Article

Cell Reports

A Network Activity Reconfiguration Underlies the Transition from Goal to Action

Graphical Abstract

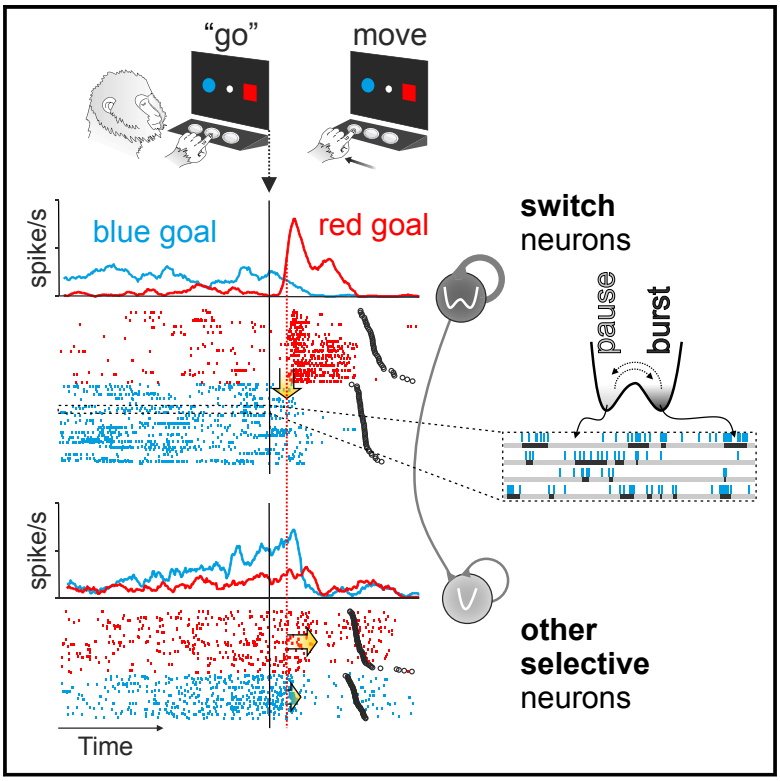

Authors

Encarni Marcos, Satoshi Tsujimoto, Maurizio Mattia, Aldo Genovesio

Correspondence

maurizio.mattia@iss.it (M.M.), aldo.genovesio@uniroma1.it (A.G.)

In Brief

Marcos et al. find a sudden activity reconfiguration of the PF network during the transition from goal in memory to action. Using a computational model, the authors propose a PF network, capable of reproducing the stability and susceptibility required for such process, composed of neurons with heterogeneous connection strengths.

Highlights

- A minority of PF neurons code the goal during both working memory and action selection
- Those neurons change or maintain goal preference with equal probability
- A network of heterogeneous cell assemblies can explain the observed dynamics

Marcos et al., 2019, Cell Reports 27, 2909–2920
June 4, 2019
https://doi.org/10.1016/j.celrep.2019.05.021

CellPress

Cell Reports
Article

A Network Activity Reconfiguration Underlies the Transition from Goal to Action

Encarni Marcos,[1,2,6] Satoshi Tsujimoto,[3,4] Maurizio Mattia,[5,*] and Aldo Genovesio[1,7,*]
[1]Department of Physiology and Pharmacology, Sapienza University of Rome, Rome, Italy
[2]Instituto de Neurociencias de Alicante, Consejo Superior de Investigaciones Científicas–Universidad Miguel Hernández de Elche, San Juan de Alicante, Spain
[3]Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University, Kyoto, Japan
[4]The Nielsen Company Pte. Ltd., Singapore, Singapore
[5]Istituto Superiore di Sanità, Rome, Italy
[6]Senior author
[7]Lead Contact
*Correspondence: maurizio.mattia@iss.it (M.M.), aldo.genovesio@uniroma1.it (A.G.)
https://doi.org/10.1016/j.celrep.2019.05.021

SUMMARY

Neurons in prefrontal cortex (PF) represent mnemonic information about current goals until the action can be selected and executed. However, the neuronal dynamics underlying the transition from goal into specific actions are poorly understood. Here, we show that the goal-coding PF network is dynamically reconfigured from mnemonic to action selection states and that such reconfiguration is mediated by cell assemblies with heterogeneous excitability. We recorded neuronal activity from PF while monkeys selected their actions on the basis of memorized goals. Many PF neurons encoded the goal, but only a minority of them did so across both memory retention and action selection stages. Interestingly, about half of this minority of neurons switched their goal preference across the goal-action transition. Our computational model led us to propose a PF network composed of heterogeneous cell assemblies with single-state and bistable local dynamics able to produce both dynamical stability and input susceptibility simultaneously.

INTRODUCTION

During natural behavior, multiple processes need to be coordinated so that specific goals can be accomplished. The prefrontal cortex (PF), by means of its connection with many other areas of the brain, plays a pivotal role in this cognitive challenge (Desimone and Duncan, 1995; Heidbreder and Groenewegen, 2003; Miller and Cohen, 2001; Squire et al., 2013; Tanji and Hoshi, 2008). Neurons in PF represent various task-related information that goes from the coding of sensory stimuli to the specific goals and actions (Averbeck et al., 2006; Falcone et al., 2016; Genovesio et al., 2006, 2008, 2012, 2014b; Genovesio and Ferraina, 2014; Genovesio and Tsujimoto, 2014; Hussar and Pasternak, 2009; Marcos and Genovesio, 2017; Saito et al., 2005). Among such information, special interest has been placed on sustained representations of stimuli and goals that are no longer present (Fuster and Alexander, 1971; Wilson et al., 1993). However, the specific neural dynamics underlying the transition from these mnemonic representations to an action are not well understood. Here, we investigate how the mnemonic trace of the goal changes or persists when the related action can be finally selected and executed.

In recent years, lateral PF functions have been investigated by adopting multiple paradigms that aimed at temporally separating the different signals that lead to action selection (Cai and Padoa-Schioppa, 2014; Markowitz et al., 2015; Sigala et al., 2008; Takeda and Funahashi, 2004; Yamagata et al., 2012). Imagine, for example, that you go to the market to buy an apple (goal). You need to hold the goal in memory until you can finally find it and select the proper action to take it. This goal-directed behavior can be roughly divided into two stages: the active maintenance of a goal in memory and the later selection of an action. Considerable effort has been made toward the description of the action selection process embedded in the context of sensory-motor transformation, in which PF neurons play an important role (Markowitz et al., 2015; Takeda and Funahashi, 2004; Zhou et al., 2016). In this case, a high proportion of PF neurons participate both in the coding of sensory-related information and in its transformation into the proper motor response (Zhou et al., 2016). Neurons in PF are also actively involved in the maintenance of goals in working memory (Genovesio and Tsujimoto, 2014; Genovesio et al., 2012; Tsujimoto et al., 2008). However, the neural dynamics underlying the transition between the two stages have remained widely unaddressed, as a conundrum challenges their full understanding. Indeed, the active maintenance of a goal in memory requires an enhanced stability of its neuronal representation as opposed to the flexibility needed to produce the network state transition behind the conversion of the goal into a proper action. A key question, then, is whether the neurons actively representing the goal in memory contribute to the goal-action transformation process or, by contrast, whether the process requires a more complex reconfiguration of the network activity.

Our recordings of the PF activity in a distance discrimination task (Genovesio et al., 2011) were ideal to investigate the neural

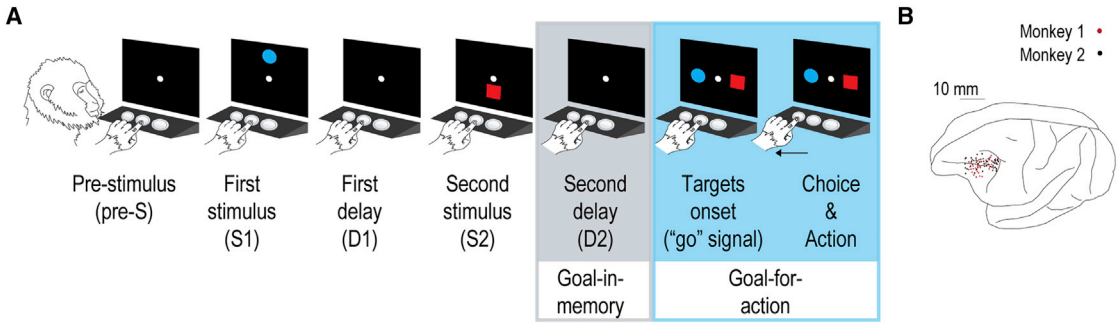

Figure 1. Experimental Setup and Neural Recordings Sites

(A) Sequence of task events within a trial. Gray and cyan rectangles indicate the phases in which we focused our analyses.

(B) Penetration sites. Composite from both monkeys, relative to sulcal landmarks.

substrate of such transition. As in our real-life example, in our experiment, the information about the ongoing task needed to be maintained in memory before the correct behavioral response could be selected and performed. Here we show that the neurons coding for the goal during the working memory period are unlikely to be involved in the subsequent goal-action transformation process. Only a minority of them code for the goal in both stages by changing or maintaining their selectivity preference between periods with equal probability. This poses a tight constraint on the possible strategy adopted by the PF network to express simultaneously both the stability and the flexibility needed to perform this task. Previous work has considered the PF network as being composed of bistable "flip-flops" (McCormick, 2005) in which two high- and low-firing attractor states are locally expressed and combined to stably encode both digital and analog information (O'Reilly, 2006). In this framework, the goal-action transformation process might include an update phase facilitated by a strong exogenous input (Churchland et al., 2010; Litwin-Kumar and Doiron, 2012) or by an increase of gain modulation of local bistable modules (Durstewitz et al., 2000b; O'Reilly, 2006). Alternative studies rely on the hypothesis that the local cortical modules composing the PF network display a wide spectrum of diverse activity levels resulting in a mixed selectivity to task-relevant information (Rigotti et al., 2010, 2013). This response heterogeneity can be the result of a suited degree of randomness in the synaptic connectivity leading to high-dimensional neuronal representations, without the need for exogenous shaping of the network dynamics. In this representational space, the network is then able to encode more than one piece of information at a time (goal, context, action to execute) in a distributed manner and support arbitrary stimulus-driven state transitions such as goal-action transformations. Here we show that besides relying on the high dimensionality of inner representations due to synaptic connectivity with a low degree of redundancy, the collective dynamics in the PF network are further enriched by the presence of local cell assemblies with heterogeneous excitability (Mattia et al., 2013). Indeed, the availability of single-state and bistable components associated to different levels of excitability can synergistically contribute to express in PF both dynamical stability and susceptibility.

