

# Neuronal population dynamics during motor plan cancellation in non-human primates

Pierpaolo Pani<sup>1</sup>, Margherita Giamundo<sup>1</sup>, Franco Giarrocco<sup>1</sup>, Valentina Mione<sup>1</sup>, Emiliano Brunamonti<sup>1</sup>, Maurizio Mattia<sup>2\*</sup>, Stefano Ferraina<sup>1§\*</sup>

<sup>1</sup> Department of Physiology and Pharmacology, Sapienza University, 00185 Rome, Italy

<sup>2</sup> Istituto Superiore di Sanità, 00169 Rome, Italy

§ Corresponding author

\* These authors contributed equally.

Address correspondence to:

Prof. Stefano Ferraina, MD PhD  
Department of Physiology and Pharmacology  
P.le Aldo Moro 5 – CU027  
00185 Rome, Italy  
+390644910936  
Email: stefano.ferraina@uniroma1.it

# Abstract

To understand the cortical neuronal dynamics behind movement generation and control most studies focused on tasks where actions were planned and then executed, using different instances of visuomotor transformations. However, to fully understand the dynamics related to movement control one must also study how movements are actively inhibited. Inhibition, indeed, represents the first level of control both when different alternatives are available and only one solution could be adopted and when is necessary to maintain the current position.

We recorded neuronal activity from a multielectrode array in the dorsal premotor (PMd) cortex of monkeys performing a countermanding reaching task that requires, in a subset of trials, to cancel a potentially planned movement before its onset. In the analysis of the neuronal state-spaces of PMd we found a subspace in which activities conveying temporal information were confined during active inhibition and position holding. Movement execution required activities to escape from the plane toward an orthogonal subspace and, furthermore, surpass a threshold associated to the maturation of the motor plan.

These results revealed further details in the neuronal dynamics underlying movement control extending the hypothesis that neuronal computation confined in an output-null subspace does not produce movements.

*Keywords:* motor control, inhibition, neuronal dynamics, premotor cortex, monkey

# Introduction

Our brain controls movements. This is true both when movement are executed but also when they are just planned, in parallel (Cisek and Kalaska 2005), and never overtly expressed. Investigating the neuronal processes behind the voluntary control of movement in the brain is helpful to understand how goal oriented behaviours are either processed in healthy subjects or altered in patients where, for example, involuntary movements are often present (Chamberlain et al., 2007; Aron, 2011; Dalley and Robins, 2017).

One fruitful approach to the study of this topic is based on the use of the countermanding task. In this task a Go signal instructs to execute a movement while a Stop signal, randomly presented in some trials during the reaction time (RT), instructs not to move. In the past two decades, the countermanding task resulted fundamental for understanding how cortical and subcortical brain regions contribute to the control of saccadic eye movements (Schall et al., 2017). By describing how neuronal activity is modulated by the appearance of a Stop signal, it was possible to conclude that saccades are generated when activity of movement neurons in both Frontal Eye Field (FEF) and Superior Colliculus (SC) reaches a fixed threshold, while the activity of fixation neurons decreases (Hanes et al., 1998; Paré and Hanes, 2003). However, the control of arm movement tested in equivalent countermanding tasks does not elicit a similar neuronal correlate, especially when tested at the single cell level (Mirabella et al., 2011; Stuphorn, 2015; Pani et al., 2013), and different cortical areas like motor and premotor cortices (especially dorsal premotor cortex, PMd) have been suggested to have a key role in movement inhibition (Mirabella et al., 2011; Mattia et al., 2012; Pani et al. 2014; Pani et al. 2018; Duque et al., 2012; Duque et al., 2017; Parmigiani and Cattaneo, 2018).

Alternatively, to uncover the neuronal correlates of arm movement generation, most of the studies focused on different version of the delayed reaching task. Here, a cue signal providing information about movement parameters is given ahead of the Go, signalling to move. The delay epoch allows to easily separate signal-related activities from activities informing on either motor plan completion or movement execution (Weinrich and Wise 1982; Churchland et al. 2010). Using the delayed reaching task, recent approaches based on multi-electrode recordings, and neuronal state-space analysis suggested PMd as a coordinated system in which, during motor preparation, a cascade of neuronal events bring the collective activity into a preferred and stereotyped (attractor-like) state that sets the initial condition for the subsequent dynamics underlying the action production (Churchland and Shenoy, 2007; Mattia et al., 2013; Churchland and Cunningham, 2014). Within this framework, movement generation is the consequence of a change in the neuronal

state, i.e., in the collective pattern of activity across the population of neurons that will move from one region of the state space to another (Kaufman et al., 2016). Recent studies have further shown that such preparatory state is well separated by the inner representation of the successive movement execution (Kaufman et al., 2014; Elsayed et al., 2016), it is resistant to perturbations (Kaufman et al., 2014; Stavisky et al., 2017) and it can unfold with flexible durations, from few hundreds of milliseconds to seconds depending on the task settings (including reaction time tasks; Ames et al., 2014; Lara et al., 2018). However, all these reaching tasks always require the generation of a movement.

But what does it happen when a planned movement is cancelled? How is the state space visited by the neuronal ensemble dynamics of motor cortices in this instance? These questions are fundamental to reveal whether the movement generated is at least a “syntactic consequence” of a specific neuronal dynamics. Indeed, only comparing conditions in which a movement is performed *versus* those in which the movement is cancelled can reveal whether the unfolding of the specific neuronal dynamics is a required prerequisite for the movement to be made. Put it simply: if a neuronal dynamics is the substrate of movement generation, this must occur before the movement is executed while it must not occur when the movement is withheld.

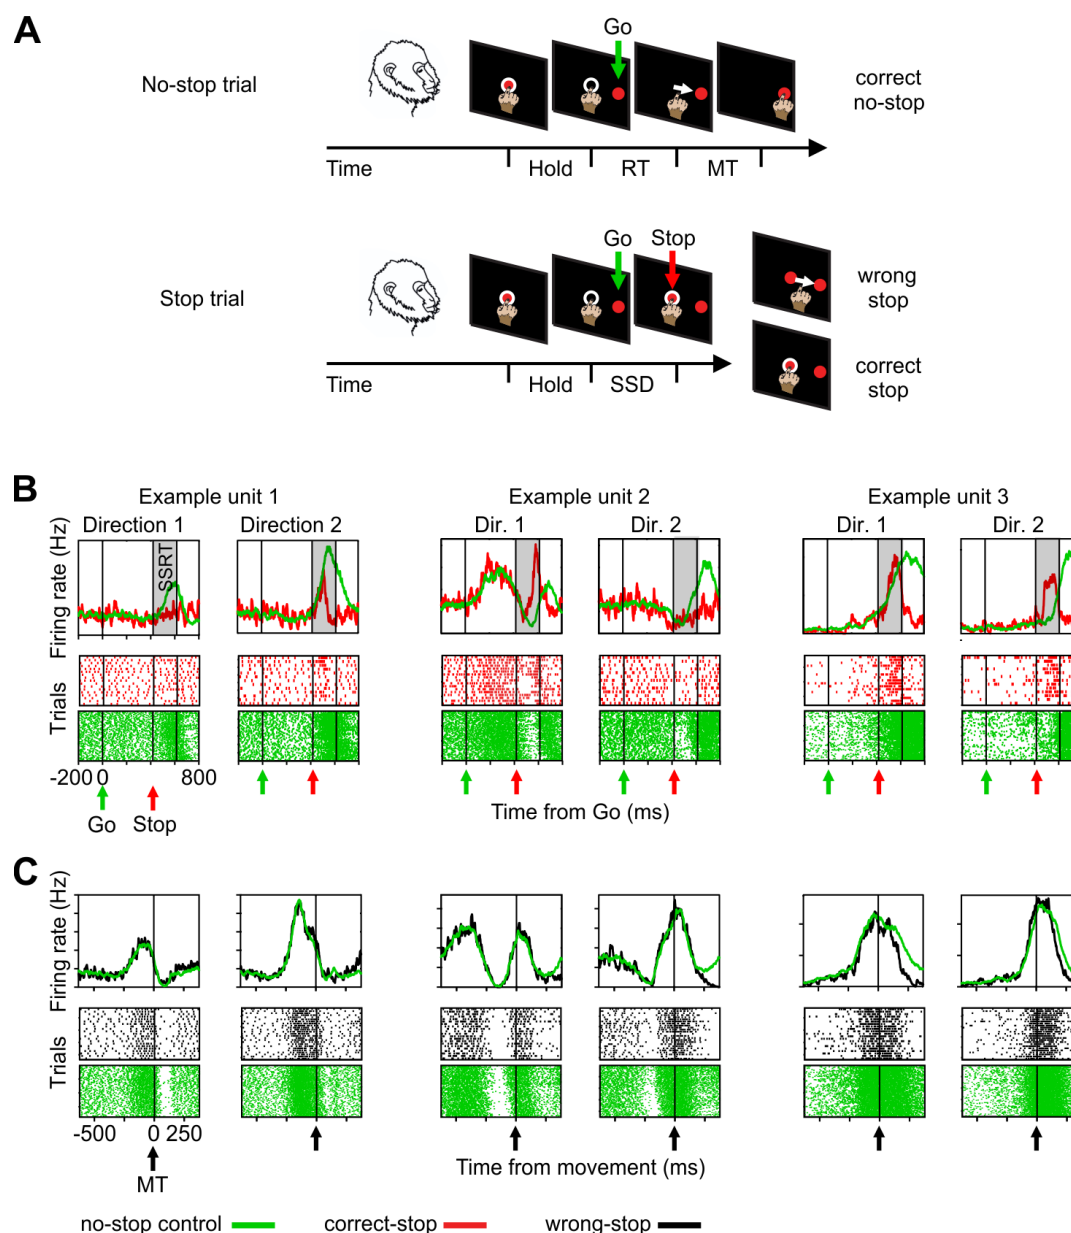
Here, we tackle these issues by uncovering the neuronal correlate of movement inhibition at the population level in the PMd of behaving monkeys. To this purpose we integrated two research lines in the study of the movement control by combining the testing capabilities of the countermanding task with the state-space analysis.

