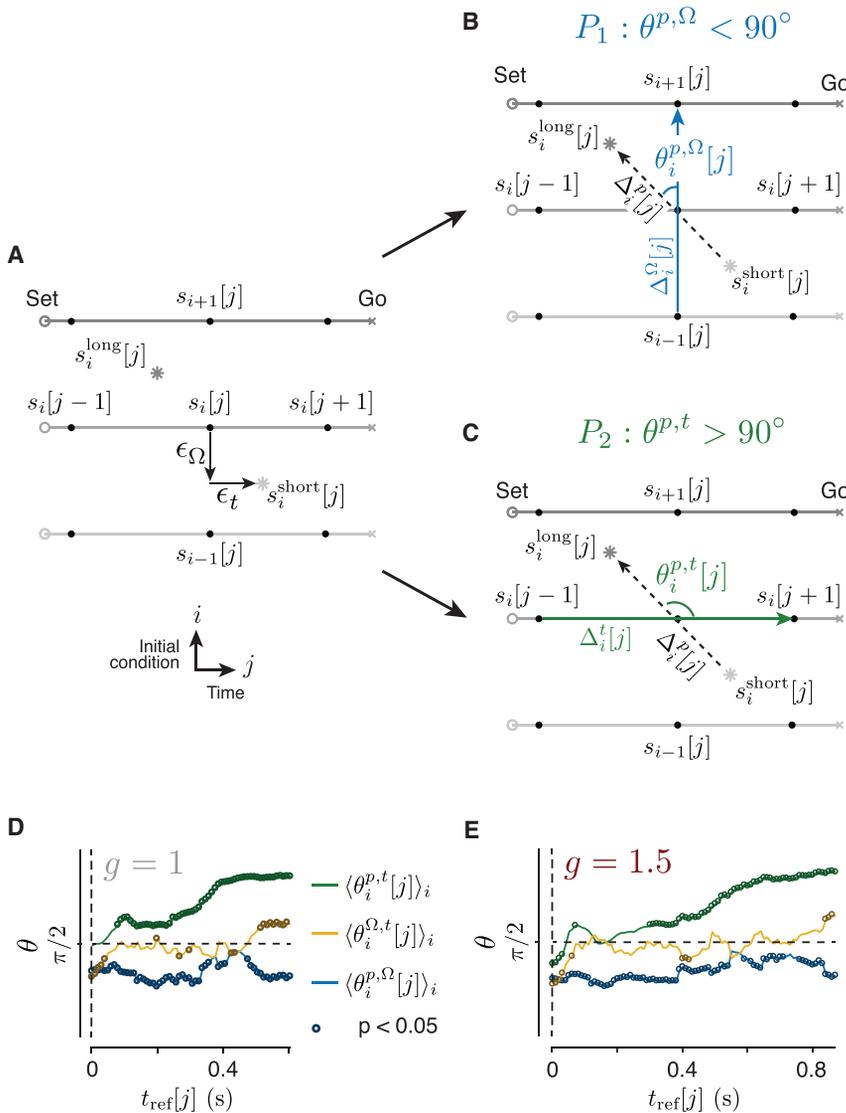


Figure 7. Relating Neural Variability to Behavioral Variability

(A) Schematic showing states (filled circles) along three neural trajectories (gray lines) between Set (circle) and Go (cross) associated with three different t_s values. The light and dark stars depict two neural states that are deviated from the average neural trajectory (the middle trajectory). The light star corresponds to trials in which t_p was shorter than median t_p ($s_i^{short}[j]$), and the dark star to trials in which t_p was longer than median ($s_i^{long}[j]$). $s_i^{short}[j]$ is deviated from the average trajectory in the direction of shorter t_s (toward $s_{i-1}[j]$ by vector ϵ^Ω) and in the direction of the Go state (toward $s_i[j+1]$ by vector ϵ^t). The opposite is true for $s_i^{long}[j]$.

(B) Prediction 1 (P_1): deviations ϵ^Ω off of one trajectory toward a trajectory associated with longer t_s should lead to longer t_p , and vice versa. The dashed arrow represents $\Delta_i^{p, \Omega}[j]$, the vector pointing from $s_i^{short}[j]$ to $s_i^{long}[j]$, and the blue arrow represents $\Delta_i^\Omega[j]$, the vector pointing from $s_{i-1}[j]$ to $s_{i+1}[j]$. P_1 is satisfied if the angle between these vectors, denoted by $\theta_i^{p, \Omega}[j]$, is acute.

(C) Prediction 2 (P_2): deviations ϵ^t along trajectories should influence the time it takes for activity to reach the Go state and should therefore influence t_p . The green arrow represents $\Delta_i^{p, t}[j]$, the vector that connects $s_i[j-1]$ to $s_i[j+1]$. If P_2 is correct, the angle $\theta_i^{p, t}[j]$ between $\Delta_i^{p, \Omega}[j]$ and $\Delta_i^t[j]$ should be obtuse.

(D and E) Testing P_1 and P_2 for the $g=1$ (D) and $g=1.5$ (E) contexts. Consistent with P_1 , average $\theta_i^{p, \Omega}[j]$ ($\langle \theta_i^{p, \Omega}[j] \rangle_i$, blue) was less than 90° from Set to Go indicating that t_p was longer (shorter) when neural states deviated toward a trajectory associated with a longer (shorter) t_s . Consistent with P_2 , $\langle \theta_i^{p, t}[j] \rangle_i$ (green) was greater than 90° , indicating that t_p was longer (shorter) when speed along the neural trajectory was slower (faster). The average angle between $\Delta_i^{p, \Omega}[j]$ and $\Delta_i^t[j]$ ($\langle \theta_i^{p, t}[j] \rangle_i$, yellow) was not significantly different than expected by chance (90°) for most time points. We determined when (at what j) an angle was significantly different from 90° ($p < 0.05$) by

comparing angles to the corresponding null distribution derived from 100 random shuffles with respect to t_p . Angles that were significantly different from 90° are shown by darker circles. See Figure S8 for individual animals and Figure S11 for different recording sites.

tonic-input RNN as a direct test of whether a network with a context-dependent tonic input could emulate the geometrical structure of responses in DMFC. The transient-input RNN, on the other hand, tested whether this structure would naturally emerge in a network without such an input.

All networks comprised synaptically coupled nonlinear units that received nonspecific background activity and were provided an input encoding Ready and Set as two brief pulses separated by t_s . We trained these RNNs to generate a linear output function after Set that reached a threshold (Go) at the desired production interval, $t_t = gt_s$ on average. Within-context analysis of successfully trained RNNs revealed that they controlled t_p by adjusting the speed of neural trajectories within a low-dimensional geometrical structure parameterized by initial conditions (Figures 8B, 8C, and S10). These results are qualitatively similar to what we found in DMFC (Figures 2C and 4).

See Figure S10 for a more detailed treatment of the networks' response properties.