RESULTS

Two monkeys were trained to discriminate which of two stimuli sequentially presented on a screen was located farther from the center (Figure 1A) (Genovesio et al., 2011). A delay period (working memory period) separated the second stimulus from the reappearance of the two stimuli (targets onset or "go" signal) randomly assigned to the right and left screen positions. Hence, monkeys had to maintain in memory the goal (blue or red stimulus) before they could reach it. Note that only after targets onset could the goal location be determined and the action planned and executed. Neurons were recorded from the PF of the two monkeys (Figure 1B), and task accuracy was high for both monkeys, with mean scores of 77% and 80% of correct responses for monkey 1 and monkey 2, respectively. Easier discrimination was associated with higher accuracy and faster reaction times (see Genovesio et al. (2011)).

Neural Coding of Goal

To study the evolution of goal coding, we first identified the neurons that represented the goal just before targets onset (pre-go period). We identified a total of 182 pre-go goal neurons (22.4%) that modulated their activity, as shown by the activity differences between preferred and nonpreferred goal trials, at least from 975 ms (start of the plot) before targets onset until 125 ms after it (Figure 2, black). Thus, these neurons ceased coding the goal shortly after the appearance of the targets. Then we asked whether that was the end of goal coding or whether another group of neurons represented the goal after targets onset. To allow for a possible delay in the transition, we analyzed the activity of the neurons in the interval from 200 ms after targets presentation to the end of movement onset. We classified 187 neurons (23%) as post-go goal neurons (see STAR Methods). In contrast to the pre-go goal neurons, the preferred and nonpreferred goal activities for post-go goal neurons (Figure 2, cyan) started to differ significantly 125 ms after targets onset. Thus, when the information about the action to perform was available, the PF network reconfigured to represent the goal through a different neural population. In other words, two different network states encoded the same goal-related information in two different epochs of a trial: the memory period

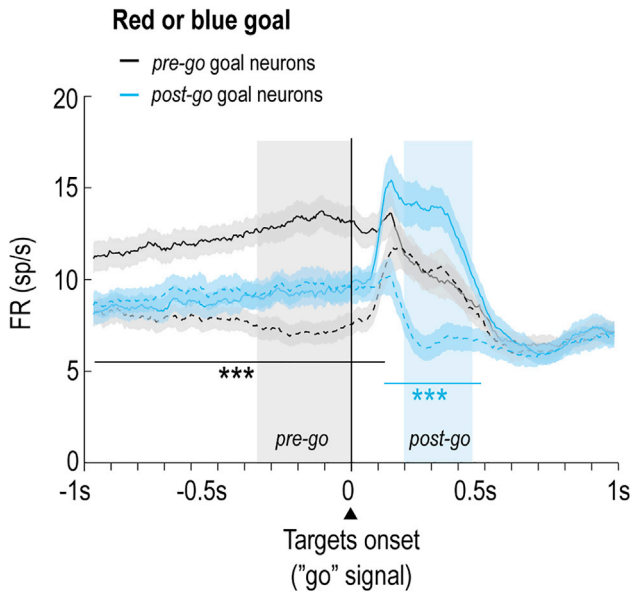

Figure 2. Pre-go and Post-go Goal Neural Population Responses

Mean firing rate of neurons (± SEM) coding the goal (blue or red) before (n = 182; black) and after (n = 187; cyan) targets onset. The time interval shown includes part of S2 presentation, the whole D2 duration (400 or 800 ms), the choice and action phase, and part of the intertrial interval. Filled gray and cyan rectangles indicate the time interval used for the statistical analyses (one-way ANOVA) performed to identify the two group of neurons. The gray rectangle shows the −350 to 0 ms period before targets onset, used to define the pre-go goal neurons, whereas the cyan rectangle indicates the time interval between 200 ms after targets onset and movement end (451.0 ± 1.7 ms [SEM]). For illustration purposes, the cyan rectangle's right border corresponds to the mean movement time calculated across sessions. Black and cyan horizontal lines indicate the consecutive time periods in which the population activity difference was significant (Wilcoxon's matched-pairs signed rank test with Bonferroni correction, ***$p < 0.001$; pre-go goal neurons, p = [0,0.00005], Z = [4.84,11.59]; post-go goal neurons, p = [0,0.0004], Z = [4.40,11.52]).

and during the goal-action transformation process. Importantly, this transition between network states was neither an artifact of the passage of time (Figure S1A) nor a consequence of the period used to identify the pre-go goal neurons (Figure S1B).

Neural Coding Dynamics

To investigate the microscopic organization of this network transition, we directly inspected single-cell activities, finding three categories of neurons (Figure 3): those that showed pre-go goal selectivity, post-go goal selectivity, or both. The neurons with goal selectivity in both periods could be further divided into neurons with either a consistent (non-switch) or an inconsistent (switch) goal preference before and after targets onset. Although the population analysis reflected the dynamics of the former two categories of neurons, it did not show any signature of switch and non-switch neurons, as their contribution averaged out because of the change of preference of some of them. Figure 3A shows an example of a pre-go goal neuron with the highest activity for the blue goal before targets onset. Its goal selectivity persisted during the D2 period, ceasing 135 ms after targets onset. The post-go goal neuron shown in Figure 3B increased its activity for the blue goal 105 ms after targets onset. Figures 3C and 3D show two neurons classified as encoding the goal during both pre-go and post-go periods. The neuron in Figure 3C maintained the same blue goal preference across trial epochs (non-switch neuron). In contrast, the neuron in Figure 3D showed a blue goal preference before targets onset that switched to the red goal 75 ms after the "go" signal (switch neuron). Altogether, this diversity in individual dynamics suggests that the apparent sequential transition from one population of neurons coding the goal in memory to another coding the same goal after targets onset was actually an oversimplification.

Next, we quantified the number of neurons belonging to each category (Figure 3E). Among the 182 goal neurons previously classified as pre-go goal neurons, 72 (39.6%) were also goal selective in the post-go period, 70 (38.5%) were also action selective (see STAR Methods), and 39 (21.4%) belonged to the three categories. On the other hand, among the 187 neurons identified as post-go goal neurons, 83 (44.4%) were also action selective. Surprisingly, almost half of the neurons with both pre-go and post-go goal selectivity (n = 72) showed a switch of their goal preference across epochs (n = 35 [48.6%]; switch neurons), whereas the other half shared the same goal preference in both periods (n = 37; non-switch neurons). Interestingly, we found that the switch in goal preference occurred only during the post-go period, as a signature of the goal-action transition (Figure S2).

From Goal to Action Coding

We have shown that there is a rapid transition between different PF network states when the information about the goal must be used to select and execute the proper action. Then we investigated when the action representation actually emerged from the neural network coding the goal after targets onset. For this purpose, we examined the neurons with a combined post-go goal and action selectivity. A paradigmatic example is the same post-go goal neuron shown in Figure 3B. In Figure 4A, its activity plotted across trials and sorted by the performed action makes apparent that the neuron had a preference for movement to the left. Such selectivity initiated about 255 ms after the "go" signal and 150 ms after the onset of the goal selectivity (Figure 4A, bottom). At the population level, the goal coding started 125 ms after targets onset and preceded by 50 ms the action coding (Figure 4B). Interestingly, the goal representation did not vanish when the action representation emerged, but rather the representation of both variables widely overlapped, decaying together around the end of the movement. Similar time latencies between goal and action coding were observed when looking at the entire population of post-go goal neurons and post-go action neurons (Figure 4C). Taking all results together, it appears that the goal information carried by the goal-selective neurons during the D2 period was not directly translated into an action. Instead, it seems to be transferred first to another network of post-go goal neurons, which are differentiated in two subpopulations of pure goal, and goal and action neurons. It is noteworthy that the same neurons coding the goal could also code the action later in time while still coding the goal. This heterogeneous set of goal- and action-selective neurons were similarly distributed

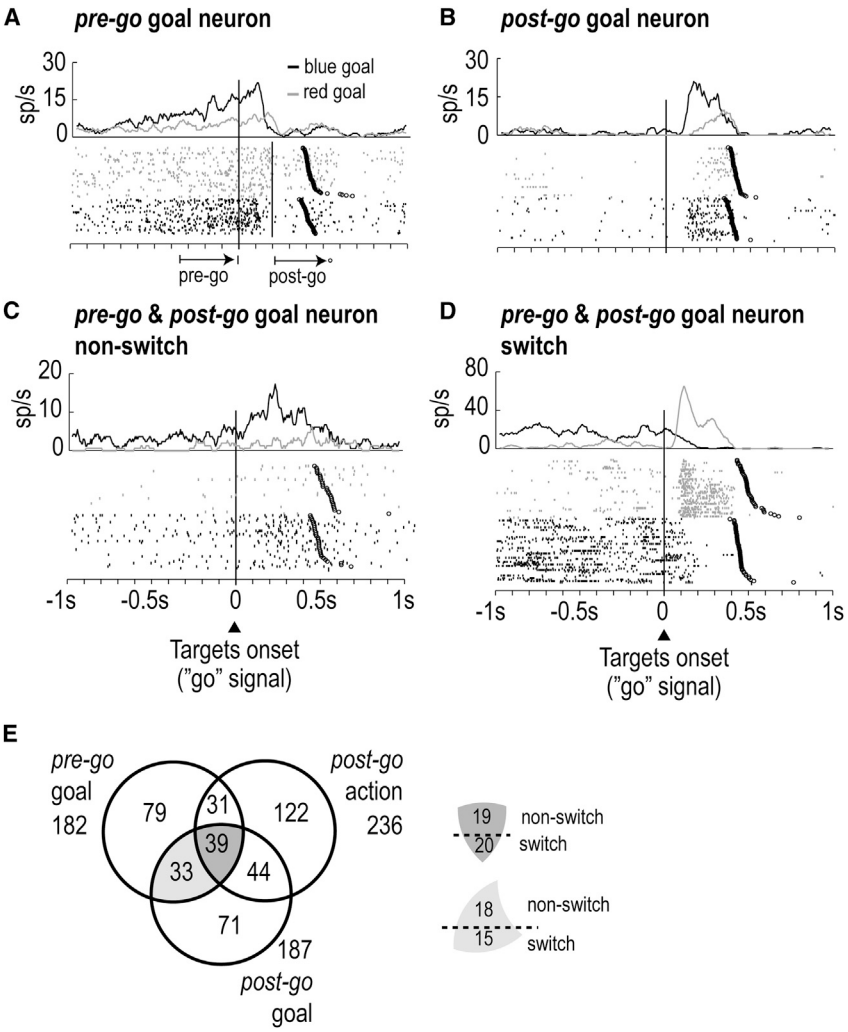

Figure 3. Raster Plot of Four Example Neurons and Venn Diagram of Classified Neurons

(A–D) Top panels inside each figure show mean firing rate activity of the neurons sorted by blue or red goal trials. Bottom panels show the spike times observed during individual trials. Black marker indicates the end of movement.

(A) Neuron showing a higher response during the period before targets onset, but not later, for the blue goal than for the red one.