As a result, we propose and test the generalization of the hypothesis that motor-related neuronal activity must be confined within an output null-state region, in order to successfully suppress the translation of motor plans into overt movements, and that only the escape from this region can lead to the movement.

## Results

Monkeys (P and C) after completing their training period performed a countermanding reaching task with 66% of **no-stop trials** and 34% of **stop trials** (Fig.1A) and targets located either to the right or to the left in the workspace. For neuronal analysis we selected two (one for each animal) of those experimental sessions in which the animal behavior was compatible with the race model hypotheses (independence assumption; Band et al., 2003; see Figure S1A, B; Table S1, and Material and Methods) such that the **stop-signal reaction time (SSRT)**; see Table S1) could be

computed. This because the SSRT is an estimate of the time necessary to suppress the movement in the task, a needed information in order to further identify the dynamics of PMd units participating to movement inhibition.



**Figure 1. Behavioral task and heterogeneity of single neuron modulation during executed and cancelled movements.** **A.** No-stop trials and Stop trials were randomly intermingled on each block. Monkeys were required to touch the central target, hold the position (Hold) and wait until the Go signal (Go) instructed them to reach the peripheral target (red circle: either one out of two possible directions; here only one position is shown). In stop trials, after the Go signal, the Stop signal (central target reappearance) instructed them to stay still (correct stop). Wrong stop: trials with failed inhibition. **B.** Examples of units recorded: spike density functions (SDFs) and raster plots for correct stop trials and latency-matched (control) no-stop trials are shown for a specific stop signal delay (SSD) separately for each movement direction. **C.** For the same units SDFs and raster plots from wrong-stop trials and latency-matched no-stop trials aligned to movement.

## **Premotor units, associated to both movement execution and inhibition, are heterogeneous**

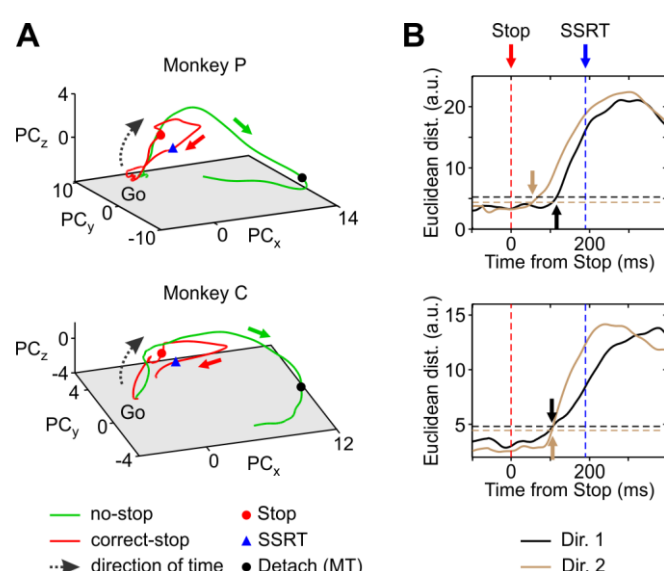
Initially we selected those units that were modulated before movement onset and/or after the Stop signal presentation in at least one movement direction compared to a control period (see Material and Methods). We then considered for further analysis only those units with a stop-related modulation for at least one movement direction. For this purpose, we searched for units predicting movement inhibition (and generation) being differently modulated in the SSRT interval. To this aim we selected units firing differently in the latency-matched no-stop trials and correct-stop trials, for the same movement direction (see Supplementary Material and Methods). No-stop latency-matched and correct-stop trials are characterized by the same maturation level of movement planning before the stop-signal, and their comparison allow to evaluate what happens when the movement generation is suppressed, being the same all the other motoric aspects.

Overall 139 units were selected, i.e. about two third of the dataset (93 out of 113 for monkey P and 46 out of 91 for monkey C). The selected units are heterogeneous because showed a variety of activity profiles in the two directions of movement during the task (for no-stop trials) and in relation to movement inhibition (either increasing or decreasing their activity in correct-stop trials when compared to latency-matched trials; see Fig. 1B) as previously described in other experiments in PMd (Kaufman et al., 2010; Mirabella et al., 2011). Moreover, we found that about 44% (61/139) of the units presented a different profile across the two directions of movement (see Figure 1B), while the others showed a similar modulation profile in the two movement directions. However, among the cells showing the same pattern, 49/78 (62%) was directionally selective (Wilcoxon rank-sum test,  $P < 0.05$ ), i.e., showed a different firing rate before movement onset for the two directions of movement. Overall, directional neurons were 110 (61 + 49; 79.1%) over 139 neurons. As expected, neuronal modulation was no different in latency-matched no-stop trials and wrong-stop trials (Figure 1 C; Figure S2 and Supplementary Material and Methods). Overall these results confirm previous evidence, i.e., that single units in PMd encode in complex ways action generation and suppression (Kaufman et al. 2010; Omrani et al. 2017)

## **Population neuronal dynamics during movement execution and inhibition**

The heterogeneity of the neuronal patterns observed is a confounding element in determining which aspect of the neuronal dynamics underlies movement inhibition and generation. To solve this issue, we analyzed data at the population level and reduced the dimensionality of the state space resorting to a principal components analysis (see Material and Methods for details).

Figure 2A shows, aligned to the Go signal, the average neuronal trajectories both in correct-stop (red) and in latency-matched no-stop (control trials; green) trials in the low-dimensional subspace determined by the first three principal components (PCs; explaining more than 60% of the variance in all conditions; see also Supplementary Material and Methods and Fig. S3) for one direction separately for each monkey. As expected, the two neuronal trajectories develop similarly until the appearance of the Stop signal (red dot). In control no-stop trials, population activity continues describing a volley until the movement occurs (black dot; detach). Conversely, after the Stop signal, the trajectory in correct-stop trials diverges from the green one moving backward to approach the initial state. Importantly, the divergence between the two trajectories occurs well before the SSRT (blue triangle).



**Figure 2. Neuronal dynamics and population estimate of the time when movements are inhibited.** **A.** Neuronal trajectories of no-stop and correct-stop trials for a single movement direction (direction 1) in the state space defined by the first three principal components. **B.** For each monkey, the average Euclidean distance between trajectories corresponding to correct-stop and latency-matched no-stop trials, aligned to the Stop signal, are represented. Arrows point out the time of significant divergence between trajectories. Dotted lines represent the reference thresholds for each movement directions (see Materials and Methods).

From these data we obtained a population estimate of the time of divergence between correct-stop trials and control no-stop trials activities (Figure 2B). To this end, we calculated the Euclidean distances between the correct-stop and the latency-matched no-stop neuronal trajectories within a time window centered around the Go signal (from 100 ms before to 200 ms after; i.e. up to the SSRT duration; see Table S1; see Materials and Methods for further details). We found that such divergence occurred between 70 and 116 ms following the Stop signal presentation for the two



directions {Dir.1, Dir.2}: Monkey C: {104, 107} ms; Monkey P: {70, 116} ms. By subtracting these values from the corresponding SSRT value (for each monkey and direction), we found that within the state space approach the neuronal modulation signaling movement suppression at the population level occurred before the one estimated from the analysis of the single unit activities separately (see Supplementary Materials and Methods; Figure S4). Population values are negatives, as preceding the movement, and compared to 95% Confidence Interval (CI) from the single-unit analysis and expressed in ms: Monkey P = {population: -126, -80} vs 95% CI= {-69.4, -51.4}; Monkey C = {population: -84, -67} vs 95% CI= {-76.7, -51.9}).

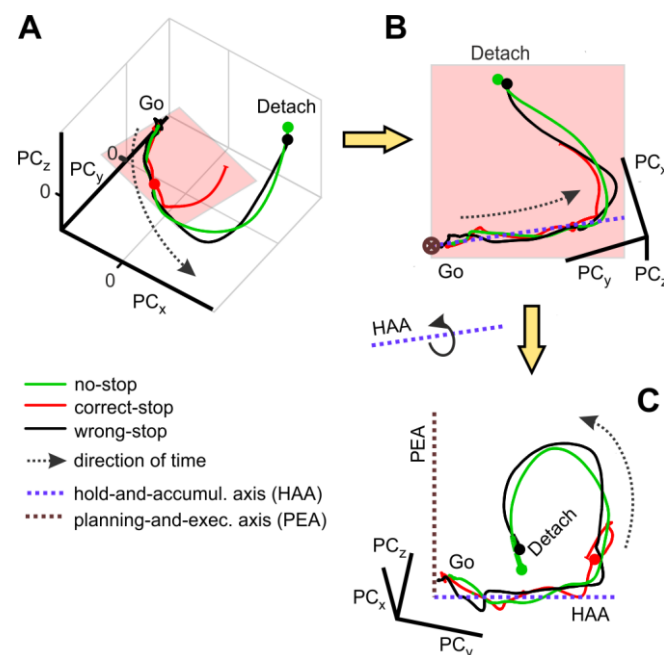
Importantly, when the population of the stop-related premotor units starts to encode the movement suppression, muscles were not activated yet (or, occasionally, just slightly activated, and then rapidly suppressed; see see Supplementary Results and Materials and Methods; Figure S5). The similarity between the population estimate of the time of divergence and the onset of muscle activity before the movement onset in no-stop trials (mean  $\pm$  SD = -82.5 ( $\pm$ 71.9) ms, CI 95% = {-103.5, -61}) indicates that suppression of the movement occurred by refraining from activating the muscles involved in reaching.