Using PCA and KiNeT to analyze activity across contexts, we found that neural trajectories in the tonic- and transient-input networks were structured differently. In the tonic-input RNN, trajectories formed two isomorphic structures separated along the dimension associated with the gain-dependent tonic input (Figure 8B). In contrast, trajectories generated by the transient-input RNN were better described as coalescing toward a single structure parameterized by initial condition (Figure 8C). To verify these observations quantitatively, we evaluated the geometry of neural trajectories in the two RNN variants using the same analyses that we performed on DMFC activity. In particular, we sorted trajectories with respect to t_p across the two gain contexts ($g=1$ and $g=1.5$) and quantified the angle between vectors connecting nearest points ($\theta_i^{p, t}[j]$). As noted in the analysis of DMFC, this

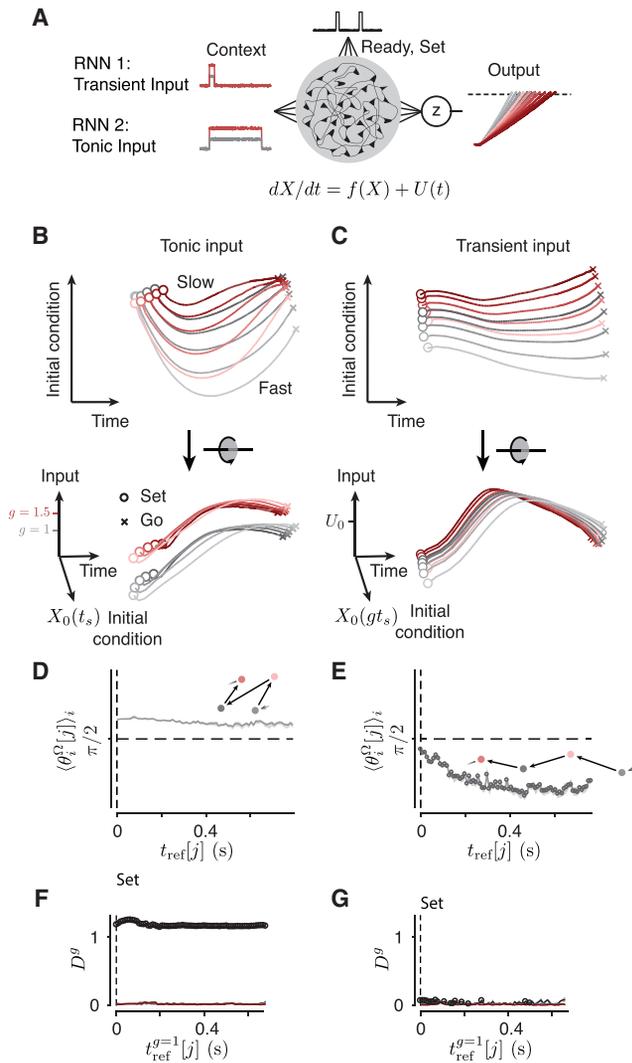


Figure 8. RNNs with Tonic but Not Transient Input Captured the Structure of Activity in DMFC

(A) Schematic illustration of the RNNs. The networks receive brief Ready and Set pulses separated in time by t_s . Additionally, each network is provided with a context-dependent “input” that either terminates prior to Ready (“Transient input,” top) or persists throughout the trial (“Tonic input,” bottom). All networks were trained so that the output (z) would generate a ramp after Set that would reach a threshold (dashed line) at the context-dependent t_t .

(B) Top: state-space projections of tonic-input RNN activity in the Set-Go epoch. The axes for 3D projections were identified using the same method as in Figure 2C.

(C) Same as (B) for the transient-input RNN.

(D) Analysis of direction in the tonic-input RNN with the same format as Figure 5B. $\langle \theta_t^g[j] \rangle_i$ was larger than 90° for the entire the Set-Go epoch. This indicates that the tonic network forms two separate context-specific sets of isomorphic neural trajectories (inset).

(E) Same as (D) for the transient-input network. $\langle \theta_t^g[j] \rangle_i$ was consistently less than 90° consistent with a geometry in which neural trajectories are organized with respect to t_p regardless of the gain context (inset).

(F and G) Trajectory separation (D^g) across contexts for the tonic-input (F) and transient-input (G) networks with the same format as Figure 6E. D^g was substantially larger through the Set-Go epoch in the tonic-input network (F). In (D)–(G), shaded regions represent 90% bootstrap confidence

angle is expected to be acute if trajectories form a single structure (A_1 : $\langle \theta_t^g[j] \rangle_i < 90^\circ$), and obtuse if trajectories form two gain-dependent structures (A_2 : $\langle \theta_t^g[j] \rangle_i \geq 90^\circ$). As predicted, the tonic-input RNNs formed two isomorphic structures (A_2) indicating that when a tonic gain-dependent input is present, RNNs rely on a solution with separate gain-dependent geometrical structures (Figure 8E). In contrast, in the transient-input RNNs, angles between consecutive trajectories were acute (A_1). This result underscores the importance of a tonic gain-dependent input in establishing separate isomorphic structures (Figure 8D).

We also compared the two RNNs in terms of the distance between trajectories across the two contexts using the same metric (D^g) we used previously for the analysis of DMFC (Figure 6E). The minimum distance between $\{\Omega_i^{g=1}\}_i$ and $\{\Omega_i^{g=1.5}\}_i$ at the time of Set was consistently smaller in the transient-input RNNs compared to tonic-input RNNs (Figures 8F and 8G). We compared values of D^g in each RNN normalized by the distance between the trajectories that correspond to the shortest and longest t_p bin for the $g=1$ context in the same RNN. In the tonic networks, the minimum normalized distance ranged between 0.4 and 1.6, which was at least 10 times larger than the that observed in the transient networks (0.003 to 0.04). Furthermore, trajectories in all transient networks gradually established a t_t -related structure consistent with A_1 . In contrast, trajectories in the tonic networks, like the DMFC data, were characterized by two separate t_p -related structures, one for each context. These results provide an important theoretical corroboration of our intuition that when gain information is provided as tonic input, a natural solution is for the system to establish distinct and isomorphic gain-dependent sets of neural trajectories.

DISCUSSION

Our results indicate that flexible control of behavior could be parsed in terms of systematic adjustments to initial conditions and simple external inputs of a dynamical system. Activity structure within each gain context was consistent with the hypothesis that the system’s initial conditions controlled t_p by parameterizing the speed of neural trajectories (Shinomoto et al., 2011; Jazayeri and Shadlen, 2015; Wang et al., 2018). The displacement of neural trajectories in DMFC state space (Shinomoto et al., 2011; Wang et al., 2018) as a function of gain, and the lack of structural representation of t_p across both gains suggested that neurons received the gain information as a context-dependent tonic input. Following recent advances in using RNNs to generate and test hypotheses about dynamical systems (Rigotti et al., 2010; Mante et al., 2013; Hennequin et al., 2014; Sussillo et al., 2015; Rajan et al., 2016; Chaisangmongkon et al., 2017; Wang et al., 2018), we corroborated this interpretation by analyzing the behavior of different RNN models trained to perform the RSG task with either tonic or transient context-dependent inputs. Although both types of networks used initial

intervals, and circles represent statistical significance ($p < 0.05$, bootstrap test, $n = 100$). See Figure S10 for a more detailed analysis of RNN results.

conditions to set the speed of neural trajectories within contexts, only the tonic-input RNNs reliably established and maintained separate structures of neural trajectories across gains, similar to what we found in DMFC.