(B) Neuron with a blue goal preference after targets onset but not before.

(C) Neuron with a preference for the blue goal consistently before and after targets onset.

(D) Neuron switching goal preference after targets onset passing from a blue to a red goal preference. Pre-go and post-go periods are indicated in (A) (bottom arrows) and were used to classify goal-selective neurons.

(E) Number of neurons classified as pre-go goal, post-go goal, or action neurons. Right subpanel: subdivision of the neurons both pre-go and post-go into cells switching and non-switching goal preference.

across the probed brain areas with only one exception: switch neurons were located predominantly in area 8 (Figure S3). This tendency could be associated with the role of this area in the selection between competing visual stimuli (Petrides, 1985).

Preference Switching Due to Neuronal Flexibility

The finding of such a high proportion of switch neurons raised a question about their role in the goal representation and in the goal-action transformation process, as well as whether they reflected any specific neuronal dynamics determining such activity profile. To address these questions, we modeled the above single-neuron experimental evidence relying on two main assumptions. First, we assumed that neurons were embedded in recurrently coupled cell assemblies, and hence their spiking activity reflected the ongoing firing rate of the network they belonged to. The other assumption relied on the observation that preference switching occurred always at an almost fixed time lag from targets onset (167 ± 12 ms, mean ± SEM for the 35 switch neurons; see Figure 3D as an example). This is suggestive of a local activity change driven by a sudden event-triggered variation of the synaptic input. Following these guidelines, we devised and simulated three model networks composed of excitatory and inhibitory integrate-and-fire neurons capable of reproducing the experimental observations but expressing different collective dynamics. These cortical modules were picked up from a continuum of model networks differing only in their excitability, here modulated by the relative strength of their glutamatergic recurrent coupling, similarly to what previously found in monkey premotor cortices (Mattia et al., 2013) (see STAR Methods). Indeed, by increasing such synaptic coupling, the susceptibility of the output firing rate to sudden changes in the input current received from other modules (i.e., areas) raised almost exponentially (Figure 5A, left, black curve). Additionally, the same parameter governed the number of available attractor states (stable activity levels) (Amit and Brunel, 1997; Wang, 1999), such that both single-state and bistable modules could populate the PF network. The same synaptic self-excitation determined also the firing rate variability, giving rise to a non-monotonic trend of the activity fluctuations, which were maximized by a suited level of the strength of the recurrent excitatory synapses (Figure 5A, left, green curve). This maximization of fluctuation size was due to the coexistence of two preferred activity levels randomly visited by the network: a dynamic regime that can be associated with some degree of flexibility of the network, as small input variations can determine relatively large changes of the firing rate. On the other hand, such flexibility can be counterbalanced by the presence of more stable modules such as those with weak and strong self-excitation. Interestingly, in the latter highly excitable module (dark gray circle), such stability underlies the capability to implement a short-term "working" memory (Figure 5A, right), as a strong enough input could

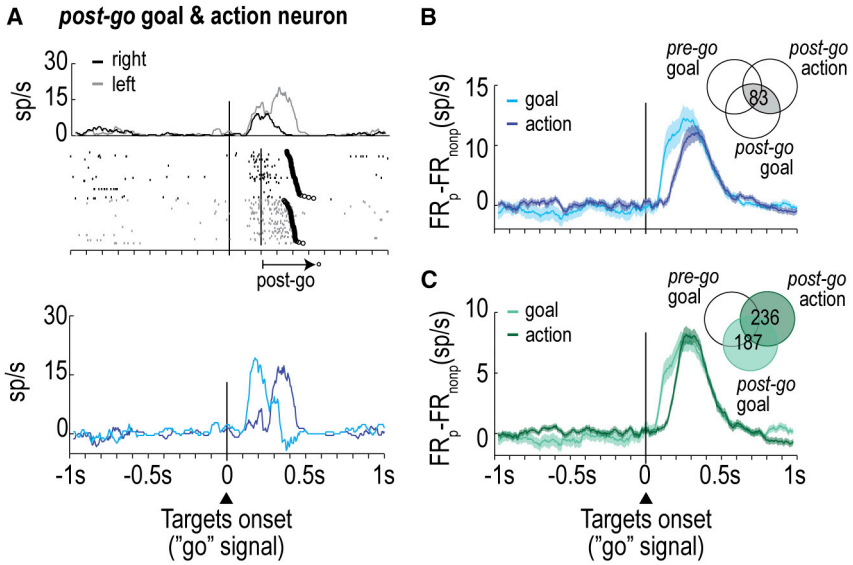

Figure 4. Neural Representation of Goal and Action

(A) Response of the example neuron shown in Figure 3B for action selection. Top: neural response for trials sorted by right (black) and left (gray; preferred) action. Same conventions as in Figure 3. Bottom: activity difference between preferred and non-preferred goals (cyan) and preferred and non-preferred actions (blue).

(B) Activity difference between preferred and non-preferred goal (cyan) and action (blue), averages across neurons selective to both post-go goal and action (n = 83). Shaded areas indicate SEM.

(C) Average activity difference between preferred and nonpreferred goal for neurons that are selective to the post-go goal (n = 187; light green) and between preferred and nonpreferred action for action-selective neurons (n = 236; dark green). Shaded areas indicate SEM.

drive the activity of the module toward a high-firing state persisting even in the absence of the external stimulation (Amit and Brunel, 1997; Wang, 1999). A cortical module such as this would keep trace of the goal during the delay period preceding the "go" signal. Downstream modules in turn could receive as input such goal-related information together with additional synaptic currents due to the onset of other relevant events such as the "go" signal (Figure 5A, middle). In the modeling framework we adopted, a cortical module was composed of several selective pools, as in Amit and Brunel (1997) and Wang (1999), and for simplicity only one pool in each module was involved in the task (see STAR Methods and Figure S4) by receiving from the pools of other modules the input changes depicted in Figure 5.

In this rather simplified multi-modular configuration (Figure 5A, middle), we investigated the pre-go dynamic features of two downstream modules endowed with different levels of excitability, with the idea that this three-modular set is a critical component of the PF network underlying the goal-action transformation (see Discussion for further details). One of these downstream modules had a moderately strong excitatory synaptic reverberation (Figure 5B, dark gray circle), chosen to spontaneously display the large activity fluctuations highlighted in Figure 5A (left) and to be not too far from the bifurcation point (white-gray border). Under suited stimulation conditions, a module such as that would react to small input changes with strong susceptibility because of its intrinsic bistability (i.e., the coexistence of two attractor states at low and high firing rates; Figure 5B, top). Here we modeled the onset of the targets as a sudden positive or negative input variation to be added to the synaptic current provided by the upstream working-memory module (Figure 5B, bottom). In this bistable cortical module, an input variation reshaped the representative energy landscape (Mattia et al., 2013) (Figure 5B, top) such that around the "go" signal the preferred and nonpreferred states (deepest and highest valleys, respectively) were exchanged.

In the second downstream module (Figure 5C, left, light gray circle), glutamatergic recurrent coupling was weaker, bringing to a linearization of the current-to-rate amplification and thus reducing its susceptibility to the input. In this case, only a single valley (i.e., attractor) state is available at a time. With the same weak input variation used in Figure 5B, no switches of firing rates resulted (Figure 5C, left). Moreover, the mean difference $FR_p - FR_{nonp}$ between the pre-go activities related to preferred and nonpreferred goals is expected to be significantly lower in this case (compare firing rate densities in the right insets of Figures 5A, 5B, and 5C, left). In this network, to produce a switch of preference, the variation of the input needed to be considerably stronger (Figure 5C, right). Thus, under weak input modulation, single-state modules are rather stiff neuronal components. Such stiffness is also present in the strong bistable "working-memory" modules (Figure 5A) when trapped in the high-firing state. As a result, under the hypothesis that in the PF network the input to different cell assemblies has similar and relatively small variations in time, only bistable modules would be capable of displaying a preference switching. This leads us to make another prediction: if for correct and incorrect trials the input received by the goal-coding modules is presumably strong and weak, respectively, the response of the modules in these two conditions would be different. During incorrect trials, which are generally those more difficult (Genovesio et al., 2011), the input received by the neural networks might be weak but still sufficient to activate the bistable modules (Figure 5B), similar to what would happen during correct, and less difficult, trials. On the contrary, same weak input might lead only to a moderate modulation of single-state modules (Figure 5C).

A further expectation suggested by this theoretical framework is that in the hypothesized condition of a relatively weak input, a bistable network wanders randomly across its double-well energy landscape, hopping by chance from high- to low-firing states and vice versa (Cao et al., 2016; Durstewitz and Deco, 2008; Litwin-Kumar and Doiron, 2012). As a result, from trial to trial the inter-spike intervals (ISIs) should display a bimodal distribution, with more frequent short and long ISIs (Cao et al., 2016; Latimer et al., 2015; Zipser et al., 1993) accumulating