### **Movement inhibition and stillness require neuronal activities to be limited into a functional subspace**

We wanted to investigate whether the dynamics we observed for no stop and correct-stop trials could be better accounted for by the existence of distinct functional subspaces. Figure 3A shows, for one animal, the average neuronal trajectories in correct-stop (red), control no-stop (green), and wrong-stop (black) trials obtained from the first three PCs.

After the Go signal, the population activity in correct stop-trials initially follows a trajectory almost indistinguishable from the ones measured in no-stop and wrong-stop trials. The late divergence of the correct-stop trajectory (as in Figure 2) could indicate the existence of a subspace in the neuronal state space where activities are confined before movement generation and during movement suppression. Such hypothesis would support the generalization of the computational strategy relying on the ‘output-null’ subspace described for PMd neurons in animals tested in a delayed reaching task (Kaufman et al., 2014; Elsayed et al., 2016).





**Figure 3. Schematic of the neuronal trajectories unfolding within the holding plane and a plane orthogonal to it**  
**A.** Average neuronal trajectories for all trial types aligned to Go onset. **B.** Top. Holding plane, where correct-stop trajectories mostly, reside highlighted by a rotation of the trajectories in A. The plum dotted line represents the axis which best fitted the trajectories of all group of trials projected of the holding plane in the first 300ms following the Go-signal (hold-and-accumulation axis, HAA). **C.** A further rotation around the blue axis allows finding the planning and execution axis (PEA, brown dotted line).

To verify the existence of such subspace of activity confinement, we focused on the collective states visited by the correct stop trials. We applied on these data, averaged and grouped by SSD for each animal and direction of (potential) movement, a Singular Value Decomposition (SVD) analysis (see Material and Methods for details). We found all trajectories of the correct-stop trials to fluctuate around a reference planar section. Figure 3B shows the view on this planar section of the average trajectories in Figure 3A here obtained, for illustrative purposes, after a suited rotation of the original 3D space. In this plane the similarity among no-stop, wrong-stop, and correct-stop trajectories is remarkable. Thus, the dimensionally-reduced population dynamics highlighted from this perspective does not contain any information about whether or not the movement will be generated.

However, as premotor cortices are known to determine the motor output in reaching tasks (Churchland et al. 2010), an orthogonal subspace must exist to allow independently and robustly reading out these motor commands from the activity evolution of the network (Elsayed et al. 2016), and thus allowing to distinguish between trials where movement is executed or inhibited. To this end, we found an axis by further fitting data from correct-stop trajectories in the 300 ms following the Go-signal. This axis was then used to extract a plane orthogonal to the first one. Figure 3C

shows (again for illustrative purposes only) the orthogonal plane, together with the corresponding dynamics of all average neuronal trajectories (bottom). From this last perspective, it is rather apparent that movement related activities (no-stop and wrong-stop trajectories) depart from the reference holding plane where correct-stop activities are instead confined for most of the time.

We called the first axis the *holding-and-accumulation axis (HAA)*, and the second subspace the *planning-and-execution axis (PEA)*, including ‘accumulation’ and ‘planning’ in the definitions for reasons that will be fully clear later.

### **Movement inhibition as an activity trap to escape from for movement generation**

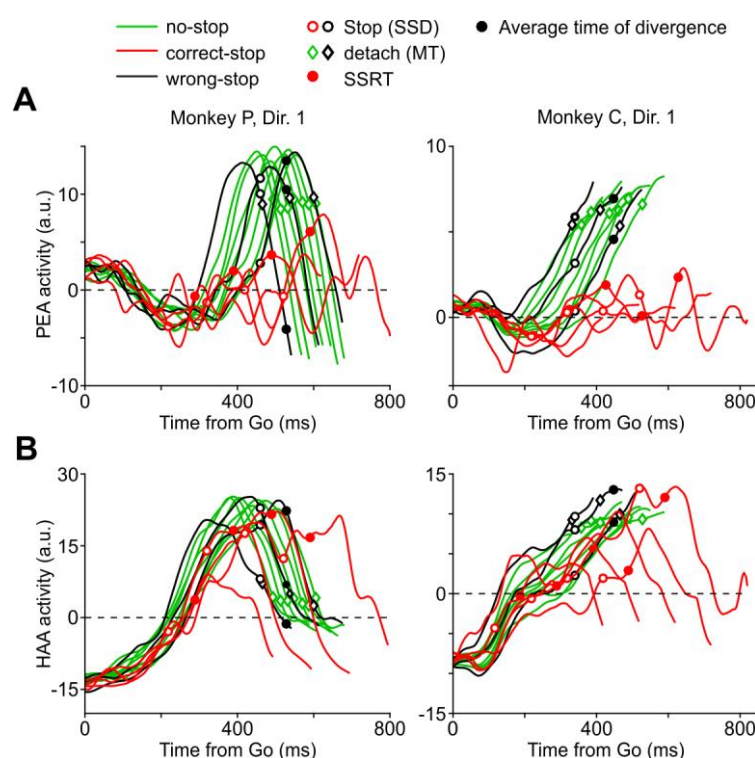
Thus a question arises: which is the mechanism underlying the departure of the movement-related activities from the reference subspace where correct stop trajectories lay? What does make a stop trial correct or wrong? To address this issue, we inspect in more detail how the projections of the neuronal trajectories onto the described axes changes in time.

Figure 4A shows (grouping trials by either RT or SSD for wrong/no-stop trials and correct-stop trials, respectively) that neuronal activities, observed as projections in the PEA, initially (after Go signal) are roughly stable, slightly fluctuating around the reference zero-baseline, and then they move towards more negative values, ‘drawing’ a trough (see below for further details on this dynamics). During the same period, neuronal activities observed as projections in the HAA (Figure 4B) show a ramp-like dynamics, similar for all trials and conditions (see also Figure S6).

In a phase successive to the troughs, Figure 4A shows that activities, in particular those related to wrong-stop (black lines) and no-stop trials (green lines), escape from the plane signalling an important change in the overall state. These transitions are described by a sudden and steep modulation of activity predictive of the movement onset time (see below). In agreement, we previously suggested that a rapid and stereotyped activity change is tightly linked to the maturation of the motor plan eventually determining the motor execution (Mattia et al., 2013).

Projections on both HAA and PEA show a strong correlation with RTs (see Fig. S7). The relationship was slightly stronger when considering projections on PEA in  $\frac{3}{4}$  of the conditions [median coefficient of correlation (IQR): monkey P, direction 1: PEA = 0.99 (0.03); HAA = 0.91 (0.38),  $p < 10^{-4}$ ; direction 2: PEA = 0.98 (0.07); HAA = 0.89 (0.43),  $p < 10^{-4}$ ; monkey C, direction 1: PEA = 0.99 (0.02); HAA = 0.81 (0.40),  $p < 10^{-4}$ ; direction 2: PEA = 0.90 (0.07); HAA = 0.95 (0.62);  $p = 0.1$ ]. However, projections onto the HAA, showed a more variable relationship with RTs, as suggested by the higher IQR.

We further asked in which of the two projections the relationship with RTs first emerged. We found that in HAA the significant relationship between neuronal activities and RTs emerged earlier than in PEA [mean (SD); monkey P, direction 1: HAA = 377 (246) ms, PEA = 482 (133) ms; direction 2: HAA = 300 (278) ms, PEA = 450 (123) ms; monkey C, direction 1: HAA = 276 (104) ms, PEA = 332 (101) ms; direction 2: HAA = 299 (182) ms, PEA = 365 (72) ms; rank-sum test  $p < 10^{-4}$  in all cases]. Thus, although in both projections a significant relationship with RTs emerged, this relationship develops as a chain first involving the HAA projections and only afterwards involving the PEA projections.