Although we do not know the constraints that led the brain to establish separate geometrical structures, we speculate about potential computational advantages associated with this particular solution. First and foremost, this may be a particularly robust solution; as the gain information was provided by a persistent visual cue, the brain could use this input as a reliable signal to modulate neural dynamics in RSG. This solution may also reflect animals' learning strategy. We trained monkeys to perform the RSG tasks with two gain contexts. At one extreme, animals could have treated these as completely different tasks leading to completely unrelated response structures for the two gains. At the other extreme, animals could have established a single parametric solution that would enable the animal to perform the two contexts as part of a single continuum (e.g., represent t_t). DMFC responses, however, did not match either extreme. Instead, the system established what might be viewed as a modular solution comprised of two separate isomorphic structures. We take this as evidence that the brain sought similar solutions for the two contexts, but it did so while keeping the solutions separated in the state space. This strategy preserves a separable, unambiguous representation of gain and t_s at the population level (Machens et al., 2010; Mante et al., 2013; Kobak et al., 2016) and provides the additional flexibility of parametric adjustments to the two parameters independently. Future extensions of our experimental paradigm to cases where context information is not present throughout the trial (e.g., internally inferred rules) might provide a more direct test of these possibilities.

Regardless of the learning strategies and constraints that shaped DMFC responses, our results highlight an important computational role for inputs that deviate from traditional views. We found that changing the level of a static input can be used to generalize an arbitrary stimulus response mapping in the RSG task to multiple contexts while preserving the computational mechanisms within contexts. Similar inferences can be made from other recent studies that have evaluated the computational utility of inputs that encode task rules and behavioral contexts (Mante et al., 2013; Song et al., 2016; Chaisangmongkon et al., 2017; Wang et al., 2018). Extending this idea, it may be possible for the system to use multiple orthogonal input vectors to flexibly and rapidly switch between sensorimotor mappings along different dimensions. Together, these findings suggest that a key function of cortical inputs may be to flexibly reconfigure the intrinsic dynamics of cortical circuits by driving the system to different regions of the state space. This allows the same group of neurons to access a reservoir of latent dynamics needed to perform different task-relevant computations.

Our results raise a number of additional important questions. First, future work should confirm the existence of and identify the source of the putative context-dependent input in the RSG task among candidate cortical and subcortical areas (Bates and Goldman-Rakic, 1993; Lu et al., 1994; Wallis et al., 2001; Wang et al., 2005; Akkal et al., 2007). The mechanism by which such putative input modulates cortical activity is also unknown.

In our RNN models, context information was provided by an external drive and the model was agnostic with respect to the exact mechanism. In the cortex, control of speed of dynamics may be mediated by an input that either drives cortical neurons toward their saturating nonlinearity (Wang et al., 2018) or modulates the synaptic properties or activation functions of a subset of cortical neurons (Harris and Thiele, 2011; Nadim and Bucher, 2014; Dipoppa et al., 2018). Second, while the signals recorded in this study were consistent with a prominent role for DMFC in RSG, other brain areas such as the thalamus (Guo et al., 2017; Schmitt et al., 2017) and prefrontal cortex (Miller and Cohen, 2001) are also likely to help maintain the observed dynamics. In particular, while we argue that the initial conditions in DMFC appear sufficient to set the speed of the dynamics within contexts, other areas may provide interval-dependent input during the Set-Go epoch (Wang et al., 2018). Recording from these areas would test this possibility. Third, although we assumed that recurrent interactions were fixed during our experiment, it is almost certain that synaptic plasticity plays a key role as the network learns to incorporate context-dependent inputs (Pascual-Leone et al., 1995; Kleim et al., 1998; Xu et al., 2009; Yang et al., 2014). Finally, it is possible that factors not measured by extracellular recording (e.g., short-term synaptic plasticity) contribute to both interval (Buonomano and Merzenich, 1995; Karmarkar and Buonomano, 2007; Murray and Escola, 2017) and contextual control (Stokes et al., 2013) in RSG. These open questions aside, our results provide a novel "computation through dynamics" framework to link neural activity to behavior.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - RSG Task
 - Recording
 - Analysis of DMFC data
 - Kinematic analysis of neural trajectories (KiNeT)
 - Organization of trajectories in state space
 - Analysis of neural trajectories across contexts
 - Using KiNeT to calculate variance explained by interval and gain
 - Calculation of $\Delta_i^{\Omega}[j]$ for comparing variability within and across neural trajectories
 - Calculation of $\Delta_i^{\dagger}[j]$ for comparing variability within and across neural trajectories
 - Recurrent neural network
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information includes 11 figures and one table and can be found with this article online at <https://doi.org/10.1016/j.neuron.2018.05.020>.

ACKNOWLEDGMENTS

We thank S.W. Egger, H. Sohn, and V. Parks for their helpful suggestions on the manuscript. D.N. is supported by the Rubicon grant (2015/446-14-008) from the Netherlands organization for scientific research (NWO) and by the Marie Skłodowska-Curie reintegration grant (PredOpt 796577) from the European Commission. M.J. is supported by NIH (NINDS-NS078127), the Alfred P. Sloan Foundation (BR-2014-102), the Esther A. and Joseph Klingenstein Fund, the Simons Foundation (542993SPI), the McKnight Endowment Fund for Neuroscience, the National Science Foundation Center for Sensorimotor Neural Engineering (UWSC6200 BPO4405), and the McGovern Institute.

AUTHOR CONTRIBUTIONS

E.D.R. and M.J. designed the main experimental paradigm. E.D.R. and E.A.H. trained the animals. E.D.R. collected neural data from both animals. E.D.R. developed KiNet. E.D.R. performed all analyses. D.N. developed the recurrent neural network models. E.D.R. and M.J. interpreted the results and wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: December 5, 2017

Revised: March 19, 2018

Accepted: May 11, 2018

Published: June 6, 2018

SUPPORTING CITATIONS

The following reference appears in the Supplemental Information: Garcia (2012).