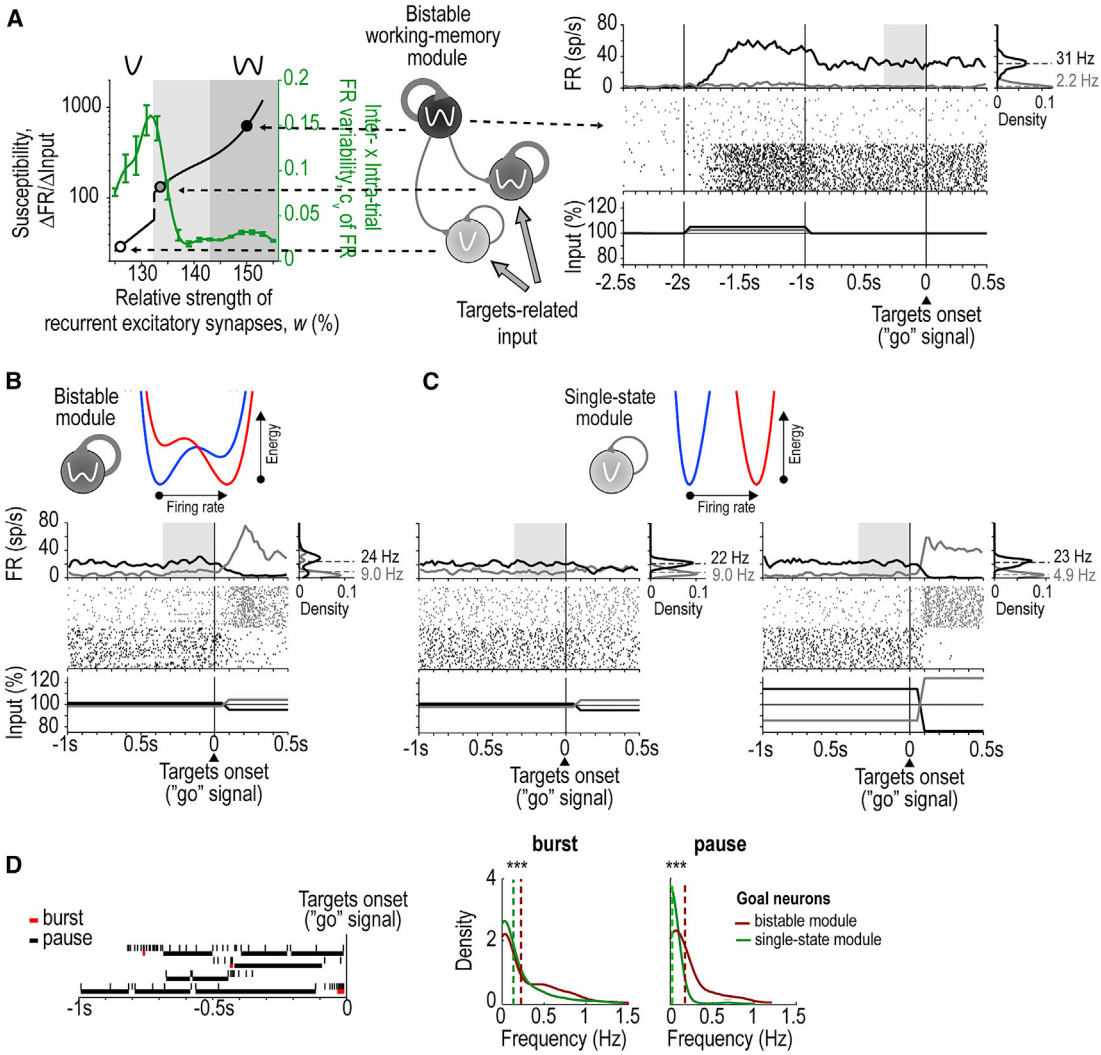

Figure 5. Neuronal Dynamics Underlying Switching of Goal Preference

(A) Left: susceptibility (black curve) to input variations (ΔInput; sudden small change in the spike rate from an external pool of excitatory neurons) of cortical modules with different strength *w* of the excitatory (glutamatergic) synaptic reverberation, and the activity variability measured in simulation as the product of the inter- and intra-trial coefficient of variability c_v (green curve) of the firing rate (FR) under the same stimulation condition (see STAR Methods). Depending on *w*, cortical modules can access only one stable activity state (white region) or can display a bistable dynamics either in the presence (light gray region) or in the absence (dark gray region) of an external stimulation. Small circles only qualitatively represent the excitability level of the cortical modules simulated. Middle: sketch of a small feedforward network of the heterogeneous modules tested in simulations and possibly composing the PF network. Right: working-memory features of the bistable cortical module with maximal synaptic self-excitation (black module). Raster plot and density of the spikes emitted by an example neuron of the network showing the persistence or not of high-firing activity when the input related to a preferred or nonpreferred goal is received before the pre-go delay period (black and gray, respectively). Histograms of firing rates on the right are from the pre-go interval (shaded area) under the two activity conditions. Error bars indicate SD.

(B) A cortical module (dark gray in A) endowed with moderately strong excitatory self-excitation (thick arrow in the top sketch) expressing bistable dynamics with two preferred activity states at low and high firing rates. An increase (decrease) of the received input (bottom) amplifies (reduces) the module excitability, deepening the high (low) activity valley (top right red and blue energy landscapes, respectively). Raster plot and density of the spikes from an example neuron of the network displaying both transitions (black and gray, respectively). Input current changes (bottom) aim at modeling the switching of goal preference observed in experiments (Figure 3D).

(C) The same as (B) for a single-state cortical module (light gray in A). This network has a relatively weak synaptic coupling (thin arrow in top sketch), and its stimulus response is quasi-linear. Under this condition, a weak input change does not result in a switch of goal preference (left). Switching of goal preference similar to (B) is obtained only as a response to a larger input change (right).

(D) Bursts and pauses detection algorithm. Left: an example of burst and pauses detected in four simulated trials. Right: density of bursts and pauses observed within 2 s before the targets onset for 100 simulated neurons randomly sampled from the cortical modules in (B) and (C) (dark red and green, respectively). Bursts and pauses are significantly more often found in switch than in non-switch neurons (Mann-Whitney U test, ***$p < 0.001$; $p = 0.00057$ and $Z = 3.25$ for bursts, $p = 0$ and $Z = 12.4$ for pauses). Dashed lines indicate the means of the distributions.

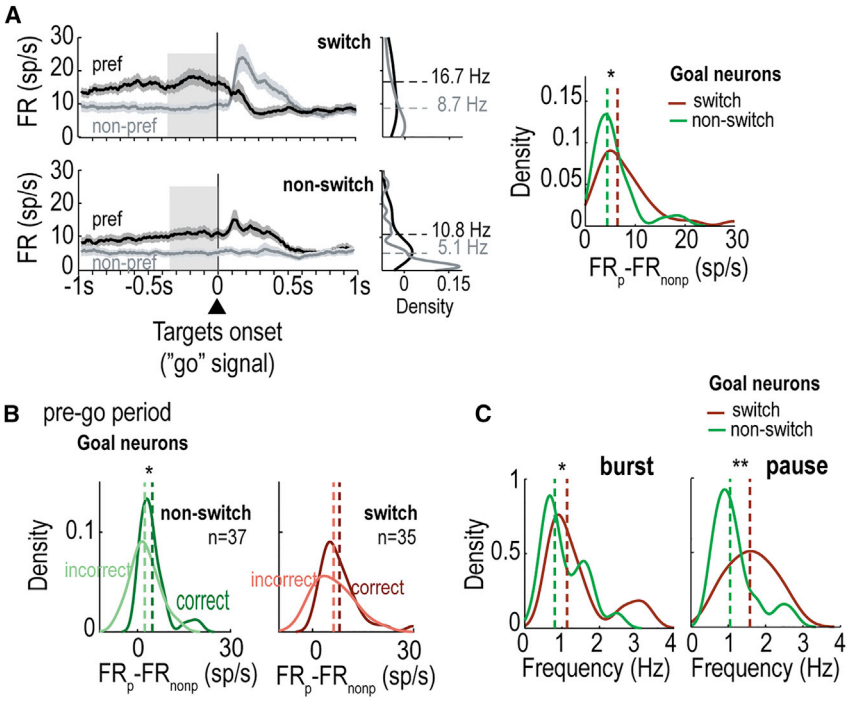

Figure 6. Evidence of Different Neuronal Dynamics Underlying Coding Flexibility

(A) Left: mean firing rate for preferred (black) and non-preferred (gray) goal of switch (n = 35; top) and non-switch neurons (n = 37; bottom panel) aligned to the "go" signal presentation. Shaded areas are SEM. Right: density of average firing rates of each population calculated in the pre-go period (gray area) and differences between preferred and non-preferred conditions computed in the pre-go period. (Mann-Whitney U test, *p < 0.05; p = 0.04 and $Z = 2.01$). Dashed lines indicate the means of the distributions.

(B) Difference between the activity in the preferred and nonpreferred goal conditions calculated in the pre-go period for non-switch and switch neurons, grouped by correct (dark colors) and incorrect (light colors) trials separately (Wilcoxon signed rank test, *p < 0.05; non-switch neurons, p = 0.028, $Z = 2.20$; switch neurons, p = 0.077, $Z = 1.77$).

(C) Frequency of bursts and pauses occurrence within 2 s before targets onset calculated for switch (n = 35; dark red) and non-switch (n = 37; green) neurons. Bursts and pauses are significantly more often found in switch than in non-switch neurons (Mann-Whitney U test, *p < 0.05, **p < 0.01; p = 0.035 and $Z = 2.11$ for bursts, p = 0.001 and $Z = 3.23$ for pauses). Dashed lines indicate the means of the distributions.

when the network was in high- and low-firing states, respectively. Such ISI variability is recognizable in the raster plot in Figure 5B, in which an excess of bursts and pauses is apparent, while they are almost absent in the spiking pattern of neurons of both single-state (Figure 5C) and working-memory (Figure 5A, right) modules. Indeed, for the same weak input, a single energy valley does not allow wide fluctuations of the firing rate (Cao et al., 2016; Mattia et al., 2013), thus leading to almost-Poissonian spike trains in which neither bursts nor pauses outnumber intermediate ISIs (Figure 5D; see STAR Methods).

Evidence of Flexible Dynamics behind Preference Switching

To test the above model predictions, we further inspected our *in vivo* recordings. According to the above theoretical framework, we found that switch neurons had a larger difference of firing rate between preferred and nonpreferred conditions in the pre-go period than the one exhibited by non-switch neurons (Figure 6A). This result supported the hypothesis that non-switch neurons are more likely picked from single-state modules rather than from working-memory modules, as the latter are expected to have the largest difference between preferred and nonpreferred conditions (Figure 5A, right). Next, we used the same measure to compare correct and incorrect trials. In accordance with the model, a significant difference in the activity was found only for non-switch neurons (Figure 6B). Last, we quantified the number of bursts and pauses detected in the activity of switch and non-switch neurons. Consistent with the model, switch neurons exhibited a significantly higher frequency of bursts and pauses in the pre-go period (Figure 6C). This behavior is rather apparent in the example switch neuron shown in Figure 3D, which exhibited ~1.8 times more bursts per second and ~3.6 more pauses per second than the non-switch cell in Figure 3C. Although firing rate and frequency of bursts and pauses are highly correlated (R = 0.76–0.83, with $p < 10^{-4}$ in all cases), firing rate alone could not explain the difference in the frequency of bursts and pauses between the neuronal types (Figure S5). The qualitative resemblance between this experimental evidence and the theoretical framework we propose supports the hypothesis that neurons switching goal preference after targets onset were those more easily activated or damped down. In principle, such susceptibility could facilitate the switching neurons and related modules to have a leading role in determining the transitions between inner mental states, such as the goal-action transformation, thus forming a reservoir of flexible units capable of shaping the whole PF network dynamics.

DISCUSSION

In the present study we examined the evolution of goal coding in PF from its maintenance (memory) to its use in the goal-action transformation process in a distance discrimination task ideally suited to disentangle in time the underlying neuronal computation. We found a peculiar reconfiguration of the network activity after the beginning of the transformation process. At the population level, the encoding of goal in memory abruptly ceased within 125 ms after the goal location was revealed. In the same time window, the goal selectivity emerged in the activity of a new ensemble of neurons, which remained active until the action was performed. Intriguingly, further inspection of the response dynamics at single-neuron level revealed that this network transition from memory to action developed as a

peculiar orchestration between distinct subsets of neurons. Indeed, although the majority of goal-related neurons (225 of 297) were selective either before or after targets onset, others were selective across both task periods, either maintaining (37 of 72) or switching their goal preference (35 of 72). The diversity in their coding scheme balanced out their contribution at the population level, supporting the abrupt population transition. We interpret all these results as a signature of an activity reconfiguration of the PF network because of a transition between different collective states sequentially coding the goal in memory and the goal in the goal-action transformation process.