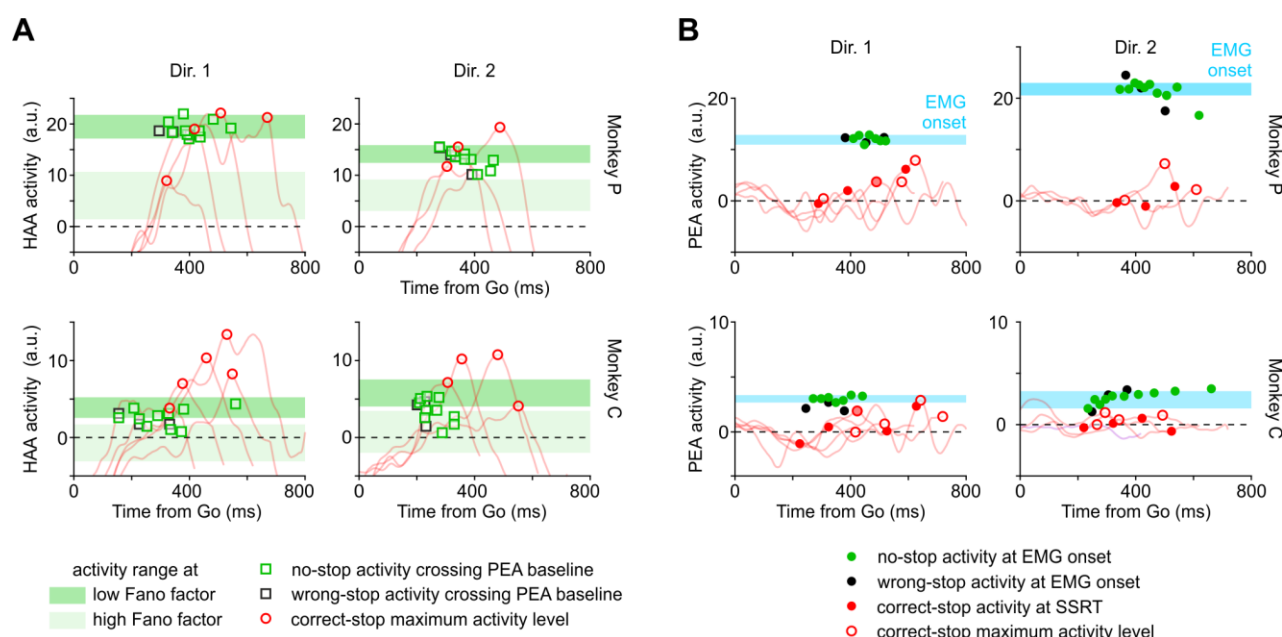


**Figure 4. Components of the neuronal dynamics determining RTs and correct stop trials.** A. Projections on the Planning and Execution Axis (PEA) of the neuronal trajectories (direction1) – namely, PEA activity – grouped by RTs (no-stop in deciles and wrong-stop trials in tertiles) or SSDs (correct-stop trials). After an initial period around the Go signal in which trajectories are overlapped, no-stop and wrong-stop trials projections diverge from correct stop trials projections. B. Projections on the Holding and Accumulation Axis (HAA) of the same neuronal trajectories – namely, HAA activity – as above: in this case the projections display an initial ramp-like dynamics and similar features for all trials and conditions.

### Neuronal activity projections onto PEA highlight a threshold-based movement generation

The ramp-like accumulation dynamics observed after the Go signal in the HAA (Figure 4B) is typically observed in decision processes or when an internal generated timing signal is employed to elaborate an input signal, suggesting the presence of a threshold (Maimon and Assad 2006; Chandrasekaran et al. 2017).

This is further suggested, aligning HAA activity to movement onset (see Figure S8) and by the reduction in neuronal variability (measured as inter-deciles Fano factor  $F = \text{variance}/\text{mean}$ ) preceding the movement (see Figure S9) often considered coincident to the termination of a decision-related accumulation process (Churchland et al. 2007, 2010). Since we were interested to see what happens around this time, we considered (before muscle activity) an interval (40 ms) when the neural variability was low versus an interval of the same duration when the neural variability was still high, similarly separated in time in the two animals (see Figure S9).



**Figure 5. Dynamics for movement generation exploits a threshold-based mechanism in the PEA only.** **A.** To construct each panel, we extracted the times when projections on PEA of no-stop and wrong-stop trials crossed the zero baseline (green and black squares, respectively). The corresponding HAA activity is the neuronal activity projected onto the HAA at the extracted times. The shadow bars represent the levels of activity (mean  $\pm$  SD; 40 ms intervals) observed in HAA and corresponding to high (light green) or low (dark green) values of the Fano factor of the same activity before movement onset (see Fig. S9). **B.** Cyan bars represent range of PEA activity (i.e. neuronal activity projected onto the PEA) when muscles are activated before movement generation (mean  $\pm$  SD, calculated on the muscle's latency CI). Maximum PEA activity in correct-stop trials are empty red circles, while the filled ones correspond to the PEA activity of the correct-stop trials at SSRT. Green (black) filled circles display the PEA activity during no-stop (wrong-stop) trials at the time of EMG onset.

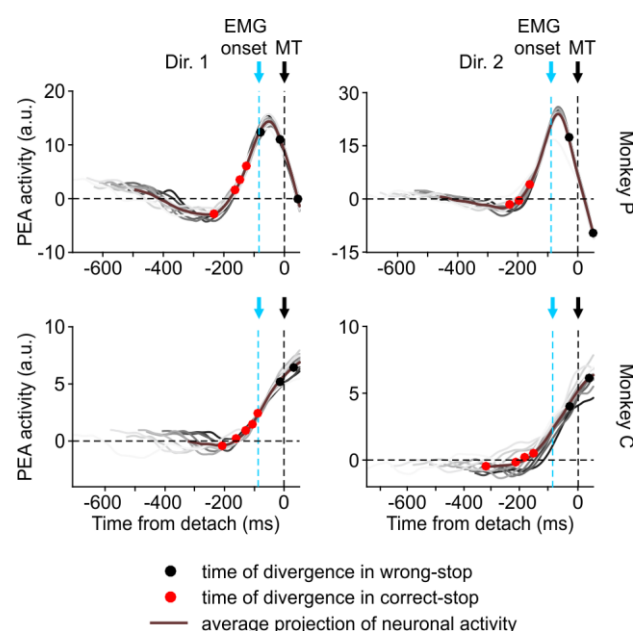
Figure 5A shows the level of activity for each RT decile in the HAA subspace of no-stop (green squares) and wrong-stop (black squares) trials when the corresponding activities projection onto the PEA start to depart from zero. The shaded green ranges indicate the level of activity in no-stop trials corresponding to the two periods of high (light green) and low (dark green) Fano factor in Figure S9. The escape of activity in PEA usually corresponds to a low variability of the HAA projections. However, this escape mechanism resulted not threshold-based when observed from the HAA

perspective. Indeed, activity in correct-stop trials (red empty circles in Figure 5A; maximum activity for each SSD; groups of trials as in Figure 4) can even pass the level corresponding to the low variability without eliciting any overt action.

Conversely, Figure 5B shows for correct-stop trials the maximum PEA activity for each SSD, together with the PEA activity at the moment of muscle activation (cyan band). Thus, when movements are inhibited, the escape of activity in the PEA remains under the level necessary for muscle activation, suggesting a threshold-based mechanism in this subspace for the movement initiation.

As can be expected from Figure 4 for no-stop and wrong-stop trials, PEA activities are also highly stereotyped in time when aligned to the movement onset (Fig. 6), similarly to what found in (Mattia et al., 2013). It is important to remark here that this dynamics occurs while specification of movement direction is already determined (i.e., well before movement onset; see Figure S10).

Contrary to what happens in wrong-stop trials, the PEA activity during correct-stop trials do not deviate too much from the baseline. This is compatible with the hypothesis that correct-stop trajectories are trapped into an intermediate subspace as in these trials movement planning is not yet mature at all or it is cancelled well before its complete maturation, compatibly with what preliminary observed in (Mattia et al., 2012; Mattia et al., 2013).

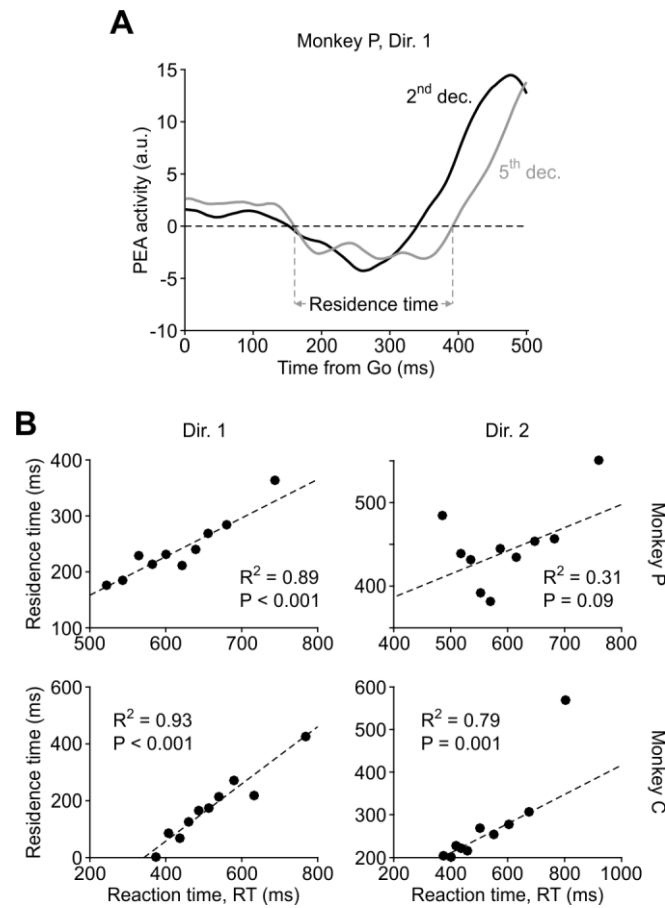


**Figure 6. Movement inhibition is effective only if Stop signal is processed before full muscle contraction.** Each panel represents the PEA activity for the no-stop trials grouped by deciles and aligned to the movement onset. Filled circles indicate the time of divergence for correct-stop trials (red; sorted by SSD) and wrong-stop trials (black; sorted by RT duration; tertiles).