REFERENCES

- Acerbi, L., Wolpert, D.M., and Vijayakumar, S. (2012). Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Comput. Biol.* *8*, e1002771.
- Afshar, A., Santhanam, G., Yu, B.M., Ryu, S.I., Sahani, M., and Shenoy, K.V. (2011). Single-trial neural correlates of arm movement preparation. *Neuron* *71*, 555–564.
- Akkal, D., Dum, R.P., and Strick, P.L. (2007). Supplementary motor area and presupplementary motor area: targets of basal ganglia and cerebellar output. *J. Neurosci.* *27*, 10659–10673.
- Bates, J.F., and Goldman-Rakic, P.S. (1993). Prefrontal connections of medial motor areas in the rhesus monkey. *J. Comp. Neurol.* *336*, 211–228.
- Brass, M., and von Cramon, D.Y. (2002). The role of the frontal cortex in task preparation. *Cereb. Cortex* *12*, 908–914.
- Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* *10*, 113–125.
- Buonomano, D.V., and Merzenich, M.M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* *267*, 1028–1030.
- Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. *Neuron* *86*, 1067–1077.
- Chaisangmongkon, W., Swaminathan, S.K., Freedman, D.J., and Wang, X.J. (2017). Computing by Robust Transience: How the Fronto-Parietal Network Performs Sequential, Category-Based Decisions. *Neuron* *93*, 1504–1517.e4.
- Churchland, M.M., Afshar, A., and Shenoy, K.V. (2006). A central source of movement variability. *Neuron* *52*, 1085–1096.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., and Shenoy, K.V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* *68*, 387–400.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. *Nature* *487*, 51–56.
- Coull, J.T., Vidal, F., Nazarian, B., and Macar, F. (2004). Functional anatomy of the attentional modulation of time estimation. *Science* *303*, 1506–1508.
- Cui, X., Stetson, C., Montague, P.R., and Eagleman, D.M. (2009). Ready...go: Amplitude of the fMRI signal encodes expectation of cue arrival time. *PLoS Biol.* *7*, e1000167.
- Dipoppa, M., Ranson, A., Krumin, M., Pachitariu, M., Carandini, M., and Harris, K.D. (2018). Vision and Locomotion Shape the Interactions between Neuron Types in Mouse Visual Cortex. *Neuron* *98*, 602–615.e8.
- Fetz, E.E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *The Behavioral and brain sciences* (Cambridge Univ Press). <https://doi.org/10.1017/S0140525X00072599>.
- Fujii, N., Mushiaki, H., and Tanji, J. (2002). 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.
- Garcia, C. (2012). A simple procedure for the comparison of covariance matrices. *BMC Evol. Biol.* *12*, 222.
- Guo, Z.V., Inagaki, H.K., Daie, K., Druckmann, S., Gerfen, C.R., and Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* *545*, 181–186.
- Halsband, U., Ito, N., Tanji, J., and Freund, H.J. (1993). The role of premotor cortex and the supplementary motor area in the temporal control of movement in man. *Brain* *116*, 243–266.
- Hanes, D.P., and Schall, J.D. (1996). Neural control of voluntary movement initiation. *Science* *274*, 427–430.
- Harris, K.D., and Thiele, A. (2011). Cortical state and attention. *Nat. Rev. Neurosci.* *12*, 509–523.
- Hennequin, G., Vogels, T.P., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* *82*, 1394–1406.
- Huerta, M.F., and Kaas, J.H. (1990). Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J. Comp. Neurol.* *293*, 299–330.
- Isoda, M., and Hikosaka, O. (2007). Switching from automatic to controlled action by monkey medial frontal cortex. *Nat. Neurosci.* *10*, 240–248.
- Isoda, M., and Tanji, J. (2003). Contrasting neuronal activity in the supplementary and frontal eye fields during temporal organization of multiple saccades. *J. Neurophysiol.* *90*, 3054–3065.
- Jazayeri, M., and Shadlen, M.N. (2010). Temporal context calibrates interval timing. *Nat. Neurosci.* *13*, 1020–1026.
- Jazayeri, M., and Shadlen, M.N. (2015). A Neural Mechanism for Sensing and Reproducing a Time Interval. *Curr. Biol.* *25*, 2599–2609.
- Karmarkar, U.R., and Buonomano, D.V. (2007). Timing in the absence of clocks: encoding time in neural network states. *Neuron* *53*, 427–438.
- Kaufman, M.T., Seely, J.S., Sussillo, D., Ryu, S.I., Shenoy, K.V., and Churchland, M.M. (2016). The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* *3*. ENEURO.0085-16.2016. <https://doi.org/10.1523/ENEURO.0085-16.2016>.
- Kim, J., Jung, A.H., Byun, J., Jo, S., and Jung, M.W. (2009). Inactivation of medial prefrontal cortex impairs time interval discrimination in rats. *Front. Behav. Neurosci.* *3*, 38.
- Kim, J., Ghim, J.-W., Lee, J.H., and Jung, M.W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.* *33*, 13834–13847.
- Kleim, J.A., Barbay, S., and Nudo, R.J. (1998). Functional reorganization of the rat motor cortex following motor skill learning. *J. Neurophysiol.* *80*, 3321–3325.
- Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016).

