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

Pierpaolo Pani[1], Margherita Giamundo[1], Franco Giarrocco[1], Valentina Mione[1], Roberto Fontana[1], Emiliano Brunamonti[1], Maurizio Mattia[2*], Stefano Ferraina[1§*]

[1] Department of Physiology and Pharmacology, Sapienza University, 00185 Rome, Italy

[2] Istituto Superiore di Sanità, 00169 Rome, Italy

§ Corresponding author

* Equally contributing senior authors.

Supplementary Results

Supplementary Figures and Tables: 17

Supplementary Material and Methods

Supplementary References

Supplementary Results

Test of the independence assumption in behavioral data

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; 1-5; see Supplemental Material and Methods). For the example sessions, we obtained that considering all the values, wrong-stop predicted RTs were shorter than wrong-stop observed RTs for both monkeys (544 ± 32 vs 574 ± 63 rank-sum test: $p < 0.0001$ monkey P; and 417 ± 35 vs 447 ± 58 $p < 0.002$, monkey P and monkey C respectively). No differences were found comparing the mean values of RT for each SSD in Monkey P ($p = 0.1$) while a difference was found for Monkey C ($p = 0.03$). Furthermore, in both cases mean RT in wrong-stop increased with increasing SSDs (Spearman corr. = 1 for monkey P and Spearman corr. = 0.8 for monkey C). Overall, the results support the independence assumption (see Figure S1 and Table S1 for the other sessions). Note, however, that the most important criteria to observe is to avoid the estimation of SSRT when RTs on wrong-stop are numerically longer than RTs of no-stop trials (5).

Time of significant contribution to movement inhibition at the single neuron level in PMd

We estimated when single neurons signal a difference between executed and inhibited movements. To this aim we focused on conditions with negative time lags between the neuronal modulations reflecting movement inhibition (differences in activity between correct-stop trials and latency-matched no-stop trials), and the corresponding SSRT (see Supplementary Material and Methods; 148/278 [2 movement directions × 139 units]), i.e., conditions for which the neuronal modulation was possibly causally related to inhibitory behavior. After excluding statistical differences between monkeys [$F(1, 146) = 0.24$; $p = 0.63$; monkey P: -60.4 ± 4.7 ms, 95% CI = {-69.7, -51.4}; monkey C: -64.3 ± 6.5 ms, 95% CI = {-77.7, -52}], we found that the population average of the time lags across all monkeys resulted -61.0 ± 3.8 ms, 95% CI={-69.5, -54.1} ms (Figure S5).

We estimated these time lags also within the state-space approach. Here we found that the neuronal modulation signaling movement suppression could be detected slightly before the one estimated from the analysis of the single unit activities: Monkey P = -103 ms (95% CI = {-126, -80}) vs -60.4 ms (95% CI = {-69.7, -51.4); Monkey C = -76 ms (95% CI = {-84, -67}) vs -64.3 ms (95% CI = {-77.7, -52}).

Overall, these results show that a subpopulation of task modulated units in PMd is capable of signaling the inhibitory reaction needed to stop a movement starting already around 80/100 ms before the end of SSRT. Importantly, when the population starts to encode the movement suppression, muscles are not activated yet (or, occasionally, just slightly activated, and then rapidly suppressed; see Supplementary Results - below - and Figure S6).

Weak neuronal modulations are observed after Stop signal in wrong-stop trials

Once we found that PMd units selectively change their activity in the proper time (i.e., before the end of the SSRT) to inhibit a movement, we also asked whether the Stop signal was capable to affect the neuronal activity in wrong-stop trials, i.e., when the movement was generated irrespective of the command to halt. A visual inspection of neuronal activity in wrong-stop and latency-matched no-stop trials shows that no difference is detected at the moment of movement generation (see unit examples in Figure S2).

To perform a quantitative analysis we considered all the 148 conditions indicated above, and we compared Contrast Indexes obtained from correct-stop vs latency matched no-stop trials to the corresponding Contrast Indexes obtained from wrong-stop vs latency matched no-stop trials. (Figure S3, Correct vs Wrong, see Supplementary Materials and Methods for details). Positive Contrast Indexes indicate a higher activity in no-stop compared to stop-trials, and vice versa.

If in wrong-stop trials the Stop signal can affect neuronal activity, we should observe non-null Contrast Indexes and their values should be close to those observed in comparisons regarding correct-stop trials. Values either close to zero or different from the Contrast indexes in correct-stop trials represent low or null effects of the Stop signal in wrong-stop trials.

We ran a mixed ANOVA with factors stop trial (correct- or wrong-) and sign of the Contrast Index (either Positive, i.e., higher activity in no-stop trials, or Negative, corresponding to lower activity in no-stop trials). We found that both Positive and Negative indexes in correct-stop trials were significantly different from the ones in wrong-stop trials (Negative Correct: -0.50 ± 0.03, 95% CI = {-0.57, -0.43}; Negative Wrong: -0.11± 0.05, 95% CI = {-0.19, -0.02}, $p < 10^{-5}$; Positive Correct: 0.48 ± 0.01, 95% CI = {0.45, 0.51}; Positive Wrong: -0.050 ± 0.019, 95% CI = {-0.09, -0.012}, $p < 10^{-5}$; $F(1, 236)=239$, $p < 10^{-4}$. Newman-Keuls post-hoc test, MSE = 0.061, df = 470).