However, according to (Kaufman et al., 2014), the maturation of the motor plan still belongs to the output-null subspace. Indeed, at this processing stage yet there is no movement of the arm, but the network state is properly positioned to be the potential initial condition for the forthcoming motor-related trajectories (Churchland et al. 2010). In this framework, a successful processing of the Stop signal could be obtained by inhibiting the translation of the motor plan into an overt movement, as in delayed reaching tasks. The departure from the holding subspace is characterized by the presence of a point-of-no-return that must be surpassed to generate a movement (Figure 6). To appreciate this, the key time is represented by the muscle activity onset (Figure 6, cyan dashed lines). In fact, in all the wrong-stop trials the expected neuronal divergences (black circles) due to the Stop-signal reaction, follow the EMG onset suggesting that the Stop signal can be effectively processed (time of divergence in correct-stop trials; red circles) only if occurring before muscle contraction. This is clearly reminiscent of a threshold-like mechanism underlying the escape from the output-null subspace. Indeed, the PEA activity is higher when a movement is performed (Figure 6, black circles before movement onset), compared to the highest peak that the fluctuating activity can reach during correct-stop trials (Figure 6 and 5B, red circles).

Figure 7A (see also Figures 4 and 6) shows that before the stereotyped escape from the holding subspace, the activities stay in a trough for a variable time interval. These durations are tightly related to the RT, as a significant correlation resulted when comparing activity residence times in the trough with the associated RT durations (Figure 7B).

Overall, the emerging picture is that after target onset the neuronal dynamics visits different states better differentiated when observed as distances from different subspaces. Movement generation is kept at bay until neuronal activity is trapped into the trough below the reference subspace (negative values observed: mean  $\pm$  SD, Monkey C: Dir.1 =  $-0.5 \pm 0.3$ ; Dir.2 =  $-0.59 \pm 0.36$ ; Monkey P: Dir.1 =  $-2.0 \pm 1.0$ ; Dir.2 =  $-1.6 \pm 1.1$ , t-test  $p < 0.01$  for all cases). The movement is generated when this activity grows reaching a threshold (see Figure 5B). Importantly this growth of activity in the PEA starts only when the accumulation process reaches a preceding threshold level in HAA, purportedly corresponding to the end of a decision process.



**Figure 7. Residence time estimate and relationship with reaction time.** **A.** For each decile of no-stop trials we estimated a residence time in a trough below the baseline of the PEA projections. **B.** Correlation between residence time duration and RTs for each movement direction, separately for each monkey.



## Discussion

In the present study we investigated the premotor neuronal dynamics underlying the control of arm movements by combining two complementary lines of research: i) neuronal population analysis, by exploiting a state-space approach (see for example Ames et al., 2014; Churchland et al., 2012; Kaufman et al., 2014), and ii) the countermanding task, an experimental design ideally suited to investigate and model volitional control of movement (Schall et al., 2017; Schmidt and Berke, 2017).

First, we confirmed that PMd units modulate their activity before the end of the SSRT thus having temporal characteristics that allow predicting whether or not a movement will be halted (Mirabella et al., 2011). Secondly, we further extended this finding by showing that neuronal modulation precedes muscle modulation in successfully inhibited trials, strongly supporting a causal role for this area in movement generation. Thirdly, and most importantly, by inspecting the premotor network activity as trajectories in an effective low-dimensional space, we uncovered the existence of a new holding subspace, which we suggest generalizing the previously introduced output-null subspace (Kaufman et al., 2014). In order to have a successful stop of an instructed movement, the state of the premotor cortex has to be trapped into this subspace where the full maturation of the motor plan is inhibited. Such novel control strategy extends the concept of the output-null subspace and it is complementary to the one adopted in delayed reaching tasks when motor plans fully develop before a Go signal instructs the movement execution.

### *Movement control in a state-space framework*

The main goal of this study was to uncover the neuronal substrate underpinning the movement control (inhibition vs generation) to reveal whether exists a specific neuronal dynamics that characterizes the movement generation with respect to movement inhibition.

We found that following the Go signal the neuronal activity is confined into a (holding) subspace where neuronal activity is initially similar between correct-stop, wrong-stop and no-stop trials. In this initial stage the neuronal activity is “trapped” into a trough visible from a specific perspective, possibly reflecting the specific task context that requires to block movement execution. As the neuronal dynamics evolves a clear differentiation emerges: in no-stop and wrong-stop trials the neuronal activity moves away from the trough towards a state-space region which almost deterministically anticipates the movement generation. Differently, in correct-stop trials the neuronal activity remains confined in the initial subspace, separate from the trajectory observed in no-stop trials. In this last condition no movement is observed. Importantly the temporal evolution in

no-stop trials (and wrong-stop trials) is strongly related to RTs, and it shows a clear stereotyped nature when aligned to the movement onset. These evidences demonstrate the existence of neuronal dynamics that anticipates the motor behavior clearly distinguishing the active inhibition versus the movement generation. Thus, we conclude that PMd expresses specific neuronal dynamics underpinning either movement generation or inhibition. This conclusion is rooted on the fact that the divergence between no-stop and correct-stop trials occurs before the end of SSRT (Hanes et al. 1998).

Our data are in line with a series of findings obtained from studies investigating the preparation and execution of reaching movements by using state-space approaches (Kaufman et al 2014; Lara et al. 2016, Lara et al. 2018, Michaels et al. 2015). In most of these experiments monkeys are provided with prior information about movement direction by a cue followed by a Go (move) signal after a delay. During the preparation phase (delay period) a cascade of neural event brings the collective activity into a prepare-and-hold (attractor) state (Churchland et al., 2010; Mattia et al., 2013; Ames et al., 2014). As proposed, the goal of this state is to set the motor plan, i.e., a neuronal state which functions as the initial condition determining the upcoming movement (Churchland et al. 2010; Michaels et al. 2015; Michaels et al. 2018). The neuronal trajectories will then evolve towards a complementary subspace, but linked, to the previous one (Elsayed et al., 2016). In this last subspace, reaching execution dynamics (rotational) unfolds (Churchland et al., 2012). Within this framework the preparatory activity does not generate *per se* the movement but it always occurs before the movement onset, even when very brief times to prepare the movement (zero delay) are available (Lara et al. 2018). Thus, it is a state that does not allow unwanted movements (Kaufman et al., 2014; Stavisky et al., 2017). The transition from preparation to execution is characterized by some degree of overlapping between preparation and movement phases (Lara et al. 2018), and by a strong condition-invariant signal that precedes the movement onset (Kaufman et al. 2016).

In summary, an important aspect of the previous studies is that the generation of the movement was always required, thus it is not clear which aspect of the dynamics can be deemed necessary for movement generation. We employed the countermanding task to address this issue. To this purpose we exploited the concept of neural manifolds (Gallego et al., 2017) finding a holding subspace where neuronal trajectories of the correct-stop trials, as well in the initial part of no-stop and wrong-stop trials, are confined. The escape of neuronal trajectories from this subspace determines the transition from the prepare-and-hold to the movement generation. This dynamics is highly stereotyped and precedes the movement onset by a fixed time-lag, similar to what observed

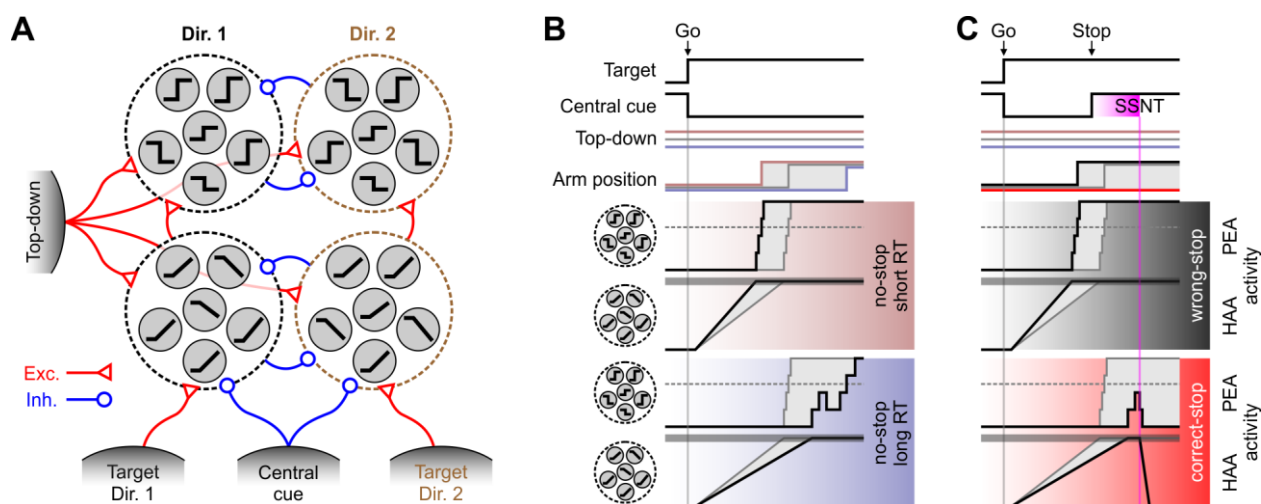
in other studies (Kaufman et al. 2016). This dynamics could correspond to the information to start the movement, an internal Go signal, that PMd can provide to other cortical and subcortical structures, as previously suggested (Kaufman et al., 2014; Elsayed et al., 2016).

This allows to avoid searching for specific inhibitory circuits that, so far, have been impossible to define in the premotor and motor cortices of primates, due to the difficulties in accounting for the connections of cortical neurons to subcortical, brainstem and spinal cord neurons (Kraskov et al. 2009), although recent attempts have been promising (Soteropoulos 2018).