- Demixed principal component analysis of neural population data. *eLife* 5, e10989.
- Kunimatsu, J., and Tanaka, M. (2012). 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.
- Kurata, K., and Wise, S.P. (1988). Premotor and supplementary motor cortex in rhesus monkeys: neuronal activity during externally- and internally-instructed motor tasks. *Exp. Brain Res.* 72, 237–248.
- Laje, R., and Buonomano, D.V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* 16, 925–933.
- Lu, M.T., Preston, J.B., and Strick, P.L. (1994). Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. *J. Comp. Neurol.* 341, 375–392.
- Macar, F., Coull, J., and Vidal, F. (2006). The supplementary motor area in motor and perceptual time processing: fMRI studies. *Cogn. Process.* 7, 89–94.
- Machens, C.K., Romo, R., and Brody, C.D. (2010). Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *J. Neurosci.* 30, 350–360.
- Maimon, G., and Assad, J.A. (2006). A cognitive signal for the proactive timing of action in macaque LIP. *Nat. Neurosci.* 9, 948–955.
- Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84.
- Matell, M.S., Meck, W.H., and Nicolelis, M.A.L. (2003). Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behav. Neurosci.* 117, 760–773.
- Matsuzaka, Y., and Tanji, J. (1996). Changing directions of forthcoming arm movements: neuronal activity in the presupplementary and supplementary motor area of monkey cerebral cortex. *J. Neurophysiol.* 76, 2327–2342.
- Matsuzaka, Y., Aizawa, H., and Tanji, J. (1992). 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.
- Meister, M.L.R., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J. Neurosci.* 33, 2254–2267.
- Merchant, H., Zarco, W., Pérez, O., Prado, L., and Bartolo, R. (2011). Measuring time with different neural chronometers during a synchronization-continuation task. *Proc. Natl. Acad. Sci. USA* 108, 19784–19789.
- Merchant, H., Pérez, O., Zarco, W., and Gámez, J. (2013). Interval tuning in the primate medial premotor cortex as a general timing mechanism. *J. Neurosci.* 33, 9082–9096.
- Michaels, J.A., Dann, B., Intveld, R.W., and Scherberger, H. (2015). Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J. Neurosci.* 35, 11415–11432.
- Michaels, J.A., Dann, B., and Scherberger, H. (2016). Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput. Biol.* 12, e1005175.
- Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
- Mita, A., Mushiaki, H., Shima, K., Matsuzaka, Y., and Tanji, J. (2009). Interval time coding by neurons in the presupplementary and supplementary motor areas. *Nat. Neurosci.* 12, 502–507.
- Miyazaki, M., Nozaki, D., and Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *J. Neurophysiol.* 94, 395–399.
- Murakami, M., Vicente, M.I., Costa, G.M., and Mainen, Z.F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* 17, 1574–1582.
- Murray, J.M., and Escola, G.S. (2017). Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* 6, e26084.
- Nadim, F., and Bucher, D. (2014). Neuromodulation of neurons and synapses. *Curr. Opin. Neurobiol.* 29, 48–56.
- Ohmae, S., Lu, X., Takahashi, T., Uchida, Y., and Kitazawa, S. (2008). Neuronal activity related to anticipated and elapsed time in macaque supplementary eye field. *Exp. Brain Res.* 184, 593–598.
- Okano, K., and Tanji, J. (1987). Neuronal activities in the primate motor fields of the agranular frontal cortex preceding visually triggered and self-paced movement. *Exp. Brain Res.* 66, 155–166.
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Harris, K.D. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*. <https://doi.org/10.1101/061481>.
- Pascual-Leone, A., Nguyet, D., Cohen, L.G., Brasil-Neto, J.P., Cammarota, A., and Hallett, M. (1995). Modulation of muscle responses evoked by transcranial magnetic stimulation during the acquisition of new fine motor skills. *J. Neurophysiol.* 74, 1037–1045.
- Pfeuty, M., Ragot, R., and Pouthas, V. (2005). Relationship between CNV and timing of an upcoming event. *Neurosci. Lett.* 382, 106–111.
- Pruszynski, J.A., Kurtzer, I., Nashed, J.Y., Omrani, M., Brouwer, B., and Scott, S.H. (2011). Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* 478, 387–390.
- Rabinovich, M., Huerta, R., and Laurent, G. (2008). Neuroscience. Transient dynamics for neural processing. *Science* 321, 48–50.
- Rajan, K., and Abbott, L.F. (2006). Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.* 97, 188104.
- Rajan, K., Harvey, C.D., and Tank, D.W. (2016). Recurrent Network Models of Sequence Generation and Memory. *Neuron* 90, 128–142.
- Rakitin, B.C., Gibbon, J., Penney, T.B., Malapani, C., Hinton, S.C., and Meck, W.H. (1998). Scalar expectancy theory and peak-interval timing in humans. *J. Exp. Psychol. Anim. Behav. Process.* 24, 15–33.
- Rao, S.M., Mayer, A.R., and Harrington, D.L. (2001). The evolution of brain activation during temporal processing. *Nat. Neurosci.* 4, 317–323.
- Ray, S., and Heinen, S.J. (2015). A mechanism for decision rule discrimination by supplementary eye field neurons. *Exp. Brain Res.* 233, 459–476.
- Remington, E., and Jazayeri, M. (2017). Late Bayesian inference in sensorimotor behavior. *bioRxiv*. <https://doi.org/10.1101/130062>.
- Rigotti, M., Ben Dayan Rubin, D., Wang, X.-J., and Fusi, S. (2010). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front. Comput. Neurosci.* 4, 24.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.
- Romo, R., and Schultz, W. (1987). Neuronal activity preceding self-initiated or externally timed arm movements in area 6 of monkey cortex. *Exp. Brain Res.* 67, 656–662.
- Romo, R., and Schultz, W. (1992). Role of primate basal ganglia and frontal cortex in the internal generation of movements. III. Neuronal activity in the supplementary motor area. *Exp. Brain Res.* 91, 396–407.
- Rossant, C., Kadir, S.N., Goodman, D.F.M., Schulman, J., Hunter, M.L.D., Saleem, A.B., Grosmark, A., Belluscio, M., Denfield, G.H., Ecker, A.S., et al. (2016). Spike sorting for large, dense electrode arrays. *Nat. Neurosci.* 19, 634–641.
- Schlag, J., and Schlag-Rey, M. (1987). Evidence for a supplementary eye field. *J. Neurophysiol.* 57, 179–200.
- Schmitt, L.I., Wimmer, R.D., Nakajima, M., Happ, M., Mofakham, S., and Halassa, M.M. (2017). Thalamic amplification of cortical connectivity sustains attentional control. *Nature* 545, 219–223.
- Scott, S.H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* 5, 532–546.
- Seely, J.S., Kaufman, M.T., Ryu, S.I., Shenoy, K.V., Cunningham, J.P., and Churchland, M.M. (2016). Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1. *PLoS Comput. Biol.* 12, e1005164.

- Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359.
- Shima, K., Mushiake, H., Saito, N., and Tanji, J. (1996). Role for cells in the presupplementary motor area in updating motor plans. *Proc. Natl. Acad. Sci. USA* 93, 8694–8698.
- Shinomoto, S., Omi, T., Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., and Tanji, J. (2011). Deciphering elapsed time and predicting action timing from neuronal population signals. *Front. Comput. Neurosci.* 5, 29.
- Shook, B.L., Schlag-Rey, M., and Schlag, J. (1991). Primate supplementary eye field. II. Comparative aspects of connections with the thalamus, corpus striatum, and related forebrain nuclei. *J. Comp. Neurol.* 307, 562–583.
- Smith, N.J., Horst, N.K., Liu, B., Caetano, M.S., and Laubach, M. (2010). Reversible Inactivation of Rat Premotor Cortex Impairs Temporal Preparation, but not Inhibitory Control, During Simple Reaction-Time Performance. *Front. Integr. Neurosci.* 4, 124.
- Song, H.F., Yang, G.R., and Wang, X.-J. (2016). Training Excitatory-Inhibitory Recurrent Neural Networks for Cognitive Tasks: A Simple and Flexible Framework. *PLoS Comput. Biol.* 12, e1004792.
- Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78, 364–375.
- Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* 18, 1025–1033.
- Sussillo, D., Jozefowicz, R., Abbott, L.F., and Pandarinath, C. (2016). LFADS - Latent Factor Analysis via Dynamical Systems. *arXiv*, arXiv:1608.06315, <https://arxiv.org/abs/1608.06315>.
- Tanaka, M. (2005). Involvement of the central thalamus in the control of smooth pursuit eye movements. *J. Neurosci.* 25, 5866–5876.
- Thura, D., and Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81, 1401–1416.
- Todorov, E., and Jordan, M.I. (2002). Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235.
- Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature* 411, 953–956.
- Wang, Y., Isoda, M., Matsuzaka, Y., Shima, K., and Tanji, J. (2005). Prefrontal cortical cells projecting to the supplementary eye field and presupplementary motor area in the monkey. *Neurosci. Res.* 53, 1–7.
- Wang, J., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* 21, 102–110.
- Werbos, P.J. (1990). Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 1550–1560.
- Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462, 915–919.
- Xu, M., Zhang, S.-Y., Dan, Y., and Poo, M.M. (2014). Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc. Natl. Acad. Sci. USA* 111, 480–485.
- Yang, S.-N., and Heinen, S. (2014). Contrasting the roles of the supplementary and frontal eye fields in ocular decision making. *J. Neurophysiol.* 111, 2644–2655.
- Yang, G., Lai, C.S.W., Cichon, J., Ma, L., Li, W., and Gan, W.B. (2014). Sleep promotes branch-specific formation of dendritic spines after learning. *Science* 344, 1173–1178.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Rhesus macaque (<i>Macaca mulatta</i>)	Alpha Genesis	N/A
Software and Algorithms		
MATLAB	MathWorks	https://www.mathworks.com/products/matlab.html
KiloSort	Pachitariu et al., 2016	https://github.com/cortex-lab/KiloSort
Other		
Cerebus system	Blackrock Microsystems	http://blackrockmicro.com/neuroscience-research-products/neural-data-acquisition-systems/cerebus-daq-system/
Plexon V-Probes	Plexon	https://plexon.com/products/plexon-v-probe/
Eyelink 1000 eye tracker	SR Research	https://www.sr-research.com/products/

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Mehrdad Jazayeri (mjaz@mit.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All experimental procedures conformed to the guidelines of the National Institutes of Health and were approved by the Committee of Animal Care at the Massachusetts Institute of Technology. Two adult monkeys (*Macaca mulatta*), one female (C) and one male (J), were trained to perform the Ready, Set, Go (RSG) behavioral task. Both subjects were pair housed and had not participated in previous studies. Monkeys were seated comfortably in a dark and quiet room. Stimuli and behavioral contingencies were controlled using MWorks (<https://mworks.github.io/>) on a 2012 Mac Pro computer. Visual stimuli were presented on a frontoparallel 23-inch Acer H236HL monitor at a resolution of 1920x1080 at a refresh rate of 60 Hz, and auditory stimuli were played from the computer's internal speaker. Eye positions were tracked with an infrared camera (Eyelink 1000; SR Research Ltd, Ontario, Canada) and sampled at 1 kHz.