These data show that even if there is some modulation during wrong-stop trials, overall it is weaker compared to correct-stop trials. Indeed, despite an attempt to inhibit the movement, the neuronal activity in wrong-stop trials is not different from the neuronal activity in no-stop trials.

In the specific task, muscles are weakly activated in stop trials and premotor units signal movement inhibition before muscles

What is the role of the stop-related activity? Is it associated with muscle activations accompanying movement inhibition? Or alternatively, does it aim at interrupting the transformation of the motor plan into an overt action?

To address these questions, we recorded the electromyographic (EMG) activity of the arm muscles

participating to the reaching movements (see Figure S6A), as done in previous studies (6-8). These sessions were different from those of neuronal recordings. However, the behavior of the EMG sessions conformed to the race model: the wrong-stop trials were faster than in no-stop trials (average difference across sessions = 38 ± 12 ms, t-test (5) = -6.6, p = 0.0012); most importantly the SSRT estimates (mean ± SD = 192 ± 25 ms) were very close to the SSRT estimates obtained in the neuronal recording sessions from the same monkeys (see Table S1).

Most of the time, when movements were successfully inhibited in both directions, the EMG activity did not show any significant modulation (for an example muscle see Figure S6B). This shows that, at least in the context of this task and setup, stopping an arm movement may not require any muscle activation. More specifically, we found that in the large majority (82; 66%) of all conditions (124: direction × muscle × session, criterion of at least 5 trials per condition) there was not any significant muscle modulation in correct-stop trials in the 300 ms-window following the Stop signal. For the other 42 conditions, we compared the EMG activity of no-stop trials and correct-stop trials to extract the latency of the first significant difference. In the conditions for which this latency could be estimated [38 out of 42, 90%; in 4 conditions the difference did not reach the criteria (see Supplementary Material and Methods)], it always preceded the behavioral SSRT. In this subset, we obtained the estimate of the electromyographic time lags by subtracting the latency of the first significant EMG difference from the behavioral SSRT. This is the electromyographic counterpart of the neuronal time lag. Muscle modulation related to inhibition followed the neuronal modulation: the muscular time lag indeed was, in absolute value, smaller than the neuronal time lag: -38.2 ± 7.1 ms, 95% CI = {-52.3, -24.1}; ($F(1, 184)=8.6$, $p = 0.004$), see Figure S6 C. Thus, the EMG signature of the stop process was either an unmodified level of activity, or when a small increase of activity was observed this was suppressed always after inhibition related neuronal modulation, and before behavioral SSRT suggesting that a proper action cancellation requires a control before muscle activation, at least in the version of the task we used.

We confirmed that the muscles we recorded from, were involved in movement generation by comparing their patterns of activations between opposite directions: all the muscles were active before movement generation and showed directional dependence (Figure S6 D-G). However, although the muscles showed directionality, they did not show clear agonistic or antagonistic activations when compared between opposite movement directions. A possible reason for this is that the arms of the monkeys - having to act on stimuli presented on the touch screen - had to contrast the force of gravity throughout the phases of the trial, so the muscles had to be always active at some level. This experimental context is different from the ones of other studies where the use of a cursor or a joystick in a 2-D plane probably helped in generating agonistic-antagonistic patterns (for example, see 9).

We also wanted to investigate the differences in muscle activation between wrong-stop trials, and correct-stop trials, aligned to the Stop signal. We expected that in wrong-stop trials the muscle activation must occur before the end of the hypothetical SSRT to start the movement. We tested this hypothesis by

calculating the latency of the difference between wrong-stop and correct-stop trials (Figure S6F and G). From a sample of 36 conditions (at least 7 trials for correct- and wrong-stop trials, with the same SSD) we found that the latency of this difference was 74 ± 40 ms (mean ± SD) from the Stop signal (maximum detected latency was of 134 ms). Thus, well before the end of the behavioral SSRT (more than 100 ms) muscle activity in wrong-stop trials was already too high to be suppressed.

In wrong-stop trials the monkeys typically started the reaching movement and then halted it before they could reach the peripheral target. Based on this observation we asked whether and how the presentation of the Stop signal might affect muscle activity during the movement phase of the wrong-stop trials. We then compared the muscle activity between the wrong-stop trials and latency matched no-stop trials. Indeed, the first phase of the movement is the same between the two trials (see Figure S6H for an example muscle), while the differences emerge only once the Stop signal is encoded in wrong-stop trials and the reaching movement is then halted. To increase the power of our analysis, we selected no-stop trials in the range of the RTs ± interquartile range (IQRT) observed for wrong-stop trials (for wrong-stop trials we considered conditions with at least 6 wrong-stop trials per SSD). Considering all 88 conditions (SSD × direction × muscles) for which the comparison was possible, we found differences between the two types of trials in the 150 ms preceding RT (comparisons between average activities in windows of 40 ms and 1 ms step, of wrong-stop and no-stop latency matched trials) only in 2 conditions. We typically observed that a clear difference between wrong-stop and no-stop trials occurred after the movement started, purportedly related to the difference in movement between the two trial types. We decided to calculate the latency of this difference relative to the movement onset, and we found a significant difference for 69/88 conditions; on average this value was 71 ± 43 ms (Figure S6I). In summary, about 70 ms after the movement onset a clear change in EMG activity probably reflected the action of the online movement stopping process in wrong-stop trials.

On average we estimated that EMG activity started to increase above baseline about -82.5 (71.9) ms; 95% CI = {-104, -61} relative to movement onset (values obtained across 48 conditions used to estimate this value; threshold was set to average baseline + 3SD).