Such PF activity reconfiguration, which does not disrupt goal information but rather transforms it, seems to be implemented by a heterogeneous cortical network. Goal-related neurons selective during the pre- or post-go period, but not both, do not directly contribute to this passage of information, while switch and non-switch neurons seem to do so relying on two different dynamics: bistable and single-state activity dynamics, respectively. This heterogeneity can result from relatively strong and weak excitatory strength of local synaptic reverberation (Amit and Brunel, 1997; Mattia et al., 2013; Wang, 1999). In this computational framework, the changes in the synaptic input representing the transitions across the different stages of the task can be relatively weak, whereby a sudden reaction can be elicited only in the most excitable modules (the bistable ones). Because of such input susceptibility, these modules are ideally suited to implement the flexibility of the PF network, needed to encode relevant variations of the environment. The least sensitive to such changes of the single-state cortical modules can instead contribute to the stability features expected to be expressed by the PF network. This heterogeneity of the local dynamics of cortical modules would solve the stability-sensitivity conundrum. This solution offers an alternative strategy compared with those in which network stability is driven exogenously by sensory-related input (Churchland et al., 2010; Litwin-Kumar and Doiron, 2012) or by dopaminergic gain modulations (Durstewitz et al., 2000a; O'Reilly, 2006). The reservoir of differently flexible units would also complement the computational advantages expected in presence of some degree of random connectivity between modules needed to implement mixed selectivity (Rigotti et al., 2010, 2013).

Role of Local Heterogeneity and Its Computational Advantages

Given the existence of cortical modules with heterogeneous excitability levels, one may ask about the advantage of having a network composed of such diverse computational units. To answer this question, here we sketch a paradigmatic network of such modules (Figure 7A) capable of autonomously solving the task performed by the monkeys. We tested two example trials having the same "go" signal with the blue and red targets on the left and right sides of the screen, respectively. The trials differ only in the goal to encode: in Figure 7B the blue circle is farther from the reference point than the red square, while in Figure 7C the opposite condition holds.

The working-memory module of the network selectively encodes the "blue circle" goal; that is, it has a preferred response to S2 in trial 7B but not in 7C and keeps it in memory during delay D2 after stimulus removal. This happens thanks to the stability of its high-firing state, such that even in the absence of the S2-related external excitatory input (red empty arrow) encoding the winning goal, a sustained activity reverberation persists (see firing rate in the top row of Figure 7B). Such memory trace is absent in Figure 7C, as the winning goal associated with a weak input (white cross) does not elicit a sufficiently high response of the working-memory module (Zipser et al., 1993; Amit and Brunel, 1997; Wang, 1999). As a result, the state of this network unit eventually biases the activity of the downstream bistable modules in a goal-dependent way. More specifically, during S2 and D2, bottom and top bistable modules have firing rates that strongly correlate and anti-correlate, respectively, with the activity of the upstream working-memory module. This occurs because the connections between the upstream and the bottom and top downstream modules in Figure 7A are chosen to be positive (red) and negative (blue), respectively. Thus, at targets onset, the whole network is in two distinct distributed states: the top bistable module is the only one active when the red goal wins (Figure 7C), whereas the bottom bistable module is the only one active in the case of blue goal selection (Figure 7B).

At targets onset, another stimulus-driven external excitatory input is delivered to the two bistable modules. This external input aims to model the specific targets presentation on the screen and is assumed to be weak, as discussed in the Results, and the same for both trials. The single-state module (Figure 7A, bottom) also receives a similar external input at that time. This module is inactive before targets onset, as the upstream inputs are anti-correlated and cancel each other because of the positive and negative feedback they provide. The single-state module is not very susceptible, and its reaction to the weak external input is almost imperceptible. However, in bistable modules, the same external input primes a chain reaction guided by their high susceptibility. For instance, in Figure 7B, the external input destabilizes the metastable low-activity state of the top bistable unit and a sharp upward transition occurs, followed by a dampening of the overshoot because of spike-frequency adaptation. Consequently, because of the inhibitory connections between the bistable modules, the bottom one is destabilized, and a sharp transition toward the low-activity state follows. This switch of activity levels and the presence of the external excitatory input eventually lead to a net increase of the synaptic input received by the single-state module, which is now large enough to elicit a significant increment of its firing rate. In Figure 7C, a similar chain of activity changes occurs. Here, the transformation is primed by the low to high firing rate transition in the bottom bistable module. This leads to an inactivation of the other bistable module and eventually to a significant increase of the firing rate of the single-state module. In this trial, the activity level reached by the latter module is lower than the one reached when the winning goal is "blue" (Figure 7B).

These chains of local transitions occurring in both trials (Figures 7B and 7C, red and blue dotted vertical arrows), and elicited by the same "go" signal, lead to different reconfigurations of the multi-modular network activity. This can be read out as two different movements once combined with the information in

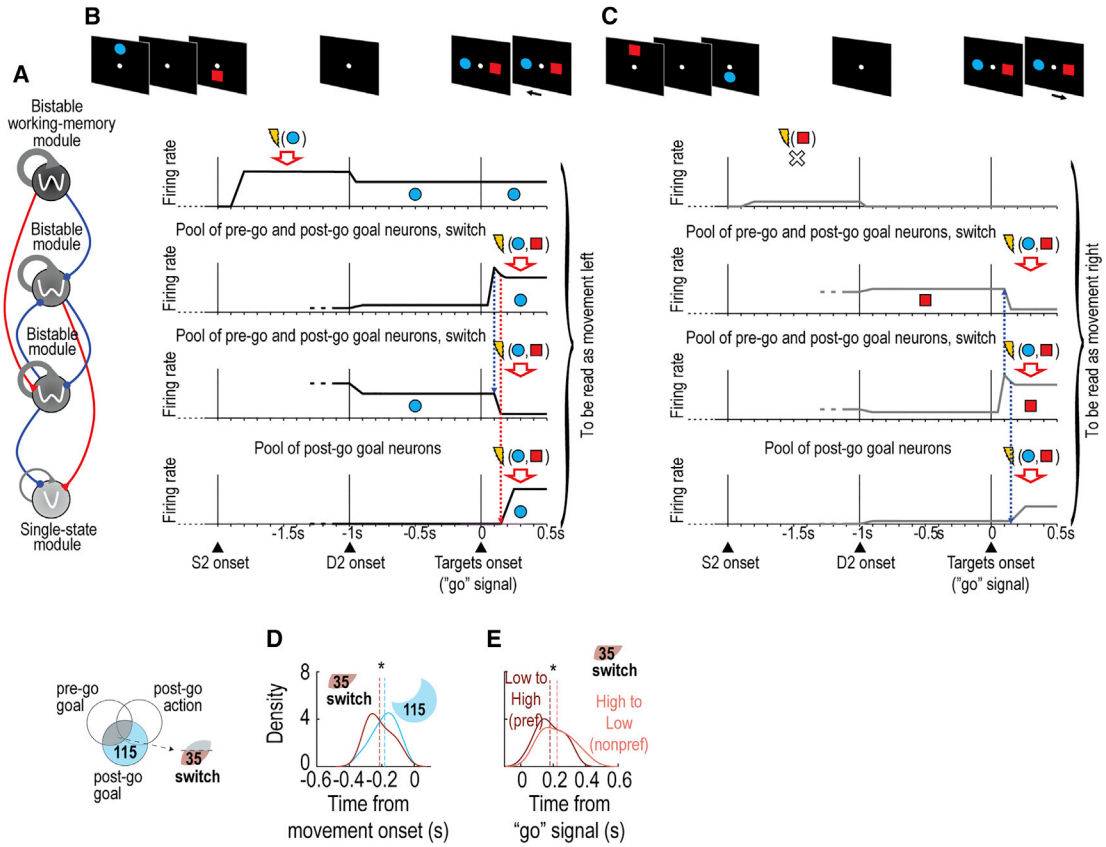

Figure 7. Plausible Neural Network of Heterogeneous Modules and Time Latency of Their Activation and Inactivation

(A) A proposal of a simple network of heterogeneous modules capable to reproduce the experimental data. Blue and red lines represent inhibitory and excitatory connections, respectively.

(B and C) Two example of trials with blue (B) or red (C) stimulus as goal. Red empty arrows represent external excitatory inputs associated to different task-related stimuli as the goal to encode and target to touch (stimulus type is represented by the bracket content following lightning). Different rows schematically represent the expected time course of the firing rates of the four modules on the left. Vertical blue and red dotted lines, inhibitory and excitatory events eliciting transitions in other modules.

(D) Goal coding latency of post-go goal, but not pre-go goal, neurons (n = 115; cyan) and switch neurons (n = 35; red) calculated with the activity aligned to movement onset (Mann-Whitney U test, *p < 0.05; p = 0.038, $Z = 2.08$).

(E) Transition times of the activity level (high and low, preferred and nonpreferred goal, respectively) in switch neurons (n = 35). Low-to-high and high-to-low transitions are dark and light red distributions, respectively (Wilcoxon signed rank test, *p < 0.05; p = 0.029, $Z = 2.19$).

the "go" signal about the position of the targets: toward the left (Figure 7B) and the right (Figure 7C) when the selected goal is "blue" and "red," respectively. The post-go action neurons should be part of the network reading out the information encoded by this newborn neuronal representation. Thus, in this theoretical framework it should not be surprising to see in Figures 4B and 4C that the activation of post-go goal neurons on average anticipates the transition time occurring in the post-go action neurons.

To conclude, the multi-modular network we sketched has the ability to autonomously translate the selected goal into the correct action to perform, without the help of other modules from different brain areas. Furthermore, the heterogeneity of the cortical modules found in our data implies a hierarchical organization of the goal-action transformation determined by the excitability of the involved units. Indeed, in the chain of local transitions we predict, the firsts are those happening in the more susceptible bistable modules found in a metastable low-activity state. Instead, the lasts are the activations of the less excitable single-state modules. Remarkably, we found such excitability-driven hierarchy of time lags also in our data. On average, transition times in switch (bistable) neurons preceded by 31 ms the activation of post-go, but not pre-go, goal (single-state) neurons (Figure 7D). Furthermore, low-to-high post-go transitions in switch neurons, on average, occurred before the inactivation (high-to-low transitions) in the same group of neurons (Figure 7E). Interestingly, a similar fine structure of the cortical network reconfiguration has been also found in dorsal premotor cortex of monkeys performing a motor decision task (Mattia et al., 2013). This seems to suggest a common computational strategy adopted by cortical networks to perform global state changes in which a subset of cortical modules, the bistable ones, can play a pivotal role in shaping the whole PF network dynamics, even if the input is small.