METHOD DETAILS

RSG Task

Task contingencies

Monkeys had to measure a sample interval, t_s , and subsequently produce a target interval t_t whose relationship to t_s was specified by a context-dependent gain parameter ($t_t = \text{gain} \times t_s$) which was set to either 1 ($g = 1$ context) or 1.5 ($g = 1.5$ context). On each trial, t_s was drawn from a discrete uniform prior distribution (7 values, minimum = 500 ms, maximum = 1000 ms), and *gain* (g) was switched across blocks of trials (101 ± 49 trials; mean \pm std).

Trial structure

Each trial began with the presentation of a central fixation point (FP, circular, 0.5 deg diameter), a secondary context cue (CC, square, 0.5 deg width, 3–5 deg below FP), an open circle centered at FP (OC, radius 8–10 deg, line width 0.05 deg, gray) and three rectangular stimuli (2.0x0.5 deg, gray) placed 90 deg apart over the perimeter of OC with their long side oriented radially (Figure 1A). FP was red for the $g = 1$ context and purple for the $g = 1.5$ context. CC was placed directly below FP in the $g = 1$ context, and was shifted 0.5 deg rightward in the $g = 1.5$ context. Two of the rectangular stimuli were presented only briefly and served as placeholders for the subsequent 'Ready' and 'Set' flashes. The third rectangle served as the saccadic target ('Go'), which together with FP, CC, and OC remained visible throughout the trial. Ready was always positioned to the right or left of FP (3 o'clock or 9 o'clock position). Set was positioned 90 deg clockwise with respect to Ready and the saccadic target was placed opposite to Ready.

Monkeys had to maintain their gaze within an electronic window around FP (2.5 and 5.5 deg window for C and J, respectively) or the trial was aborted. After a random delay (uniform hazard), first the Ready and then the Set cues were flashed (83 ms, white). The two flashes were accompanied by a short auditory cue (the "pop" system sound), and were separated by t_s . The produced interval t_p was

defined as the interval between the onset of the Set cue and the time the eye position entered a 5-deg electronic window around the saccadic target. Following saccade, the response was deemed a “hit” if the error $\varepsilon = |t_p - t_t|$ was smaller than a t_t -dependent threshold $\varepsilon_{\text{thresh}} = \alpha t_t + \beta$ where α was between 0.2 and 0.25, and β was 25 ms. The exact choice of these parameters were not critical for performing the task or for the observed behavior; instead, they were chosen to maintain the animals motivated and willing to work for more trials per session. On hit trials animals received juice reward and the target turned green. The reward amount, as a fraction of maximum possible reward, decreased with increasing error according to $((\varepsilon_{\text{thresh}} - \varepsilon/\varepsilon_{\text{thresh}})^{1.5})$, with a minimum fraction 0.1 (Figure 1B). Trials in which t_p was more than 3.5 times the median absolute deviation (MAD) away from the mean were considered outliers and were excluded from further analyses.

As an initial analysis of whether monkeys learned the RSG task across gains, we fit linear regression models to the behavior separately for each gain to quantify the difference in slopes between the two contexts:

$$t_p = \beta_0 + \beta_1 t_s$$

We also fit models with an interaction term across both contexts:

$$t_p = \beta_0 + \beta_1 t_s + \beta_2 g + \beta_3 g t_s$$

If the animals successfully learned to apply the gain, β_3 should be positive.

We further applied a Bayesian observer model (Miyazaki et al., 2005; Jazayeri and Shadlen, 2010, 2015; Acerbi et al., 2012), which captured the behavior in both contexts (Figure 1E). Full details of the model can be found in previous work (Jazayeri and Shadlen, 2010, 2015). Briefly, we assumed that both measurement and production of time intervals are noisy. Measurement and production noise were modeled as zero-mean Gaussian with standard deviation proportional to the base interval (Rakitin et al., 1998), with constant of proportionality of w_m and w_p , respectively. A Bayesian model observer produced t_p after deriving an optimal estimate of t_t from the mean of the posterior. To account for the possibility that the mental operation of mapping t_s to t_t according to the gain factor might be noisier in the $g = 1.5$ context than in the $g = 1$ context (Remington and Jazayeri, 2017), we allowed w_m and w_p to vary across contexts.

Recording

We recorded neural activity in dorsomedial frontal cortex (DMFC) with 24-channel linear probes (Plexon). Recording locations were selected according to stereotaxic coordinates with reference to previous studies recording from the supplementary eye field (SEF; Schlag and Schlag-Rey, 1987; Huerta and Kaas, 1990; Shook et al., 1991; Fujii et al., 2002) and presupplementary motor area (Pre-SMA; Matsuzaka et al., 1992; Fujii et al., 2002), and the existence of task-relevant modulation of neural activity. We sampled DMFC starting from the center of the stereotaxically identified region and moved outward a quasi-systematic fashion looking for regions with strong modulation during either epoch of the RSG task. In monkey C, recordings were made using single probes in seven sessions between 3 mm to 7 mm lateral of the midline and 0.5 mm posterior to 5.5 mm anterior of the genu of the arcuate sulcus. In monkey J, we recorded using paired probes (48 total channels, 1.5 mm separation between probes) in three sessions from between 3 mm to 4.5 mm lateral of the midline and 0.75 mm to 5 mm anterior of the genu of the arcuate sulcus. We conservatively refer to the recorded regions as dorsomedial frontal cortex (DMFC), potentially comprising SEF, Pre-SMA and dorsal SMA (no recordings were made in the medial bank). Recording sites are shown in Figures S2 and S11. Data were recorded and stored using a Cerebus Neural Signal Processor (NSP; Blackrock Microsystems). Preliminary spike sorting was performed online using the Blackrock NSP, followed by offline sorting using the Phy spike sorting software package using the spikdetekt, klusta, and kilosort algorithms (Pachitariu et al., 2016; Rossant et al., 2016). Sorted spikes were then analyzed using custom code in MATLAB (The MathWorks).