In agreement with Logan (10) an incoming movement can be stopped only if the stop process acts before the point of no return, i.e., the point when neuronal and muscle activities supporting movement generation are too advanced to be suppressed by it. This last case is observed in wrong-stop trials, in which the Stop signal is presented when the movement planning-execution is too advanced, and thus the movement is generated despite its presentation. By applying the race model, we can consider the end of SSRT as coincident to the movement onset (10). In this view, to be effective the Stop signal must be presented around 190 ms before the movement starts. In this study we found that muscle activity starts to rise about 85ms before movement onset and that the modulation of muscle activity following a Stop-signal can be detected in some cases around 40 ms before the end of SSRT. Thus, neuronal activity must either impede the rise of or suppress a low level of muscle activity before and not later than 85-40 ms before movement onset to yield an

effective inhibition

We found that the latency of neuronal population signal reflecting suppression can start already around 80 ms before the end of SSRT, suggesting that the information to stop arrives on average when the EMG activity is - possibly - only beginning to rise, as confirmed by our observations on EMG activity.

HPA activity reflects movement planning

To better understand the role of the HPA activity in movement generation we implemented a **Context informed Go-NoGo task**. This task was characterized by the following sequence of stimuli and events: at each trial a visual Cue defining the context (Cue-Context-signal, see below) was presented for 200-300 ms as soon as the monkey touched the central target. Then a Target appeared either to the left or to the right. After a variable delay (800-1200 ms) a Command Signal appeared, instructing the monkey either to reach towards the Target (Go signal), or to keep the hand in position (No-Go signal).

The Cue-Context signal allowed us to modulate the activity of movement planning. Indeed, 3 different Contexts were defined: the **Advanced Go** (see the Figure S12), informed the monkey that at the end of the trial the Command Signal was a Go signal; the **Advanced NoGo** informed the monkey that at the end of the trial the Command signal was a NoGo signal; and the **Go-NoGo** context that informed the monkey that in 50% of the trials a Go signal would appear as Command signal, while in the others a NoGo signal would.

Considering all the trials, 37% were Advanced Go, 21% Advanced NoGo, 21% Go trials, 21% NoGo trials. Thus, four classes of trials were collected: Advanced Go in which the monkey was sure since the beginning of the trial that a movement was required; Advanced NoGo in which the monkey knew since the beginning of the trial that no movement was required (and, as such, no planning was required); Go, in which the monkey decided to move after the Command signal; NoGo in which the monkey decided not to move after the NoGo signal. Advanced NoGo trials were conceptually similar to the NoGo trials of the Go/NoGo task previously used to study the advanced cancellation of motor action (11-13).

Using this task, we extracted data from 113 units. Trials were ordered by RT and grouped in deciles and tertiles for Advanced Go and Go classes. For Advanced NoGo and NoGo were ordered by delay durations and grouped in tertiles. We performed the same analysis done on the countermanding (CMT) task on this Go-NoGo task (see Material and Methods). In this case the holding space was estimated by considering the activity in the NoGo trials, because they required movement preparation but the Command Signal was a NoGo signal. Somehow, this condition can be considered as a Stop condition with SSD = 0. The monkey did not make errors in this condition.

Figure S12 shows the activity projection on PEA and HPA of the different trial classes. In the PEA (Figure S12A), only activities associated with movement generation (Advanced Go and Go; gray and black

traces respectively) emerged above the dotted line representing a threshold level. Interestingly this activity starts even before the Command Signal presentation for the Advanced Go trials. In the HPA all the classes but one (the Advanced NoGo; red traces), show an increase in activity that overlaps up to 200 ms following the Command Signal. These classes are associated with movement preparation and planning, because a Command signal requiring movement is certain or probable. Their activities are similar to the ones observed in the CMT for all the trial types (no-stop, wrong-stop, and correct-stop). This is because in all these instances the generation of a movement was possible. In this new task only the Advanced NoGo shows a constant low level of activity in the HPA. Indeed, the cue presentation makes clear to the monkey that no movement will be required. As such, we can hypothesize that the HPA activity is associated with the first stage of movement planning. As such it is necessary for movement generation (because it is always present before movement generation or in cases when a movement is possible), but its presence does not define whether there will be a movement. The generation of the movement occurs only when activity increases in the PEA. This is further supported by the linear relationship between RTs and Advanced Go (gray) and Go (black) trials (Figure S12B), similar to the one observed for the CMT task (see Figure 6 in the main text).