Coding Flexibility and Coding Independence

Our study extends the understanding of the flexible goal coding in PF (Fuster and Alexander, 1971; Hoshi et al., 2000; Rainer et al., 1999). Previous studies have shown that PF neurons change their selectivity from retrospective visual stimulus to prospective goal objects (Rainer et al., 1999) and from visual cues to object goal signals (Hoshi et al., 2000), but neither study reported switches of preference. We had previously assumed that the neurons coding prospectively the goal lead directly to the action by maintaining the same goal selectivity in the action period (Genovesio et al., 2009). Contrary to such assumption, by examining the goal coding at the time of the goal-action transformation, we observed a remarkable dissociation between goal-maintenance and goal-action transition.

It has been shown that the same prefrontal neurons can encode multiple types of information even within the same task period (Cai and Padoa-Schioppa, 2014; Genovesio et al., 2005; Hoshi et al., 2000; Yamagata et al., 2012) and use different coding schemes depending on the task (Rao et al., 1997). With our task, we present a further type of coding flexibility by investigating the coding dynamics of exactly the same information, as represented by the goal in different task periods, and showing how it can flexibly change as a result of a reconfiguration of the network activity. The goal representation is transmitted to a new neuronal ensemble, as in a relay race when the baton is passed between runners. Information transmission between subnetworks can be seen as an effective way to implement working memory, similar to the sequence-based circuit dynamics recently found in rodent parietal cortex (Harvey et al., 2012).

Although PF coding flexibility is well documented, it is still not clear what can generate it or to what degree it must be considered a general attribute of PF neurons. In a recent study, Yamagata et al. (2012) investigated the transformation of the relative position to choose one of two targets into a planned action. The results showed that different group of neurons coded the relative position before and after the targets were presented. However, it was not clear whether the dissociation represented a special transition or if, in contrast, it simply reflected the passage of time or the appearance of a new event (Marcos et al., 2016).

Our study shows that the neurons coding the goal in memory and the neurons involved in the goal-action transformation are similarly distributed in the prefrontal areas of our recordings. However, by dividing switch and non-switch neurons, we observed a tendency for the switch neurons to be located especially in area 8. Monkeys with lesions in the periarcuate cortex, including areas 8 and 6 but not 46, fail to choose one of two visual stimuli on the basis of a conditional rule (Petrides, 1985), and inactivation of the frontal eye field (FEF) generated a deficit to direct the saccades to a location instructed by the color of a cue (Keller et al., 2008). In these experiments, a stimulus instructed the future goal to select in the presence of two or more alternatives. Considering our results in the context of these neuropsychological studies, the presence of switch neurons predominantly in area 8 might have a key role in the process of visual stimulus selection on the basis of the goal color (the farthest visual stimulus) (Passingham and Wise, 2012).

Goal-Action Transformation

One of the main results we report is the fundamental role of the post-go goal coding in the goal-action transformation. However, we cannot identify the specific function carried out by the neurons, that is, whether it represented attention for action (Lau et al., 2004; Lebedev et al., 2004), target selection (Hoshi et al., 2000), motor planning (Marcos et al., 2015; Takeda and Funahashi, 2004), or confidence about the decision (Marcos et al., 2013). The action usually involved congruent eye and hand movements, and thus we could not investigate the possible interaction between the two in the signal carried out by the neurons in PF (Thura et al., 2008).

We have shown that a network of cortical modules, with heterogeneous sensitivity to the input, gives rise to an augmented dynamic richness that allows simultaneously for both change of the global network state as well as retention of the information about previous events. This can be achieved without the need of unspecific excitability modulation because of the dopaminergic system (Durstewitz et al., 2000a; O'Reilly, 2006; Seamans and Yang, 2004). Nevertheless, one could argue that activity fluctuations of switching neurons in bistable modules are due to fluctuations of the input they receive. In the theoretical framework we propose (Figures 5 and 7), these input changes could be provided only from outside the probed PF network, as single-state modules are passive units unable to autonomously produce activity fluctuations. This would imply that both bistable and single-state modules fluctuate in the same way, driven by the same external input changes. However, our data show that this is not the case, further supporting the hypothesis that activity fluctuations before goal-action transformation are endogenously produced.

Altogether our findings show that goal information kept in working memory did not directly lead to action, hence pointing away from a unitary view of goal representation. Rather, the representation of the "goal in memory" is reconfigured into another network activity, which is perhaps specifically suited for being read out to produce specific actions (i.e., "goal for action"). Recently, Brincat and Miller (2016) suggested a shift of PF networks from external to internal modes along with the progress of associative learning. Our findings may concur with this shift but differ in at least two ways: (1) the opposite direction (i.e., shifting from internal to external stages) and (2) with a shorter time window (i.e., within a trial). The PF may have a general scheme of reconfiguration between external and internal coding networks supported by a mix of single-state and bistable cell assemblies, thereby allowing us to adapt to a complex environment.

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 SUBJECTS DETAILS
- METHODS DETAILS
  - Surgery and data collection
  - Behavioral task

- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Neural analyses
  - Additional classification of neurons
  - Goal selectivity index
  - Models and simulations
  - Statistical analysis

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.celrep.2019.05.021.

ACKNOWLEDGMENTS

We thank Paolo Del Giudice for comments on the manuscript. This work was supported by Division of Intramural Research of the National Institute of Mental Health (grant Z01MH-01092), by Sapienza University of Rome (grant "Avvio alla Ricerca 2016"), and by a Juan de la Cierva-Incorporación scholarship (IJCI-2016-27864 from the Spanish Ministry of Science, Innovation and Universities).

AUTHOR CONTRIBUTIONS

A.G. and S.T. collected the data. E.M. and A.G. designed the study and analyzed the data. M.M. developed the model and designed the model-driven analyses. E.M. and M.M. prepared the figures. All authors interpreted the results and wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 20, 2017
Revised: March 10, 2018
Accepted: May 3, 2019
Published: June 4, 2019

SUPPORTING CITATIONS

The following references appear in the Supplemental Information: Jolliffe (1986); Petrides and Pandya (2002).

REFERENCES

Amit, D.J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex *7*, 237–252.

Averbeck, B.B., Sohn, J.W., and Lee, D. (2006). Activity in prefrontal cortex during dynamic selection of action sequences. Nat. Neurosci. *9*, 276–282.

Brincat, S.L., and Miller, E.K. (2016). Prefrontal cortex networks shift from external to internal modes during learning. J. Neurosci. *36*, 9739–9754.

Cai, X., and Padoa-Schioppa, C. (2014). Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. Neuron *81*, 1140–1151.

Cao, R., Pastukhov, A., Mattia, M., and Braun, J. (2016). Collective activity of many bistable assemblies reproduces characteristic dynamics of multistable perception. J. Neurosci. *36*, 6957–6972.

Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat. Neurosci. *13*, 369–378.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. *18*, 193–222.

Durstewitz, D., and Deco, G. (2008). Computational significance of transient dynamics in cortical networks. Eur. J. Neurosci. *27*, 217–227.

Durstewitz, D., Seamans, J.K., and Sejnowski, T.J. (2000a). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. J. Neurophysiol. *83*, 1733–1750.

Durstewitz, D., Seamans, J.K., and Sejnowski, T.J. (2000b). Neurocomputational models of working memory. Nat. Neurosci. *3* (*Suppl*), 1184–1191.

Falcone, R., Brunamonti, E., Ferraina, S., and Genovesio, A. (2016). Neural encoding of self and another agent's goal in the primate prefrontal cortex: human-monkey interactions. Cereb. Cortex *26*, 4613–4622.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652–654.

Genovesio, A., and Ferraina, S. (2014). The influence of recent decisions on future goal selection. Philos. Trans. R. Soc. Lond. B Biol. Sci. *369*, 20130477.

Genovesio, A., and Tsujimoto, S. (2014). From duration and distance comparisons to goal encoding in prefrontal cortex. Adv. Exp. Med. Biol. *829*, 167–186.

Genovesio, A., Brasted, P.J., Mitz, A.R., and Wise, S.P. (2005). Prefrontal cortex activity related to abstract response strategies. Neuron *47*, 307–320.

Genovesio, A., Brasted, P.J., and Wise, S.P. (2006). Representation of future and previous spatial goals by separate neural populations in prefrontal cortex. J. Neurosci. *26*, 7305–7316.

Genovesio, A., Tsujimoto, S., and Wise, S.P. (2008). Encoding problem-solving strategies in prefrontal cortex: activity during strategic errors. Eur. J. Neurosci. *27*, 984–990.

Genovesio, A., Tsujimoto, S., and Wise, S.P. (2009). Feature- and order-based timing representations in the frontal cortex. Neuron *63*, 254–266.

Genovesio, A., Tsujimoto, S., and Wise, S.P. (2011). Prefrontal cortex activity during the discrimination of relative distance. J. Neurosci. *31*, 3968–3980.

Genovesio, A., Tsujimoto, S., and Wise, S.P. (2012). Encoding goals but not abstract magnitude in the primate prefrontal cortex. Neuron *74*, 656–662.

Genovesio, A., Tsujimoto, S., Navarra, G., Falcone, R., and Wise, S.P. (2014a). Autonomous encoding of irrelevant goals and outcomes by prefrontal cortex neurons. J. Neurosci. *34*, 1970–1978.

Genovesio, A., Wise, S.P., and Passingham, R.E. (2014b). Prefrontal-parietal function: from foraging to foresight. Trends Cogn. Sci. *18*, 72–81.

Glass, G.V., Peckham, P.D., and Sanders, J.R. (1972). Consequences of failure to meet the assumptions underlying the fixed effects analysis of variance and covariance. Rev. Educ. Res. *42*, 237–288.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62–68.

Heidbreder, C.A., and Groenewegen, H.J. (2003). The medial prefrontal cortex in the rat: evidence for a dorso-ventral distinction based upon functional and anatomical characteristics. Neurosci. Biobehav. Rev. *27*, 555–579.

Hoshi, E., Shima, K., and Tanji, J. (2000). Neuronal activity in the primate prefrontal cortex in the process of motor selection based on two behavioral rules. J. Neurophysiol. *83*, 2355–2373.

Hussar, C.R., and Pasternak, T. (2009). Flexibility of sensory representations in prefrontal cortex depends on cell type. Neuron *64*, 730–743.

Jolliffe, I.T. (1986). Principal Component Analysis (Springer).

Keller, E.L., Lee, K.M., Park, S.W., and Hill, J.A. (2008). Effect of inactivation of the cortical frontal eye field on saccades generated in a choice response paradigm. J. Neurophysiol. *100*, 2726–2737.

Ko, D., Wilson, C.J., Lobb, C.J., and Paladini, C.A. (2012). Detection of bursts and pauses in spike trains. J. Neurosci. Methods *211*, 145–158.