### ***Nonlinear network dynamics underpinning the holding subspace***

HAA and PEA are orthogonal axis, meaning that activity projection on such axes is mainly due to different subset of units. In other words, ramping activity along the HAA is probably due to a subset of units displaying a post-Go quasi-linear increase/ decrease of the firing rate (as example units 2 and 3 in Fig. 1B, sketched at the bottom of Fig. 8A). Instead, the sudden increase of PEA activity tightly locked to the movement onset is determined by a pool of units displaying sharp transitions from low-to-high activity levels and *vice versa* (see example units 1 and 3 in Fig. 1B, Fig. 8A-top). These two subsets only in part overlap, and the heterogeneity of their roles and behaviours appears to be a universal feature in cortical areas involved in arm movement planning (Seidemann et al., 1996; Maimon & Assad, 2006; Mattia et al., 2013, Chandrasekaran et al. 2017). In the specific framework investigated here, the relatively slow and linear drifting along the HAA determines a post-Go “refractory phase” where presumably target-related information accumulates and no movement occurs. The passage across a region (see Fig. 5A, low Fano factor) after a variable time elicits the rise of the PEA activity. In those trials with short RTs (Fig. 8B-middle), the related sharp transitions of that activity are strongly stereotyped even in the Go-centred profiles (Fig. 6), meaning that units displaying a fast switching of activity are rather synchronized in time (Mattia et al., 2013).

In correct-stop trials and in no-stop trials with long RTs more gradual ramps and less stereotyped activations of the switching units are visible (Fig. 8A and B-bottom). Intriguingly the latter appears to fluctuate up and down (Fig. 4A), and under this condition, motor program maturation can be then interrupted if it is not fully developed, i.e. the activity is confined below a trigger threshold (Fig. 8B-bottom). Indeed, only when planning has reached a sufficient degree of maturation it can be translated into an overt movement.



**Figure 8. Schematic proposal of PMd organization underlying movement execution and inhibition.** **A.** Pools of ramp-like and switch-like units (grey circles) respond differently depending on the target to be reached signaled by a selective excitatory input coming from other brain areas (grey shaded units at the bottom). Target related representations are mutually exclusive due to cross inhibition. Ramp-like units contributing to the rising of HAA activity elicit sharp-transitions in switch-like units. The latter contributes to the changes of PEA activity that indirectly react to the target-related input. An unspecific top-down input (left) making differently excites all the PMd units from trial-to-trial. Onset and offset of the central cue signaling to Stop and Start, respectively has an inhibitory effect on the PMd units. Red and blue arrows represent excitatory and inhibitory connections, respectively. **B.** Behavioral, environmental (top) and neuronal changes in no-stop trials with long (bottom) and short (middle) RTs, which are determined by a relatively weak (bluish) and strong (reddish) top-down excitation, respectively. Grey traces represent a reference trial with average RT. Black traces, pooled activity of switch-like and ramp-like units are plotted in black side by side with the respective unit icons. These traces are those resulting from projecting the high-dimensional population activity on the PEA and HAA, respectively. Horizontal dashed grey lines, activity level when movement onset is irreversibly elicited when crossed in the PEA. Simultaneously, when in the HAA the pooled ramp-like units reach the horizontal gray strip, the switch units are facilitated to change state. **C.** As in panel B, for those trials in which the Stop signal is presented. Black and red shaded subpanels correspond to example wrong- and correct-stop trials, respectively.

RTs are both determined by the time needed to reach an accumulation threshold (dashed line in Fig. 8B) and by a variable time in which switching units complete the transformation of the cortical state into the one encoding the motor plan. The latter phase has a more prominent role only for relatively long RTs, and the Stop signal can be successfully processed both if SSRT fall after the end of the first accumulation stage. In possible short-RT trials the onset of Stop signal must be early in correct-stop trials such that the SSRT could precede the beginning of the activity rise along the PEA. If one of these two conditions is not fulfilled, a wrong-stop trial is performed (Fig. 8C-middle).

A mechanistic implementation of the discussed interplay between ramping and switching units on one hand requires that the former provide input to the latter, such that once the input received by the switching units crosses a trigger threshold, a cascade of activity switches occurs (Mattia et al., 2013; Marcos et al., 2019). On the other hand, one should also explain the trial-by-trial variability of such chain of reactions, which underlies both the capability to inhibit an

instructed movement and the variability of the RTs. To this aim, an additional unspecific input has to be taken into account. Indeed, a tonic input modulating the excitability of both unit types would allow changing from trial to trial the switching rate and the ramp slope (Cao et al., 2016). This input could be provided by a proactive tonic top-down control, which varies across trials (top-down signal in Fig. 8).

The existence of such a proactive control mechanism in models of saccadic countermanding tasks is not new (Lo et al., 2009; Wiecki & Frank, 2013). For instance, in (Lo et al., 2009) a top-down control of frontal origin implemented a post-Go phasic reduction of the excitatory input to a specific subset of neurons. This sudden change in the input resembled the effects of a homunculus intervening at specific random times aiming at recovering the RTs statistics and inhibition function (Schall et al., 2017). To work around such limitation, in (Wiecki & Frank, 2013) is suggested that another proactive mechanism is at work leading to a modulation of the “response caution” governed by the baseline activity of the infero-frontal gyrus (IFG). This in turn, slows down or accelerates the movement production via the IFG-subthalamic nucleus hyperdirect pathway. Our results are compatible with this second scenario, as each trial has its own degree of excitability/stability of the ramping and switching units that we suggest to be associated to a different top-down tonic input (i.e., the IFG baseline activity). Differential stability of the no-movement condition can be interpreted as a different speed of the Go process in the race model (Logan & Cowan, 1984; Verbruggen & Logan, 2009), and we predict it is related to the high-beta state found in (Leventhal et al., 2012) associated to the capability of the cortical-basal ganglia circuits to stabilize selected motor plans.

Under this hypothesis, the permanence in the newly found holding subspace appears to be effectively separated from the boundaries delimiting the previously introduced output-null space (Kaufman et al., 2014). This in our view is obtained by making the holding subspace a more “attracting” region. In the framework we propose, the attracting force is modulated by an unspecific top-down control, while the premotor network can autonomously implement the machinery needed to produce and inhibit an instructed movement. In turn, the stability of this holding state allows having complex and computationally relevant population dynamics such as the accumulation process preceding the full maturation of the motor plan, without the danger to cross the boundary of the output-null subspace triggering the movement onset.

*Neuronal correlates of movement control include also subcortical structures*

Premotor cortex is not the only structure to be involved in movement inhibition. In this framework, an important role is also played by the basal ganglia. More specifically, it has been argued that in rodents the suppression of instantiation movement in the basal ganglia occurs through different mechanisms. Mallet and collaborators (2016) suggested that, when a movement is halted, a first pause general mechanism acting through sub-thalamic nucleus (STN) and substantia nigra pars reticulata (SNR) will make more difficult for the movement generation to unfold (possibly increasing the level of neuronal activity that must be reached), while a subsequent, selective mechanism, mediated by the activation of archypallidal cells, will abolish movement preparation occurring in the striatum (Mallet et al., 2016). Effects of these modulations during movement inhibition will be also visible at the motor cortical level. Importantly, at the cortical level these effects will be heterogeneous, thus suggesting that cortical premotor activity would display these effects in a neuronal state space, and not through simple push-pull patterns as observed at the basal-ganglia level (Oldenburg and Sabatini, 2015). Although there is no direct evidence in primates, this is a possible mechanism. The presence of modulation in basal ganglia during the stop task, leads to a question: which are the relative contributions of premotor cortical and basal ganglia structures in controlling the movement? To solve this issue simultaneous recording from cortical and subcortical structures seems necessary.

# Material and Methods

## Subjects

Two adult male rhesus macaque monkeys (*Macaca mulatta*; P, C) weighing 7–9.5 kg served as subjects. All experimental procedures, animal care, housing, and surgical procedures conformed with European (Directive 86/609/ECC and 2010/63/UE) and Italian (D.L. 116/92 and D.L. 26/2014) laws on the use of nonhuman primates in scientific research and were approved by the Italian Ministry of Health. In both monkeys a 96 channels Utah arrays (BlackRock Microsystem, USA) was implanted using anatomical landmarks (arcuate sulcus - AS - and pre-central dimple - pCD) after dura aperture on the dorsal premotor cortex (PMd) contralateral to the arm employed during the experiments (Figure S1C)

## Apparatus

Stimuli were presented on a 17-inch LCD monitor (800×600 resolution) equipped with a touch-screen (MicroTouch, USA). Stimuli consisted of red circles with a diameter of 2.8 cm on a dark background. During the performance of the task eye movements were monitored by using a non-invasive Eye-tracker (Arrington Research Inc, AZ).

## Behavioral task

Monkeys performed a countermanding task (Fig 1A). They had to touch the stimulus with their fingers, hold and fixate it (Holding Time range: 400-900 ms). Thereafter the central stimulus disappeared and, simultaneously, a target appeared (Go signal) randomly at one of two opposite peripheral positions (left and right). In no-stop trials to get a juice reward, monkeys had to start the arm movement within a maximum time (1000 ms) to discourage from adopting a procrastination strategy due to the presence of stop trials, and to maintain their fingers on it for a random time (400-800 ms, 100ms step). In stop trials at a variable delay (stop signal delay, SSD) after the go signal was presented, the central stimulus reappeared (stop signal) instructing the monkey to keep the hand on the starting position (additional holding time; 400-1000 ms interval, 100ms step) to perform a correct-stop trial and earn the juice. If the monkey moved the hand during stop trials, the trial was considered a wrong-stop trial and no reward was given. No-stop and stop trials were randomly intermingled in such a way that the no-stop trials were more frequent (66%).