Supplementary Figures and Tables

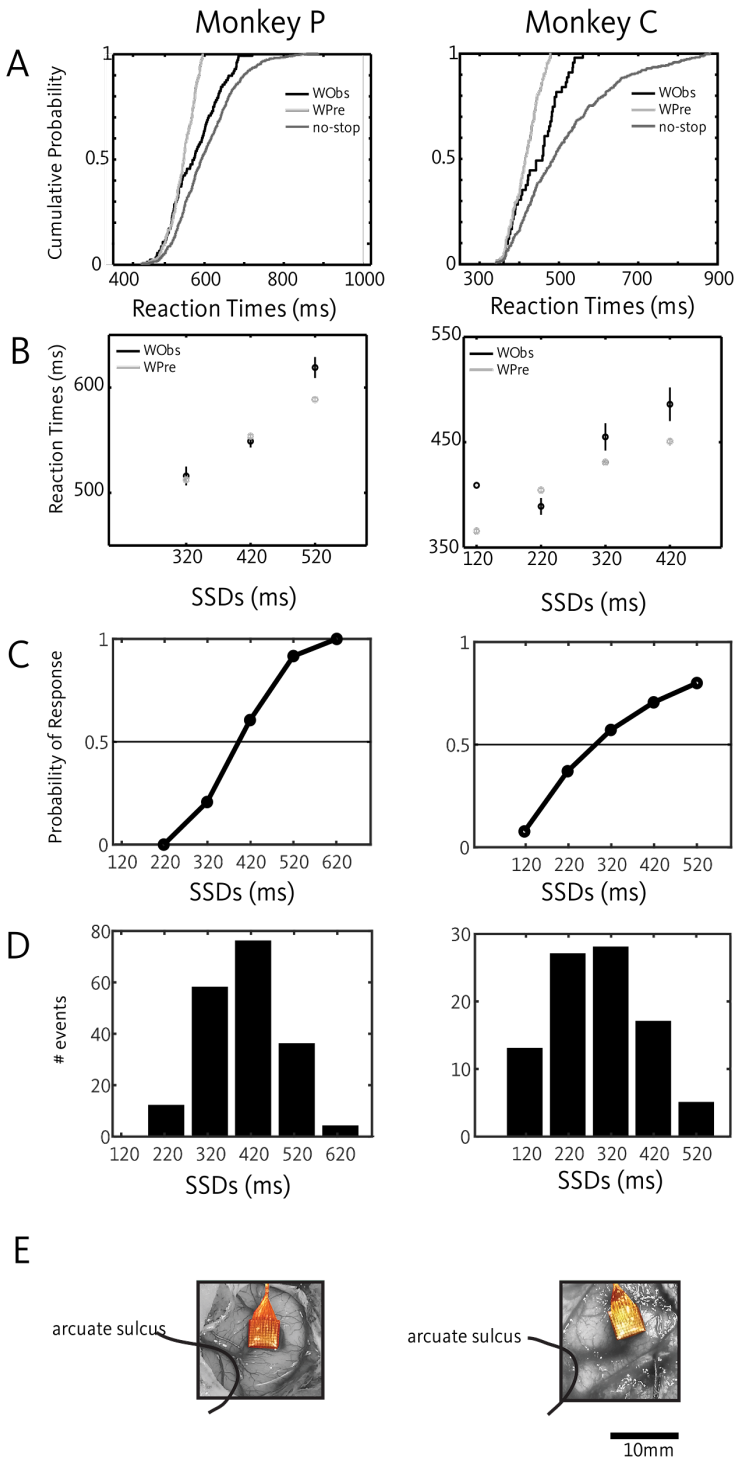

Figure S1. Behavioral performance in the countermanding task for sessions in the main text and recording sites.

A. Validation of the race model: for each monkey no-stop, observed wrong-stop (WObs), and predicted wrong-stop (WPre) cumulative RTs distributions are plotted (top row). **B.** The same data are also shown as functions of the SSDs (bottom row). C. Inhibition functions: probability of responding to a Stop-signal as a function of SSDs. D. Distribution of SSDs presented by employing the staircase procedure. **E.** Pictures of the array's location for monkey P and C.

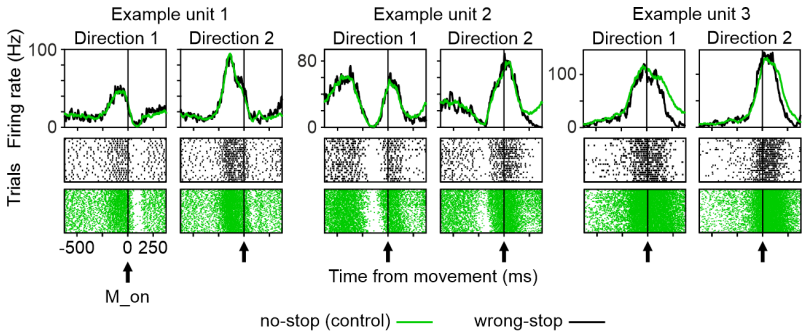

Figure S2. Heterogeneity of single unit modulation during executed and wrongly not cancelled movements.

For the same units of Figure 1, SDFs and raster plots from wrong-stop trials (black) and latency-matched no-stop trials (green) aligned to movement onset (M_on).

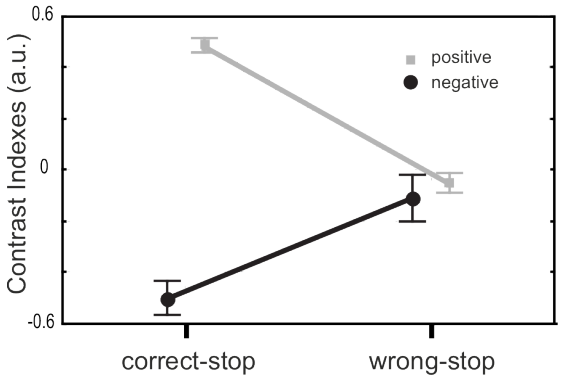

Figure S3. Values of contrast indexes for inhibited (correct-stop trials) vs not-inhibited (wrong-stop trials) movements.

Error bars, 95% confidence interval. (see Supplementary Material and Methods for further details).