Analysis of DMFC data

Average firing rates of individual neurons were estimated using a 150 ms smoothing filter applied to spike counts in 1 ms time bins. We used PCA to visualize and analyze recorded activity patterns. For both animals, neural trajectories (Figure S2) and the KiNeT results (Figures S5–S8) were qualitatively similar. We thus performed the main analyses on the combined population of neurons across animals. Combined analysis would not be warranted if results were substantially different across animals. PCA was applied after a soft normalization: spike counts measured in 10 ms bins were divided by the square root of the maximum spike count across all bins and conditions. The normalization was implemented to minimize the possibility of high firing rate neurons dominating the analysis.

When binning data according to increasing values of t_p , we ensured that all bins had equal number of trials, independently for each session. To average firing rates across trials within a group, we truncated trials to the median t_p and averaged firing rates with attrition. Analyses of neural data were applied to all 10 sessions across both monkeys. For most analyses, we included neurons for which at least 15 trials were recorded in each condition and which had a minimum unsmoothed modulation depth of 15 spikes per second. Because the comparison of t_s - versus t_p -related structure (Figure 7) required grouping trials into substantially more bins than the other analyses (14 versus 7 or 5), we reduced the minimum number of trials required to 10 for this analysis (273 units; 95 from monkey C and 178 from monkey J). We did not find that the results of any of the analyses were dependent on the specific threshold chosen, and results were similar in individual subjects. We did not separately analyze trials immediately following context switches due to the low number of context switches per session (mean = 6.8 switches).

For visualization of neural trajectories in state space, we projected population firing rates onto axes along which responses were maximally separated with respect to context (“gain axis,” Figures 2B and 2C and “input,” Figure 3C) and additionally with respect to interval within context (“interval axis,” Figures 2B and 2C, and “initial condition,” Figure 3C) for the Set-Go epoch. For the Ready-Set epoch (Figure 2B), the gain axis was defined as the vector connecting population firing rates averaged over t_s across contexts at the time of Set. For the Set-Go epoch, the interval axis was defined as the vector between neural activity for the shortest and longest t_p averaged over time and context. This component of the activity was then subtracted away from the full activity, and the gain axis was defined as the vector between neural activity for the two contexts averaged over time and t_p . For the Ready-Set epoch, we additionally plotted PCs 1 and 2 with the context component removed (Figure 2B). For the Set-Go epoch, we additionally plotted PC 1 with the context and t_p components removed (Figures 2C and 3C). Axes used to plot the RNN activity were calculated using the same method (Figures 8B and 8C).

Kinematic analysis of neural trajectories (KiNeT)

We developed KiNeT to compare the geometry, relative speed, and relative position along a group of neural trajectories that have an orderly organization and change smoothly with time. To describe KiNeT rigorously, we developed the following symbolic notations. Square and curly brackets refer to individual items and groups of items, respectively.

The algorithm for applying KiNeT can be broken down into the following steps: 1) Choose a Euclidean coordinate system to analyze the neural trajectories. We chose the first 10 PCs in the Set-Go epoch, which captured 89% of the variance in the data. 2) Designate one trajectory as reference, Ω_{ref} . We used the trajectory associated with the middle t_p bin as reference. 3) On each of the non-reference trajectories Ω_i ($i \neq \text{ref}$), find $\{s_i[j]\}_j$ with minimum Euclidean distance to $\{s_{\text{ref}}[j]\}_j$ and their associated times $\{t_i[j]\}_j$ according to the following equations:

$$t_i[j] = \arg \min_{\tau} \|\Omega_i(\tau) - s_{\text{ref}}[j]\|$$

$$s_i[j] = \Omega_i(t_i[j])$$

Organization of trajectories in state space

The distances $\{D_i[j]\}_j$ were used to characterize positions in neural state space of each Ω_i relative to Ω_{ref} . The magnitude of $D_i[j]$ was defined as the norm of the vector connecting $s_i[j]$ to $s_{\text{ref}}[j]$, which we refer to as $\Delta_i^{\text{ref}}[j]$. The sign of $D_i[j]$ was defined as follows: for the trajectory Ω_1 associated with the shortest t_p , and Ω_5 associated with the longest, $D_i[j]$ was defined to be negative and positive, respectively. For all other trajectories, $D_i[j]$ was positive if the angle between $\Delta_i^{\text{ref}}[j]$ and $\Delta_5^{\text{ref}}[j]$ was smaller than the angle between $\Delta_i^{\text{ref}}[j]$ and $\Delta_1^{\text{ref}}[j]$, and negative otherwise.

Analysis of neural trajectories across contexts

We initially applied the KiNeT procedure to the trajectories ordered by duration irrespective of context (Figure 5). Next, we analyzed the relationships between the sets of trajectories for the two contexts. This required aligning the activity between the two contexts in time. To do this, we started with the aligned times found within each context, and using successive groups of neural states in the $g = 1$ context indexed by $\{t_{\text{ref}}^{g=1}[j]\}_j$, found the reference time $\{t_{\text{ref}}^{g=1.5}[j]\}_j$ in the $g = 1.5$ context for which the mean distances between neural states in paired trajectories (i.e., the first t_p bins of both gains, second t_p bins, etc.) were smallest. This resulted in an array of times from $\{t_{\text{ref}}^{g=1.5}[j]\}_j$, indexed by $\{t_{\text{ref}}^{g=1}[j]\}_j$, such that the trajectories across gains were aligned in time for subsequent analyses (Figure 6C). We also measured the distance D^g between the structures using the across-context time alignment. For successive j , we measured the minimum distance between line segments connecting consecutive trajectories within each context. For five t_p bins, this meant four line segments for each context, and $4^2 = 16$ distances. We chose the minimum of these distance values as the value of D^g between the two structures. As a point of comparison, we generated set of “null” distances by splitting trajectories from each context into odd- and even- numbered trajectories and calculating the minimum distance between the sets of connecting line segments (Figure 6E).

Using KiNeT to calculate variance explained by interval and gain

To quantify the fraction of variance explained by task parameters, we used a marginalization procedure described previously (Kobak et al., 2016). Briefly, this is done by successively subtracting away neural activity averaged over each task parameter to create centered signals, then measuring the variance of these centered signals. This marginalization procedure, however, cannot be directly applied to neural activity when trajectories are of different duration. In the Ready-Set epoch, we tackled this problem by concatenating trajectories across t_s to create one trajectory for each gain.

For the Set-Go epoch, we first used KiNeT to transform all trajectories to sequences of states ($\{s_i[j]\}_j$) of equal length, and applied the marginalization on the resulting same-length sequences. With this procedure, we were able to compute variance associated with different factors. Interval (i.e., the five t_p bins) accounted for 21% of the variance in the full dataset, and 18% of the variance in PCs 1-10 used for KiNeT-related analyses. Gain accounted for 8% of the variance in the full dataset, and 7% of

the variance in PCs 1-10. The interaction between Interval and Gain accounted for 6% and 1.5%, respectively, while condition-independent (i.e., time-related) variability accounted for 65% and 73% of the variance. This is consistent with previous studies which have found time varying, condition-independent signals to account for a large percentage of variance of neural responses in frontal cortex (Kaufman et al., 2016; Kobak et al., 2016).