Latimer, K.W., Yates, J.L., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2015). Neuronal modeling. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science *349*, 184–187.

Lau, H.C., Rogers, R.D., Haggard, P., and Passingham, R.E. (2004). Attention to intention. Science *303*, 1208–1210.

Lebedev, M.A., Messinger, A., Kralik, J.D., and Wise, S.P. (2004). Representation of attended versus remembered locations in prefrontal cortex. PLoS Biol. *2*, e365.

Litwin-Kumar, A., and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. Nat. Neurosci. *15*, 1498–1505.

Marcos, E., and Genovesio, A. (2017). Interference between space and time estimations: from behavior to neurons. Front. Neurosci. *11*, 631.

Marcos, E., Pani, P., Brunamonti, E., Deco, G., Ferraina, S., and Verschure, P. (2013). Neural variability in premotor cortex is modulated by trial history and predicts behavioral performance. Neuron *78*, 249–255.

Marcos, E., Cos, I., Girard, B., and Verschure, P.F.M.J. (2015). Motor cost influences perceptual decisions. PLoS ONE *10*, e0144841.

Marcos, E., Tsujimoto, S., and Genovesio, A. (2016). Event- and time-dependent decline of outcome information in the primate prefrontal cortex. Sci. Rep. *6*, 25622.

Markowitz, D.A., Curtis, C.E., and Pesaran, B. (2015). Multiple component networks support working memory in prefrontal cortex. Proc. Natl. Acad. Sci. U S A *112*, 11084–11089.

Mattia, M., and Del Giudice, P. (2000). Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. Neural Comput. *12*, 2305–2329.

Mattia, M., Pani, P., Mirabella, G., Costa, S., Del Giudice, P., and Ferraina, S. (2013). Heterogeneous attractor cell assemblies for motor planning in premotor cortex. J. Neurosci. *33*, 11155–11168.

McCormick, D.A. (2005). Neuronal networks: flip-flops in the brain. Curr. Biol. *15*, R294–R296.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. *24*, 167–202.

O'Reilly, R.C. (2006). Biologically based computational models of high-level cognition. Science *314*, 91–94.

Passingham, R.E., and Wise, S.P. (2012). The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight (Oxford University Press).

Petrides, M. (1985). Deficits in non-spatial conditional associative learning after periarcuate lesions in the monkey. Behav. Brain Res. *16*, 95–101.

Petrides, M., and Pandya, D.N. (2002). Comparative cytoarchitectonic analysis of the human and the macaque ventrolateral prefrontal cortex and corticocortical connection patterns in the monkey. Eur. J. Neurosci. *16*, 291–310.

Rainer, G., Rao, S.C., and Miller, E.K. (1999). Prospective coding for objects in primate prefrontal cortex. J. Neurosci. *19*, 5493–5505.

Rao, S.C., Rainer, G., and Miller, E.K. (1997). Integration of what and where in the primate prefrontal cortex. Science *276*, 821–824.

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.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Saito, N., Mushiake, H., Sakamoto, K., Itoyama, Y., and Tanji, J. (2005). Representation of immediate and final behavioral goals in the monkey prefrontal cortex during an instructed delay period. Cereb. Cortex *15*, 1535–1546.

Seamans, J.K., and Yang, C.R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. Prog. Neurobiol. *74*, 1–58.

Sigala, N., Kusunoki, M., Nimmo-Smith, I., Gaffan, D., and Duncan, J. (2008). Hierarchical coding for sequential task events in the monkey prefrontal cortex. Proc. Natl. Acad. Sci. U S A *105*, 11969–11974.

Squire, R.F., Noudoost, B., Schafer, R.J., and Moore, T. (2013). Prefrontal contributions to visual selective attention. Annu. Rev. Neurosci. *36*, 451–466.

Takeda, K., and Funahashi, S. (2004). Population vector analysis of primate prefrontal activity during spatial working memory. Cereb. Cortex *14*, 1328–1339.

Tanji, J., and Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. Physiol. Rev. *88*, 37–57.

Thura, D., Hadj-Bouziane, F., Meunier, M., and Boussaoud, D. (2008). Hand position modulates saccadic activity in the frontal eye field. Behav. Brain Res. *186*, 148–153.

Tsujimoto, S., Genovesio, A., and Wise, S.P. (2008). Transient neuronal correlations underlying goal selection and maintenance in prefrontal cortex. Cereb. Cortex *18*, 2748–2761.

Wang, X.-J. (1999). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J. Neurosci. *19*, 9587–9603.

Wilson, F.A.W., Scalaidhe, S.P., and Goldman-Rakic, P.S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. Science *260*, 1955–1958.

Yamagata, T., Nakayama, Y., Tanji, J., and Hoshi, E. (2012). Distinct information representation and processing for goal-directed behavior in the dorsolateral and ventrolateral prefrontal cortex and the dorsal premotor cortex. J. Neurosci. *32*, 12934–12949.

Zhou, X., Qi, X.-L., and Constantinidis, C. (2016). Distinct roles of the prefrontal and posterior parietal cortices in response inhibition. Cell Rep. *14*, 2765–2773.

Zipser, D., Kehoe, B., Littlewort, G., and Fuster, J. (1993). A spiking network model of short-term active memory. J. Neurosci. *13*, 3406–3420.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Primate: Rhesus macaque	NIH	N/A
Software and Algorithms		
MATLAB	MathWorks	N/A
Custom-made spiking neural network simulator (MATLAB, C)	Mattia and Del Giudice, 2000	N/A
Cortex	NIMH Cortex	N/A
Off Line Sorter	Plexon	N/A

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, Aldo Genovesio (aldo.genovesio@uniroma1.it).

EXPERIMENTAL MODEL AND SUBJECTS DETAILS

Two adult male rhesus monkeys (*Macaca mulatta*), 8.5 and 8.0 kg, served as subjects in this study. All procedures followed the Guide for the Care and Use of Laboratory Animals (1996, SBN 0-309-05377-3) and were approved by the National Institute of Mental Health Animal Care and Use Committee.

METHODS DETAILS

Surgery and data collection

Recording chambers were implanted over the exposed dura mater of the left frontal lobe, along with head restraint devices, using aseptic techniques and isofluorane anesthesia (1–3%, to effect). Monkey 1 was implanted with two 18-mm-diameter chambers, one placed over the caudal PF cortex, the other over the dorsolateral PF cortex; Monkey 2 had a single 27 × 36 mm chamber encompassing both areas (Figure 1B). No monkeys had received prior surgery. The animals were individually or pair-housed and kept on a 12 h/12 h light/dark cycle. The experiments were conducted during the light cycle of the day.

We used an infrared oculometer (Arrington Recording) to record eye position. We recorded single-cell activity using quartz insulated platinum-iridium electrodes (0.5–1.5M_at 1 kHz) positioned by a 16-electrode drive assembly (Thomas Recording). The electrodes were arranged in a concentric recording head with 518 mm spacing. We discriminated single-unit potentials online with the Multichannel Acquisition Processor (Plexon) and confirmed each isolated waveform carefully using the Off Line Sorter (Plexon). The offline verification of unit isolation was based on principal component analysis, minimal interspike intervals, and clearly differentiated waveforms inspected individually for every isolated neuron. Inadequately isolated potentials were eliminated from the dataset before performing any additional analysis. Eye position was not placed under experimental control because of the tendency of both monkeys to saccade to each stimulus when it appeared.

Near the end of recordings, we made electrolytic marking lesions (15 μA for 10 s). Ten days later, the monkeys were deeply anesthetized and perfused through the heart with 10% formol saline. After sectioning the brain and staining the section for Nissl substance, we plotted the recording sites by reference to the marking lesions, pins inserted at the time of the perfusion, and structural magnetic resonance images (MRI) taken periodically before and between recording sessions.

Behavioral task

Neurons from PF were recorded from two male rhesus monkeys while they performed a spatial discrimination task (Figure 1A; see Genovesio et al., 2011). The neurons were predominantly recorded from area 8, area 46 and a small population of area 12 (Figure 1B). Monkeys were required to select from two visual stimuli, sequentially displayed on the screen, the one that had been presented farther from the screen's central point. They sat on a primate chair 29 cm from a screen, with their head fixed. Three 3x2 cm infrared switches were used as an interface between the monkeys and the experimental task. Each trial started when the monkeys touched with their left hand the central switch that led to the presentation of a central reference point for 400-800 ms after which a first stimulus (S1), either a blue circle or a red square, appeared for 1.0 s above or below the reference point. After a first delay (D1) of 400 ms or

800 ms, a second stimulus (S2) appeared for another 1.0 s. S2 was a blue circle if S1 was a red square and a red square otherwise. If the first stimulus was presented above the reference point, the second was presented below and vice versa when the first was below. A second delay (D2) of either 0 ms, 400 ms or 800 ms followed S2 offset and preceded the simultaneous representation of the two stimuli as targets, in this case acting as potential goals, which served as a "go" signal. The monkeys had to select, within 6.0 s, the one that was presented farther from the reference point by pressing either the left or the right switch based on the position of the chosen goal on the screen (7.8° to the left or to the right of the reference point). The distance between the reference point and S1 and S2 ranged from 8 mm to 48 mm, in steps of 8 mm. This range is equivalent to values of 1.6° to 9.4° of visual angle. Correct responses were rewarded with 0.1 mL of juice whereas incorrect responses were followed by an acoustic signal. An intertrial interval of 700–1000 ms separated the end and the start of two consecutive trials. All the variables of the experiment, such as color and shape of the stimuli and their left or right position during their reappearance, were pseudorandomly assigned. On average, there were 37.69 trials (±0.44, SEM) with the blue goal as the correct choice and 35.20 trials (±0.43, SEM) with the red goal as the correct choice.

QUANTIFICATION AND STATISTICAL ANALYSIS

Neural analyses

From the original database (Genovesio et al., 2011), we selected the neurons that had at least a mean of 1 spike/s when considering the time period between the beginning of the trial and the end of the movement onset (N = 814/974).

To select the neurons that encoded the chosen goal (red square or the blue circle) in the D2 period before targets onset, we considered a *pre-go* period from 350-0 ms before targets onset. We performed a one-way ANOVA using the red and blue goals as factor. Same analysis but in the period between 200 ms after targets onset until movement end, designated as *post-go* period, allowed us to identify the neurons classified as *post-go* goal neurons. A one-way ANOVA with action direction (left and right) in the *post-go* period as a factor was performed to identify the neurons modulated by action selection. Unless otherwise specified the analyses did not include trials with D2 of 0 ms duration and were performed on correct trials. An alpha level of 0.05 was used for all ANOVAs.