Data were collected using a staircase tracking procedure to change the SSD from one stop trial to the next according to the behavioral performance: if the monkey succeeded in withholding the response (correct-stop trial), the SSD increased by a fixed amount (100 ms); if it failed (wrong-stop trial), the SSD decreased



by the same amount of time. The goal of the tracking procedure is to determine a SSD for which the probability of response (i.e. the probability to have a wrong-stop trial) is 0.5 or very close to this value. This allows estimating a reliable stop signal reaction time (SSRT, Osman et al., 1986, 1990; Band et al., 2003; Verbruggen and Logan, 2009, see below for further details).

## **Behavioral analysis**

The countermanding task permits to estimate the stop signal reaction time (SSRT) by extracting three main variables: the reaction time (RT) distributions of no-stop trials and wrong-stop trials, and the probability to respond [p(R)] by error to the stop signal. These data are modelled according to the race model (Logan and Cowan, 1984) to establish first if the assumptions of the model are respected (see below) and, in this instance, to estimate the SSRT. Briefly the model states that in stop trials two stochastic processes race toward a threshold: the Go process, started by the Go signal, and the Stop process, started by the Stop signal. The behavioural result of this race, either movement generation in wrong-stop trials or movement suppression in correct-stop trials, will depend on which of these process first will reach his own threshold. In correct-stop trials the Stop process wins over the Go process and vice versa. By changing the SSD the output of the race is affected: the longer the SSD, the higher the probability to facilitate the GO process.

An important assumption of the model is that the Go process in the stop trials is the same as in the no-stop trials (independence assumption, Logan and Cowan, 1984; Band et al., 2003; Verbruggen and Logan, 2009; Verbruggen et al., 2013; see Supplemental Material and Methods).

We calculated the SSRT by using the integration method because it is the most reliable (Verbruggen et al., 2013). This method requires subtracting a given SSD from the finishing time of the STOP process. The finishing time of the STOP process is calculated by integrating the no-stop trial RT distribution from the onset of the Go signal until the integral equals the corresponding observed proportion of wrong-stop trials for the given SSD (Logan, 1994). We obtained the estimate first by calculating the mean SSD presented in each session and then subtracting it from the finishing time obtained by considering the overall p(R) (Ridderinkhof et al., 1999; Verbruggen et al., 2004; Verbruggen et al., 2013).

## **Neural recordings**

The unit activity was first isolated online, and then controlled off-line by using specific software (Open Sorter, Tucker Davies Technologies). In this work we included both single unit as well multi-unit activity. Multi-unit was defined as spiking activity that could not definitively attributed to a single unit. In the text we use the term unit without distinction. This is because the main analyses we performed are based on population activity and dimensionality reduction techniques typically produce very similar results regardless of the employment of single or multi-unit (Kaufman et al. 2016; Trautmann EM et al. 2019) Array data in

this paper come from a single recording session for both monkey P and monkey C. The use of data coming from a single, not repeated recordings, allowed us to completely avoid risks of unit replication across different sessions.

## Neuronal analysis

To perform some of the neuronal analysis, we represented the neuronal activity by a spike density function (SDF) obtained by convolving the spike train with an exponential function mimicking a postsynaptic potential: we used the same equation described in Scangos and Stuphorn (2010) where the convolution kernel  $K(t)$  of the neuronal activity is:

$$K(t) = [1 - \exp(-t/\tau_g)] \cdot \exp(-t/\tau_d)$$

where  $\tau_g = 1$  ms corresponds to the growth phase of the synaptic potential, and  $\tau_d = 20$  ms is the decay phase (Sayer et al., 1990).

We selected for analysis units that showed an increase/decrease of the average firing rate before movement onset (pre-RT, from -200 to -50ms before movement onset), compared to the 200 ms before Go signal (pre-Go) for at least one movement direction (Wilcoxon rank-sum test  $P < 0.01$ ), and/or on units that showed a significant difference between correct-stop trials and latency-matched (see below) no-stop trials after the Stop signal. For this last case average neuronal activities in 50 ms not overlapping bins were compared in the (+100, +400) ms interval after Stop signal presentation. A neuron was included for further analysis if it showed a significant difference for at least one bin (Wilcoxon rank-sum test  $P < 0.01$ ). These analyses were performed by using the exact spike counts. Directional selectivity was evaluated by comparing the pre-RT activity between movement directions (Wilcoxon rank-sum test  $P < 0.01$ ).

For correct-stop trials the latency-matched no-stop trials were those with RTs longer than SSD+SSRT for each specific SSD and session; for wrong-stop trials latency-matched no-stop trials were those with RTs shorter than SSD+SSRT.

## Neuronal population analysis in a low-dimensional state space

To study the neuronal dynamics at the population level, we resorted to a standard principal component analysis (PCA) to reduce the dimensionality of the recorded state space. The approach we developed consists in two main steps. Firstly, we characterized the neuronal dynamics of the movement suppression process and the latency of its neuronal correlate at the population level. Then, we evaluate the planning and movement generation process across different trials types and RTs.

In the first step we calculated average spike density functions in 1-ms time bins by aligning the trials

to the Go and to the Stop signals, separately. For no-stop trials where the Stop signal was absent, we considered the hypothetical SSD corresponding to the average SSD obtained for correct-stop trials. As results, for each trial type and movement direction the spike densities from all the different units composed a matrix of samples. We then concatenated these three matrices preserving the number of rows ( $N$  = number of units). The activity of each unit was then normalized by subtracting mean activity and dividing the activity standard deviation computed across all conditions. The obtained neuronal activities (matrix rows) were eventually smoothed with a Gaussian kernel of 50-ms width. In this framework, the population state of the probed cortex at any time was represented by a  $N$ -dimensional vector (a column of the aforementioned activity matrix).

At this stage, we computed the distance between average no-stop and correct-stop trials trajectories to estimate the time of divergence at population level. More specifically, in the above  $N$ -dimensional state space we carried out the Euclidean distance between no-stop and correct-stop trials at each 1-ms time bin in the interval (-100, 200) ms centred around the Go signal. This was done for each movement direction. We then calculated the mean and the SD of these distances and we set a threshold level corresponding to mean + 3 SD. Finally, the Euclidean distance was also computed between the average population activity in no-stop and correct-stop trials aligned to the Stop signal. The latency of the Stop process onset was the first time from the Stop signal onset when the population activity distance exceeded the threshold level.

In order to capture the inter-trial variability associated to different behaviors and at the same time to reduce the unavoidable fluctuations of the estimated spike densities, we further divided the no-stop trials into ten groups (deciles) and the less numerous wrong-stop trials into three groups (tertiles) based on the ordered RTs (neuronal activity spanned from 50ms before the Go signal up to 50 ms after movement onset) (see Figs. 4-6). The correct-stop trials were grouped based on the length of the SSDs and the neuronal activity was shown from 50 ms before Go signal. The group number of correct-stop trials varied from 3 to 5 depending on the condition. Three-dimensional trajectories were then obtained from grouped trials by averaging in each group the population activity eventually smoothing it with a Gaussian kernel (50 ms).

Starting from these grouped spike densities, in the second stage of the analysis we performed the PCA on the concatenated matrix of such average population activities with dimensions  $N$  (neurons)  $\times$  Ctg (conditions  $\times$  time bins per trial group  $\times$  number of groups). As illustrated in the main text (Figs. 4-6), we further singled out some relevant subspaces. Firstly, we looked for a plane where correct-stop trajectories for a given condition resided. This was defined as the holding plane (Fig. 3 B-top) best fitting (Singular Value Decomposition; SVD) the cloud of points composing all the three-dimensional trajectories associated to the average population activity for each group of correct-stop trials. We then looked for the axis which best fitted the trajectories of all group of trials projected of the planning-null plane in the time window between 0 and 300 ms from the Go signal. This holding-and- accumulation axis (HAA) is represented by a dashed plum dotted line in Fig. 3B, and we found to be the ideal subspace to represent the accumulation of information before the beginning of the related motor plan maturation (always occurring after 200 ms from the Go signal

for all conditions and animals), as shown Fig. 4A. Finally, we defined the planning-and-execution axis (PEA; brown dotted line in Fig 3B, bottom) as the one orthogonal to the planning-null plane. The trajectory projections onto this direction optimally highlighted the maturation process of the motor plan and the neuronal correlate of the movement execution. The time course of such projections for each trial group is shown in Fig. 4-bottom. The three-dimensional average trajectories for the no-stop, correct-stop and wrong-stop trials in Fig. 3B-bottom are shown on a plane after a suited rotation around the HAA.

We also performed another analysis to establish in which of the two subspaces the relationship between neuronal projections and RTs started first. To this aim we considered two samples of data for each subspace, movement direction and monkey: one composed by the times of all the points crossing the thresholds (exclusively for significant regressions), the other composed by the times of the deciles from 4th to 6th. We then compared, for each direction and monkey, the corresponding distributions relative to the PEA to the HAA (Wilcoxon rank-sum test,  $P < 0.01$ ). In both cases the significant relationship started first in the HAA in all conditions considered. In the main text we report the results from the analysis performed on the deciles from the 4th to the 6th.

## **Electromyographic recordings and analysis**

Electromyographic (EMG) signal was obtained by needle electrodes (monopolar derivation) inserted into the targets muscles with the monkey calm. Muscles recorded were: pectoralis, deltoid (anterior, middle, posterior), infraspinatus, biceps, triceps (Fig. S3A).