Calculation of $\Delta_i^\Omega[j]$ for comparing variability within and across neural trajectories

For the analyses in Figure 7, we binned trials according to t_s , rather than t_p , and redefined $\Delta_i^\Omega[j]$ such that there was one vector for each $\Delta_i^\Omega[j]$ trajectory Ω_i , instead of one fewer (Figure 4A). For most values of i , this vector was calculated as:

$$\Delta_i^\Omega[j] = s_{i+1}[j] - s_{i-1}[j]$$

Because there was a finite number of t_s values, we could not compute $\Delta_i^\Omega[j]$ for $i = 1$ and $i = 7$ ($s_{i-1}[j]$ was not defined for $i = 1$ and $s_{i+1}[j]$ was not defined for $i = 7$). Therefore, for the shortest t_s , we defined $\Delta_i^\Omega[j]$ as $s_2[j] - s_1[j]$ (instead of $s_2[j] - s_0[j]$), and for the longest t_s , as $s_7[j] - s_6[j]$ (instead of $s_8[j] - s_6[j]$).

Calculation of $\Delta_j^t[j]$ for comparing variability within and across neural trajectories

For the analyses in Figure 7, $\Delta_j^t[j]$ was calculated similarly to $\Delta_i^\Omega[j]$, with:

$$\Delta_j^t[j] = s_i[j+1] - s_i[j-1]$$

Because there was a finite number of time points, we could not compute $\Delta_j^t[j]$ for $j = 1$ and $j = N$ (where N is the number of states in Ω^{ref} ; $s_i[j-1]$ was not defined for $j = 1$ and $s_i[j+1]$ was not defined for $j = N$). Therefore, for the first time point, we defined $\Delta_j^t[j]$ as $s_i[2] - s_i[1]$ (instead of $s_i[2] - s_i[0]$), and for the last time point, to $s_i[N] - s_i[N-1]$ (instead of $s_i[N+1] - s_i[N-1]$).

Recurrent neural network

We constructed a firing rate recurrent neural network (RNN) model with $N = 200$ nonlinear units. The network dynamics were governed by the following differential equation:

$$\tau \dot{\mathbf{x}}(t) = -\mathbf{x}(t) + \mathbf{J}\mathbf{r}(t) + \mathbf{B}\mathbf{u} + \mathbf{c}_x + \rho_x(t)$$

$$\mathbf{r}(t) = \tanh[\mathbf{x}(t)]$$

$\mathbf{x}(t)$ is a vector containing the activity of all units, and $\mathbf{r}(t)$ represents the firing rates of those units by transforming \mathbf{x} through a tanh nonlinearity. Time t was sampled every millisecond for a duration of $T = 3300$ ms. The time constant of decay for each unit was set to $\tau = 10$ ms. The unit activations also contain an offset \mathbf{c}_x and white noise $\rho_x(t)$ at each time step with standard deviation in the range [0.01-0.015]. The matrix \mathbf{J} represents recurrent connections in the network. The network received multi-dimensional input \mathbf{u} through synaptic weights $\mathbf{B} = [\mathbf{b}_c, \mathbf{b}_s]$. The input \mathbf{u} was comprised of a gain-dependent context cue $u_c(t)$ and an input $u_s(t)$ that provided Ready and Set pulses. In $u_s(t)$ Ready and Set were encoded as 20 ms pulses with a magnitude of 0.4 that were separated by time t_s .

Two classes of networks were trained to perform the RSG task with multiple gains. In the tonic-input RNNs, the gain-dependent input $u_c(t)$ was set to a fixed offset for the entire duration of the trial. In contrast, in the transient-input RNNs, $u_c(t)$ was active transiently for 440 ms and was terminated 50-130 ms before the onset of the Ready pulse. The amplitude of $u_c(t)$ was set to 0.3 for $g = 1$ and 0.4 for $g = 1.5$. The transient network received an additional gain-independent tonic input of magnitude 0.4, similar to the tonic networks. Both types of networks produced a one-dimensional output $z(t)$ through summation of units with weights \mathbf{w}_o and a bias term.

$$z(t) = \mathbf{w}_o^T \mathbf{r}(t) + c_z$$

Network Training

Prior to training, model parameters (θ), which comprised \mathbf{J} , \mathbf{B} , \mathbf{w}_o , \mathbf{c}_x and \mathbf{c}_z were initialized. Initial values of matrix \mathbf{J} were drawn from a normal distribution with zero mean and variance $1/N$, following previous work (Rajan and Abbott, 2006). Synaptic weights $\mathbf{B} = [\mathbf{b}_c, \mathbf{b}_s]$ and the initial state vector \mathbf{x}_0 and unit biases \mathbf{c}_x were initialized to random values drawn from a uniform distribution with range [-1,1]. The output weights, \mathbf{w}_o and bias \mathbf{c}_z , were initialized to zero. During training, model parameters were optimized by truncated Newton methods using backpropagation-through-time (Werbos, 1990) by minimizing a squared loss function between the network output $z_i(t)$ and a target function $f_i(t)$, as defined by:

$$H(\theta) = \frac{1}{|T|} \sum_i \sum_t (z_i(t) - f_i(t))^2$$

Here i indexes different trials in a training set ($l = \text{different gains } (g) \times \text{intervals } (t_s) \times \text{repetitions } (r)$). The target function $f_i(t)$ was only defined in the Set-Go epoch (the output of the network was not constrained during the Ready-Set epoch). The value of $f_i(t)$ was zero during the Set pulse. After Set, the target function was governed by two parameters that could be adjusted to make $f_i(t)$ nonlinear, scaling, non-scaling or approximately linear:

$$f_i(t) = A \left(e^{\frac{t}{\alpha t_s}} - 1 \right)$$

For the networks reported, $f_i(t)$ was an approximately linear ramp function parametrized by $A = 3$ and $\alpha = 2.8$. Variable t_t represents the transformed interval for a given t_s and gain g . Solutions were robust with respect to the parametric variations of the target function (e.g., nonlinear and non-scaling target functions). In trained networks, the production time, t_p was defined as the time between the Set pulse and when the output ramped to a fixed threshold ($z_i = 1$).

During training, we employed two strategies to obtain robust solutions. First, we trained the networks to flexibly switch between three gain contexts, the two original values ($g = 1$ and $g = 1.5$) and an additional intermediate value of $g = 1.25$ for which the amplitude of $u_c(t)$ was set to 0.35. However, the behavior of networks trained with the two original gains were qualitatively similar. Second, we set $\rho(\mathbf{t})$ to zero, and instead, the context-dependent input, $u_c(t)$ received white noise with standard deviation of 0.005, per unit time ($\Delta t = 0$).

QUANTIFICATION AND STATISTICAL ANALYSIS

Confidence intervals for KiNeT performed on trajectories binned according to t_p were computed by a bootstrapping procedure, randomly selecting trials with replacement 100 times. To test for statistical significance of metrics generated through the KiNeT procedure, we used bootstrap tests, where p was the fraction of bootstrap iterations for which the metric was consistent with the null hypothesis. For analyses relating neural variability to behavioral variability, we performed permutation tests in which null distributions were derived from 100 random shuffles with respect to t_p . Unless otherwise stated, significance of a measure for individual time points was set to $p < 0.05$. The results of KiNeT were similar for different methods of data smoothing.