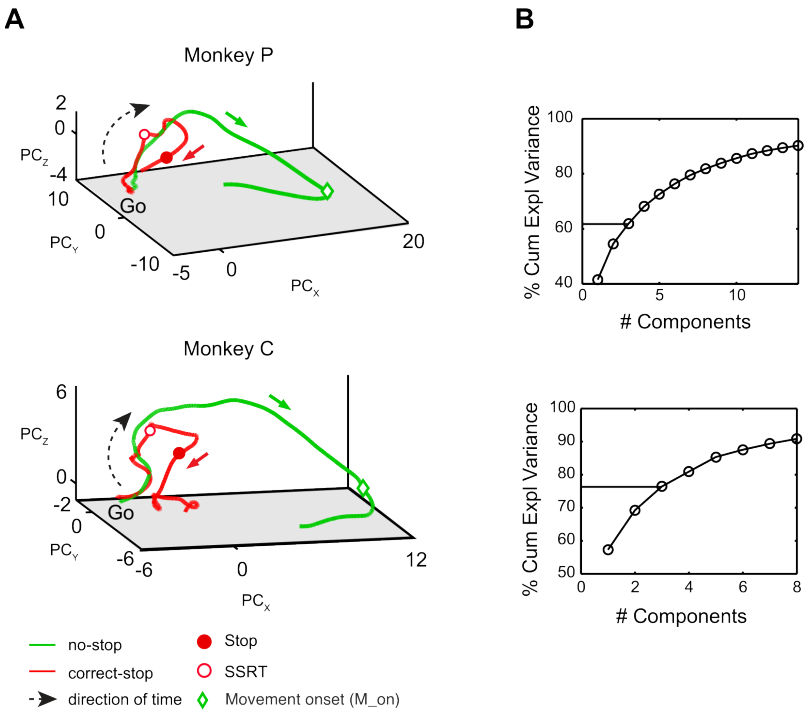

Figure S4. Neuronal dynamics of movement inhibition. Direction 2.

A. Neuronal trajectories in the state space defined by the first three principal components associated to direction 2 as in Figure 2A of the main text. **B.** Cumulative explained variance by the principal components capable of explaining up to 90%. The horizontal line delimits the first three components used to plot the neuronal trajectories. Across all monkeys and sessions, the first 3 components could explain at least 60% of the variance.

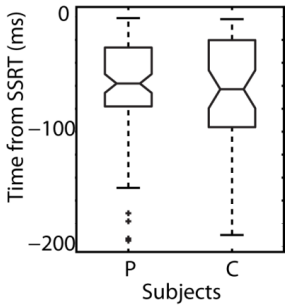

Figure S5. Neuronal modulation occurs before behavioral SSRT. Boxplots illustrate the timing of neuronal modulations (obtained from single unit activities for sessions in the main text) with respect to the behavioral SSRT for both monkeys.