In all cases, the preferred goal condition was identified as the one with the highest mean activity and the nonpreferred condition as the other one. The average neural firing rate of the different neural populations was calculated as the mean firing rate of each individual mean neural activity. We used a temporal window of 50 ms and a sliding window of 5 ms. The significance test (Wilcoxon's matched-pairs signed-rank test with Bonferroni correction) between preferred and nonpreferred conditions was computed using a non-overlapping window of 50 ms. The time reported refers to the middle point of the temporal window in which the statistical test was performed.

Bursts and pauses were detected using the Robust Gaussian Surprise (RGS) method (Ko et al., 2012) which simultaneously searches for burst and pause in a given spike train. In brief, the RGS method identifies the interspike intervals within the spike trains of an individual neuron that are significantly ($p < 0.05$) lower (burst) or higher (pause) than a central distribution calculated with the pooled data of all trials of the neuron. To calculate the central distribution, we divided the data by preferred and nonpreferred conditions.

For individual neurons, we estimated the goal coding latency as the time at which the mean activity difference between preferred and nonpreferred goal conditions reached the 20% of the maximum difference within the period of analysis. The goal coding latencies of the *post-go* goal, but not *pre-go* goal, neurons (n = 115) and the switch neurons (n = 35) were calculated using the neural activity observed in the period from the "go" signal to movement onset with the data aligned to movement onset. We restricted the analyses to only those neurons that were recorded during more than 20 trials with blue and red goal conditions (73% of post-go goal, but not pre-go goal, neurons and 90% of switch neurons) in order to ensure a good estimate of the neural activity difference between conditions. We used a temporal window of 50 ms moving in steps of 5 ms. These same windows were used to plot the mean firing rate activity of individual neurons.

Besides, we estimated the time at which the mean activity of switch neurons for the preferred goal condition changed from a low to a high level and the time at which the activity of the same neurons for the nonpreferred condition changed from a high to a low level. To do so, we calculated the time when the activity reached 80% of its maximum and when the mean activity of a neuron fell to the 20% of its maximum, respectively. In both cases, we used the activity observed during the RT period with the data aligned to the "go" signal.

Additional classification of neurons

Apart from the neural classification presented in the main text, we further identified two additional groups of neurons to control the main results: neurons coding the goal in the *early pre-go* period (from 750 ms to 400 ms before targets onset) and neurons coding the goal in both the *pre-S2* offset (from 350 ms to S2 offset) and the *post-S2* offset (from 50 ms to 400 ms after S2 offset). We used the neural activity observed in those periods to perform a One-way ANOVA using the red and blue goals as factor.

Goal selectivity index

The goal selectivity index (GI) of each neurons was calculated as:

$$GI = \frac{\overline{N}_{blue} - \overline{N}_{red}}{\overline{N}_{blue} + \overline{N}_{red}}$$

where $\overline{N}$ is the mean count of spikes for blue goal or red goal selection conditions. The GI is a value within −1 and 1, with 0 indicating non-selectivity. The GI values were calculated using the neurons coding the goal in the *pre-go* and *post-go* periods (n = 72), in the *early pre-go* and *pre-go* periods (n = 96) and in the *pre-S2* offset and *post-S2* offset periods (n = 102).

Models and simulations

Simulated cortical modules were networks of leaky integrate-and-fire (LIF) neurons similar to those introduced in (Cao et al., 2016; Mattia et al., 2013). Briefly, a cortical module was composed of 1000 LIF neurons (80% excitatory and 20% inhibitory) fully connected. As stationary baseline background noise, the neurons received a Poissonian spike train at 2400 Hz frequency (modeling the firing of 800 external neurons firing at rate FR_{ext} = 3 Hz), which was further modulated in order to simulate the different task conditions (see below) and to affect neuronal membrane potential via synapses with efficacy $J_{E,ext}$ = 0.429 mV and $J_{I,ext}$ = 0.560 mV (E and I stay for excitatory and inhibitory post-synaptic neurons, respectively). Membrane potential decay constants were τ_E = 20 ms and τ_I = 10 ms, and refractory periods were 2 ms and 1 ms for excitatory and inhibitory neurons, respectively. Emission thresholds V_{thr} = 20 mV and reset potential after spike emission V_{res} = 15 mV were the same for both neuron types. Spike-frequency adaptation was included in excitatory neurons, as they received an after hyper-polarizing (AHP) current increasing by 20 mV/s and 90 mV/s at every emission of a spike in the networks of Figures 6A–6C, respectively, eventually decaying to 0 mV/s with a time constant τ_{AHP} = 100 ms.

Excitatory neurons were structured in 8 equally sized pools of 100 cells each (see inset in Figure S4), by setting the intrapool synaptic efficacy between neurons to J_{intra} = {0.652, 0.644, 0.541} mV and the inter-pool one to J_{inter} = {0.397, 0.397, 0.413} mV in the bistable, working-memory and single-state modules, respectively. Although the former bistable model had a stronger synaptic self-excitation (J_{intra}) than the working-memory one, its excitability level was smaller due to a larger aforementioned self-inhibitory (AHP) current (90 mV/s and 20 mV/s, respectively). Indeed, network excitability can be reduced either by decreasing the glutamatergic synaptic strength or by increasing the ionotropic inhibitory currents. We chose a mixed strategy (i.e., an increased self-inhibition to change working-memory modules into weakly bistable ones, and a reduced synaptic self-excitation to transform the weakly bistable modules into the single-state ones) as a stronger activity-dependent AHP current leads to an overshoot of activity in response to a sudden increase of the external input both in bistable and single-state modules, a temporal activity profile that better matches the post-go firing rate of many recorded neurons (see for instance Figures 2 and 3). The settings for the simulated bistable and working-memory modules allowed us to have a winner-take-all network in which an unspecific inactive Down state (all neurons fire at low firing rate) coexisted with other 8 stable Up states where only the neurons belonging to the same pool reverberated at a relatively high firing rate (Amit and Brunel, 1997; Wang, 1999). The other weaker synaptic couplings allowed us to set up a network with only one attractor state at a time, as shown in Figure S4A. All the networks had instantaneous synaptic transmission from excitatory to inhibitory neurons with an average efficacy J_{IE} = 0.560 mV, while those from inhibitory neurons were $J_{EI} = J_{II} = -1.50$ mV. Other network settings were as in Mattia et al. (2013).

To produce the activity patterns emulating a switch neuron both in bistable and single-state network, the baseline external Poissonian input (corresponding to 100%) to the stimulated pool (the first of the 8 available) was set to have spike rate of 2340 Hz and 2726 Hz, respectively. In working-memory modules this rate was reduced to 2 kHz in order to keep the low-firing state stable in all excitatory pools (see pre-stimulus activity between −2.5 and −2.0 s in Figure S4). Finally, for each cortical module, burst and pause detection analysis was performed on 100 randomly sampled neurons of the stimulated pool in 5 replicas of bistable and single-state cortical modules. Replicas of the same module differed only from a random selection of the synaptic efficacies, keeping unchanged mean and standard deviations of the Js. Responses to preferred and nonpreferred stimuli for the associated selective pool (E_1 in Figure S4) was simulated by delivering relatively large and small increases in the frequency of external spike rates, respectively. A preferred stimulus, elicited a large response of the selective pool (Figure S4A) amplified by the strong glutamatergic synaptic reverberation, while a nonpreferred stimulus (Figure S4B) did not lead to the threshold activity capable to elicit the nonlinear amplification of the pool. Different repetitions of this stimulation protocol represented different trials, and the spiking activity shown in Figure 5 was from a single neuron chosen between those composing the selective pool. Before the beginning of each simulated trial, the network was prepared in a specific activity state (initial condition). For instance, in the working-memory module this initial condition led all the 8 excitatory pools to fire at the same low average firing rate. In the bistable module including the switching neurons, the selective pool was alternatively prepared to stay in the asymptotic firing rate elicited by the input level set at the beginning of the trial (low and high firing rate states, light and dark gray conditions, respectively). Different stimulations are shown in Figure 5, aiming at modeling different trial stages/conditions. In Figure 5A (right) the input change to the selective pool modeled the S2 phase, while in Figures 5B and 5C, the input change modeled the "go" signal. In the multi-modular setting sketched in Figure 7, the synaptic connections between modules represent effective excitatory (red) and inhibitory (blue) couplings between local pools selectively engaged in the task. In other words, a link between an upstream (presynaptic) module and a downstream (postsynaptic) one actually can represent a potentiated or a depressed synaptic connection between selective pools.

The susceptibility plotted in Figure 5A was measured directly from the simulations as the ratio $\Delta FR/\Delta FR_{ext}$, where ΔFR was the firing rate increase in an excitatory pool due to an increment $\Delta FR_{ext} = 0.05\ FR_{ext}$ of the external neuron spike frequency. We administer the same stimulation protocol to each of the 16 cortical modules devised with different excitatory synaptic strengths. The protocol consisted of 20 trials with the same stimulation time course as the one shown in Figure 6A (right). External firing rate increase occurred after 500 ms from the beginning of the activity sampling and persisted for one second. Activity fluctuations were computed

for each module as the product between the coefficient of variation c_v (standard deviation over mean) of the firing rate across and within trials sampled in the last 500 ms of stimulation. More specifically, $c_v = c_v^{(intra)}\, c_v^{(inter)}$ where $c_v^{(intra)}$ was the the average across trials of the firing rate c_v, while $c_v^{(inter)}$ was the cv of the dataset composed of the mean firing rates in the different trials. In this way c_v was maximal (minimal) when fluctuations occurred (were dampened) both within and across trials. Tested networks were the same as the aforementioned working-memory module but with different strengths in the recurrent excitatory synapses $J_{intra} = w\ 0.361$, where w was the relative synaptic efficacy ranging from 125% to 155%. J_{inter} was set accordingly in order to keep unchanged the activity level of the low-firing state. More specifically, if $\Phi(FR, J_{intra}, J_{inter})$ is the input-output gain function of the excitatory neurons in the network (Amit and Brunel, 1997) $J_{inter}(J_{intra})$ is the implicit function satisfying the self-consistency equation $\Phi(FR_{Down}, J_{intra}, J_{inter}) = FR_{Down}$, where the Down state firing rate was FR_{Down} = 3 Hz. As a final step needed to avoid spontaneous escapes from the Down state due to the endogenous activity fluctuations, the frequency FR_{ext} was also decreased from 100% to 80% within the spanned range of w.

All simulations results were obtained using custom-made MATLAB (Mathworks) and C code (Mattia and Del Giudice, 2000).

Statistical analysis

Our sample size was not pre-determined but it is similar to the ones generally used in the field. Consistent with our previous work (Genovesio et al., 2011, 2014a), a one-way ANOVA was performed to identify the neurons with goal or action selectivity. An ANOVA is not sensitive to slight deviation from normality (Glass et al., 1972). For the rest of analyses, we used non-parametric tests because normality was not satisfied. All the analyses were performed using MATLAB (MathWorks).