Signal from each trial was sampled at 3052 Hz, then band-pass filtered (60-1000 Hz), rectified and down sampled to 1000 Hz. Finally, it was smoothed by using a moving window of 30 ms (in 1 ms steps) and averaged across trials. Data reported are from 6 recording sessions. For further details, see Supplemental Material and Methods.

**Author contributions:** P.P. and S.F. designed research; P.P, E.B., V.M., M.G, F.G. performed research; P.P., M.M. analyzed data; P.P., S.F., M.M., wrote the paper. All authors revised critically the work.

## **Acknowledgments**

Sapienza H2020 2017 grant to SF.

# References

- Ames KC, Ryu SI, Shenoy KV. 2014. Neuronal dynamics of reaching following incorrect or absent motor preparation. *Neuron* 81:438-451.
- Aron AR. 2011. From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol Psychiatry* 69:e55-68.
- Band GP, Van der Molen MW, Logan GD. 2003. Horse-race model simulations of the stop-signal procedure. *Acta Psychol* 112:105-142.
- Buneo CA, Soechting JF, Flanders M. 1994. Muscle activation patterns for reaching: the representation of distance and time. *J Neurophysiol* 71:1546-1558.
- Cao R, Pastukhov A, Mattia M, Braun J. 2016. Collective activity of many bistable assemblies reproduces characteristic dynamics of multistable perception. *J Neurosci* 36(26), 6957–6972.
- Chamberlain SR, Fineberg NA, Menzies LA, Blackwell AD, Bullmore ET, Robbins TW, Sahakian BJ. 2007. Impaired cognitive flexibility and motor inhibition in unaffected first-degree relatives of patients with obsessive-compulsive disorder. *Am J Psychiatry* 164:335-338.
- Chandrasekaran C, Peixoto D, Newsome WT, Shenoy KV. 2017. Laminar differences in decision-related neuronal activity in dorsal premotor cortex. *Nat Commun* 8:614.
- Churchland MM, Cunningham JP. 2014. A dynamical basis set for generating reaches. *Cold Spring Harb Symp Quant Biol* 79:67-80.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012. Neuronal population dynamics during reaching. *Nature* 487:51-56.
- Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. 2010. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* 68:387-400.
- Churchland MM, Shenoy KV. 2007. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol* 97:4235-4257.
- Cisek P, Kalaska JF. 2005. Neuronal correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45:801-814.
- d'Avella A, Portone A, Fernandez L, Lacquaniti F. 2006. Control of fast-reaching movements by muscle synergy combinations. *J Neurosci* 26:7791-7810.
- Dalley JW, Robbins TW. 2017. Fractionating impulsivity: neuropsychiatric implications. *Nat Rev Neurosci* 18:158-171.
- Duque J, Greenhouse I, Labruna L, Ivry RB. 2017. Physiological markers of motor inhibition during human behavior. *Trends Neurosci* 40:219-236.
- Duque J, Labruna L, Verset S, Olivier E, Ivry RB. 2012. Dissociating the role of prefrontal and premotor cortices in controlling inhibitory mechanisms during motor preparation. *J Neurosci* 32:806-816.
- Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. 2016. Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun* 7:13239.

- Gallego JA, Perich MG, Miller LE, Solla SA. 2017. Neural manifolds for the control of movement. *Neuron* 94:978–984.
- Hanes DP, Patterson WF, Schall JD. 1998. Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J Neurophysiol* 79:817–834.
- Hanes DP, Schall JD. 1996. Neuronal control of voluntary movement initiation. *Science* 274:427–430.
- Kaufman MT, Churchland MM, Shenoy KV. 2013. The roles of monkey M1 neuron classes in movement preparation and execution. *J Neurophysiol* 1, 10:817–825.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2014. Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* 17:440–448.
- Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, Shenoy KV. 2010. Roles of monkey premotor neuron classes in movement preparation and execution. *J Neurophysiol* 104:799–810.
- Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM. 2016. The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* 3:1–25.
- Vigneswaran G, Kraskov A, Lemon RN. 2011. Large identified pyramidal cells in macaque motor and premotor cortex exhibit “thin spikes”: implications for cell type classification, *J Neurosci* 31(40), 14235–14242
- Lara AH, Cunningham JP, Churchland MM. 2018. Different population dynamics in the supplementary motor area and motor cortex during reaching. *Nat Commun* 9:2754.
- Lara AH, Elsayed GF, Zimnik, AJ, Cunningham J, Churchland MM. 2018. Conservation of preparatory neuronal events in monkey motor cortex regardless of how movement is initiated. *Elife* 7,189035.
- Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD. 2012. Basal ganglia beta oscillations accompany cue utilization. *Neuron* 73(3), 523–536.
- Lo CC, Boucher L, Paré M, Schall JD, Wang XJ. 2009. Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. *J Neurosci* 29(28), 9059–9071.
- Logan GD. 1994. On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In: *Inhibitory processes in attention, memory, and language*. (Dagenbach D, ed; Carr TH, ed), San Diego, CA: Academic Press. 189–239.
- Logan GD, Cowan WB. 1984. On the ability to inhibit thought and action: A theory of an act of control. *Psychol Rev* 91:295–327.
- Maimon G, Assad JA. 2006. A cognitive signal for the proactive timing of action in macaque LIP. *Nat Neurosci* 9(7), 948–955.
- Mallet N, Schmidt R, Leventhal D, Chen F, Amer N, Boraud T, Berke JD. 2016. Arky pallidal cells send a stop signal to striatum. *Neuron* 89:308–316.
- Mattia M, Pani P, Mirabella G, Costa S, Del Giudice P, Ferraina S. 2013. Heterogeneous attractor cell assemblies for motor planning in premotor cortex. *J Neurosci* 33:11155–11168.
- Mattia M, Spadacenta S, Pavone L, Quarato P, Esposito V, Sparano A, Sebastiano F, Di Gennaro G, Morace R, Cantore G, Mirabella G. 2012. Stop-event-related potentials from intracranial electrodes reveal a key role



of premotor and motor cortices in stopping ongoing movements. *Front Neuroeng* 5:12.

Michaels JA, Dann B, Intveld RW, Scherberger H. 2015. Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J Neurosci* 35(32):11415-11432.

Michaels JA, Scherberger H. 2018. Population coding of grasp and laterality-related information in the macaque fronto-parietal network. *Sci Rep* 8, 1710.

Mirabella G, Pani P, Ferraina S. 2011. Neuronal correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys. *J Neurophysiol* 106:1454-1466.

Oldenburg IA, Sabatini BL. 2015. Antagonistic but not symmetric regulation of primary motor cortex by basal ganglia direct and indirect pathways. *Neuron* 86:1174-1181.

Osman A, Kornblum S, Meyer DE. 1986. The point of no return in choice reaction time: controlled and ballistic stages of response preparation. *J Exp Psychol Hum Percept Perform* 12:243-258.

Osman A, Kornblum S, Meyer DE. 1990. Does motor programming necessitate response execution? *J Exp Psychol Hum Percept Perform* 16:183-198.

Pani P, Di Bello F, Brunamonti E, D'Andrea V, Papazachariadis O, Ferraina S. 2014. Alpha- and beta-band oscillations subserve different processes in reactive control of limb movements. *Front Behav Neurosci* 8:383.

Pani P, Giarrocco F, Giamundo M, Montanari R, Brunamonti E, Ferraina S. 2018. Visual salience of the stop signal affects the neuronal dynamics of controlled inhibition. *Sci Rep* 8(1):14265.

Pani P, Menghini D, Napolitano C, Calcagni M, Armando M, Sergeant JA, Vicari S. 2013. Proactive and reactive control of movement are differently affected in attention deficit hyperactivity disorder children. *Res Dev Disabil* 34:3104-3111.

Paré M, Hanes DP. 2003. Controlled movement processing: superior colliculus activity associated with countermanded saccades. *J Neurosci* 23:6480-6489.

Parmigiani S, Cattaneo L. 2018. Stimulation of the Dorsal Premotor Cortex, But Not of the Supplementary Motor Area Proper, Impairs the Stop Function in a STOP Signal Task. *Neuroscience* 394:14-22.

Ridderinkhof KR, Band GP, Logan GD. 1999. A study of adaptive behavior: effects of age and irrelevant information on the ability to inhibit one's actions. *Acta Psychol* 101:315-337.

Robinson DA. 1963. A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Bio-Med Eng* 10:137-145.

Schall JD, Palmeri TJ, Logan GD. 2017. Models of inhibitory control. *Philos Trans R Soc Lond B Biol Sci* 372(1718).

Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. 1996. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16(2), 752-768.

Soteropoulos DS. 2018. Corticospinal gating during action preparation and movement in the primate motor cortex. *J Neurophysiol* 119(4):1538-1555.

Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2017. Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neuronal state space dimensions. *Neuron* 95:195-208.

- Stuphorn V. 2015. Neuronal mechanisms of response inhibition. *Curr Opin Behav Sci* 1:64-71.
- Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O'Shea DJ, Vyas S, Sun X, Ryu SI, Ganguli S, Shenoy KV. 2019. Accurate Estimation of Neural Population Dynamics without Spike Sorting. *Neuron* 103:292-308.
- Verbruggen F, Logan GD. 2009. Models of response inhibition in the stop-signal and stop-change paradigms. *Neurosci Biobehav Rev* 33:647-661.
- Weinrich M, Wise SP. 1982. The premotor cortex of the monkey. *J Neurosci* 2(9):1329-45.
- Wiecki TV, Frank MJ. 2013. A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychol Rev* 120(2), 329–355.