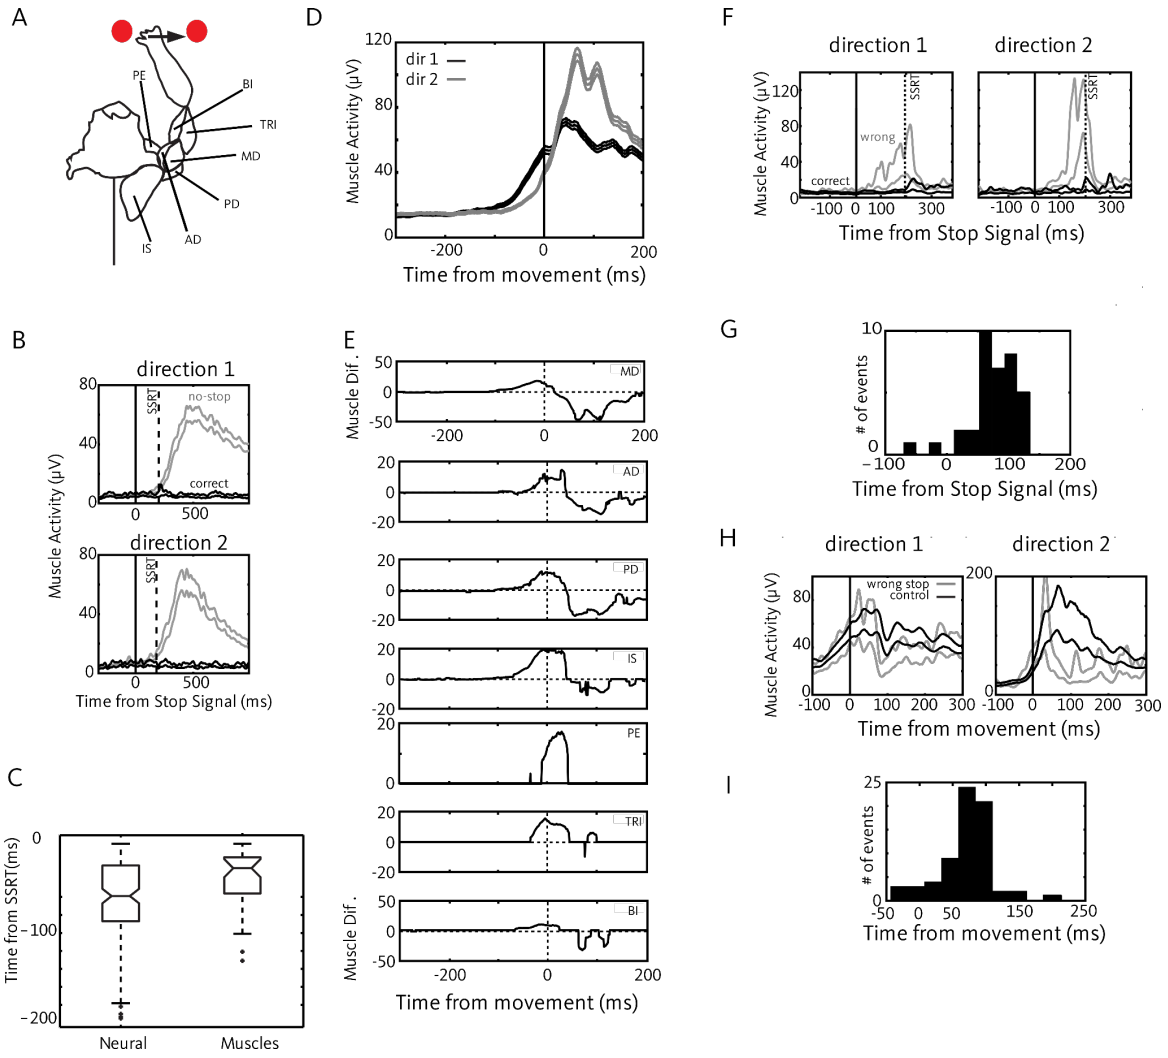

Figure S6. Muscle modulation in the task.

A. Recorded muscles and schematic of the limb postures during the task. Six muscles were recorded: infraspinatus (IS), anterior deltoid (AD), middle deltoid (MD), posterior deltoid (PD), biceps (BI), triceps (TRI) and pectoralis (PE). **B.** Comparison between muscle activities in correct-stop and corresponding latency matched no-stop trials for the infraspinatus in one experimental session (Monkey C). Each panel represents muscle activity for a single movement direction (mean ± SE). **C.** Boxplots of time lags between stop-related neuronal modulations and SSRT (Neural), and time lags between muscle stop-related neuronal modulations and SSRT (Muscles) for both monkeys. **D**. Average activity for the MD across different sessions (n=4; mean (SE)) and directions. **E**. Activity differences between the two movement directions aligned to movement onset for all the muscles recorded across sessions. Values different from 0 represent significant differences (see Supplementary Materials and Methods). **F**. Comparison between correct-stop and wrong-stop trials for the same SSD in the two movement directions (sess. #3, medium deltoid). **G**. Distribution of the latencies of EMG divergences between correct-stop and wrong-stop trials relative to stop signal presentation. **H**. Comparison of an example muscle activity between wrong-stop and control no-stop trials (posterior deltoid, one session), aligned to movement onset. **I**. Distribution of the latencies of the onset of the EMG differences between wrong-stop and no-stop trials after movement onset.

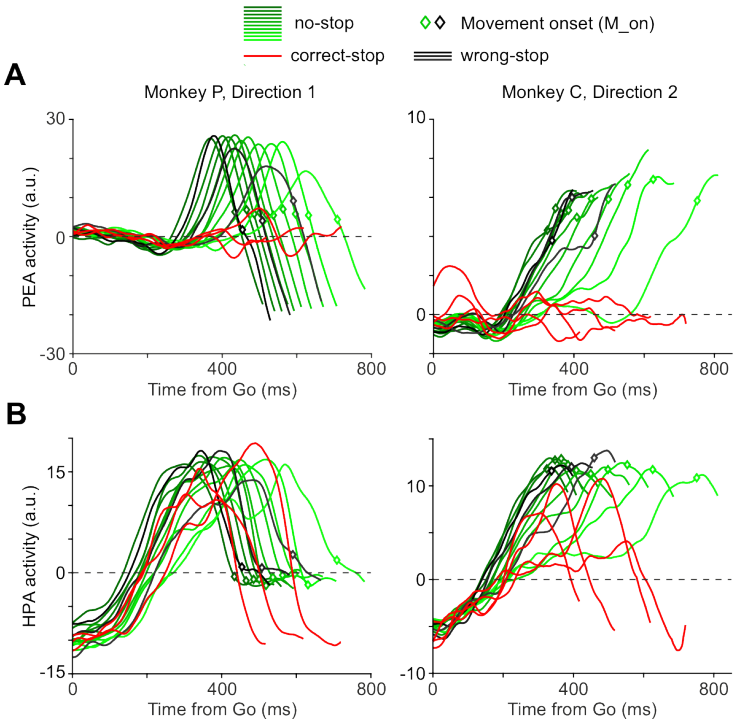

Figure S7. Projections of neuronal activities on the HPA and PEA for the sessions in the main text, Direction 2.

A. Projections on the planning-and-execution axis (PEA) of the neuronal trajectories (for direction 2) grouped by RTs (no-stop in deciles and wrong-stop trials in third intervals) or SSDs (correct-stop trials), for each monkey. **B.** Projections on the holding-and-planning axis (HPA) of the same neuronal trajectories as above. This figure refers to Figure 4 in the main text.

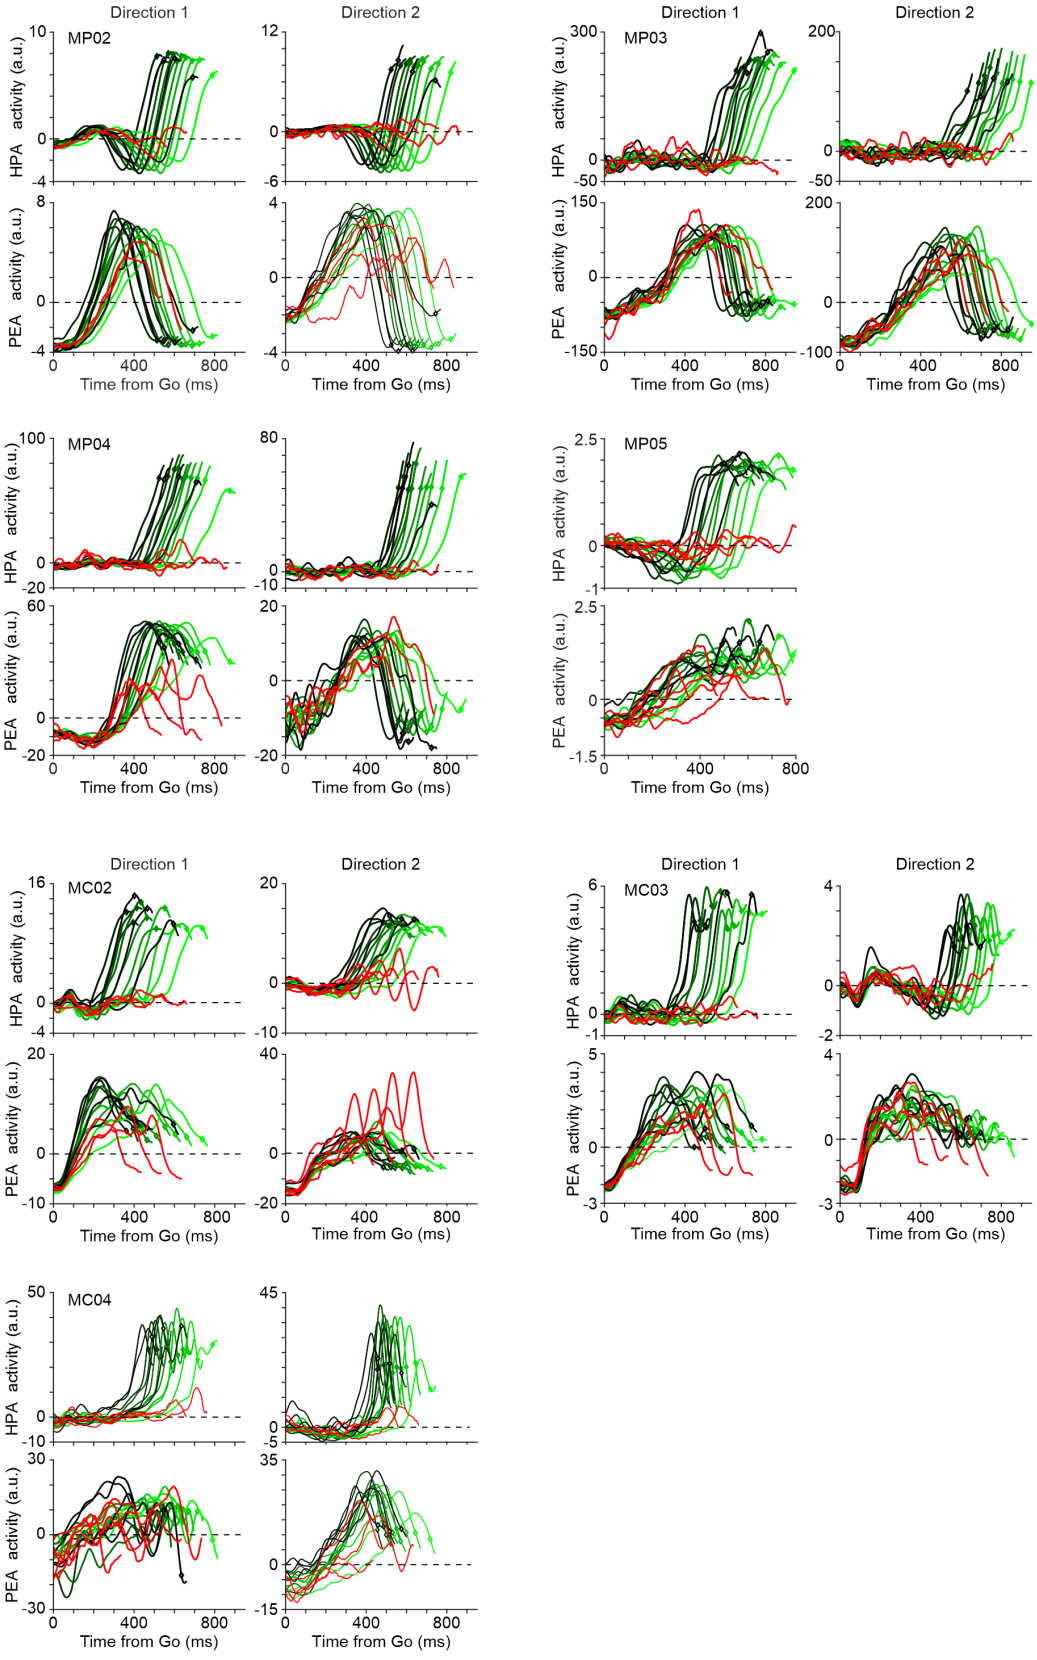

Figure S8. Projections of neuronal activities on the HPA and PEA for all the other sessions. Sessions are labeled as in Table S1. Other conventions and symbols as in Figure 4 and Figure S7-8. Only one direction of movements was used in Session MP05.

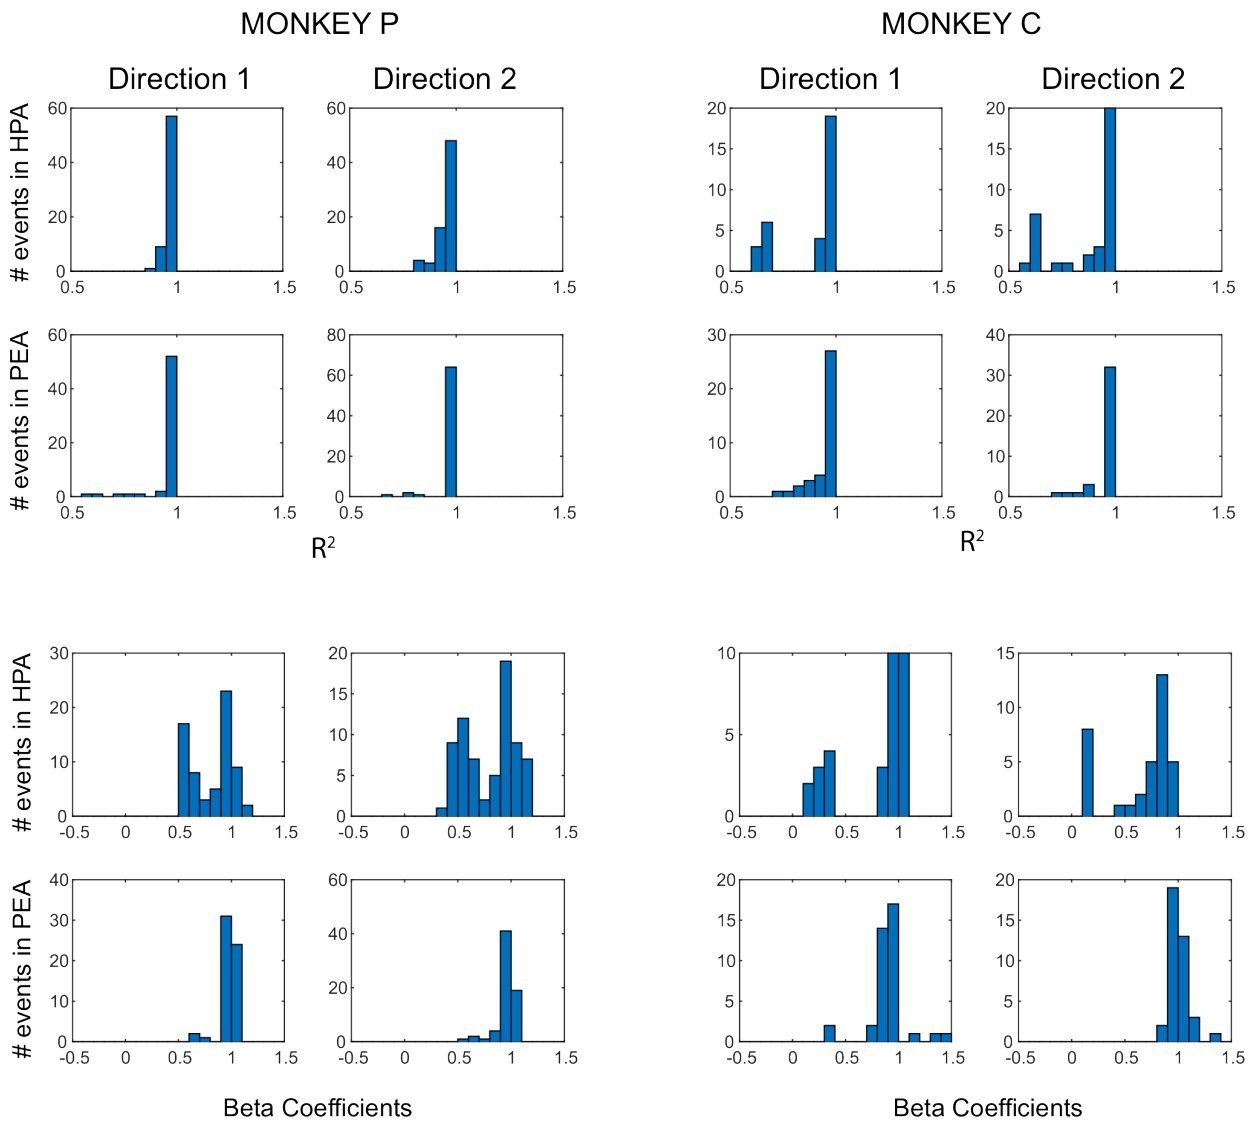

Figure S9. Details of the relationship between PEA and HPA levels of activity, and RTs for the example sessions of the main text.

For each monkey and direction of movement values of R^2 and Beta coefficient of the regression analysis between the values extracted by using the thresholding method (see Supplementary Materials and Methods) and RTs are reported. Mean beta coefficient (interquartile range, IQR): monkey P, direction 1: PEA = 0.99 (0.03); HPA = 0.83 (0.38), $p < 10^{-4}$; direction 2: PEA= 0.97 (0.05); HPA = 0.82 (0.44), $p < 10^{-3}$; monkey C, direction 1: PEA = 0.97 (0.04); HPA = 0.65 (0.44), $p < 10^{-4}$; direction 2: PEA= 0.85 (0.07); HPA = 0.89 (0.41); p = 0.02.

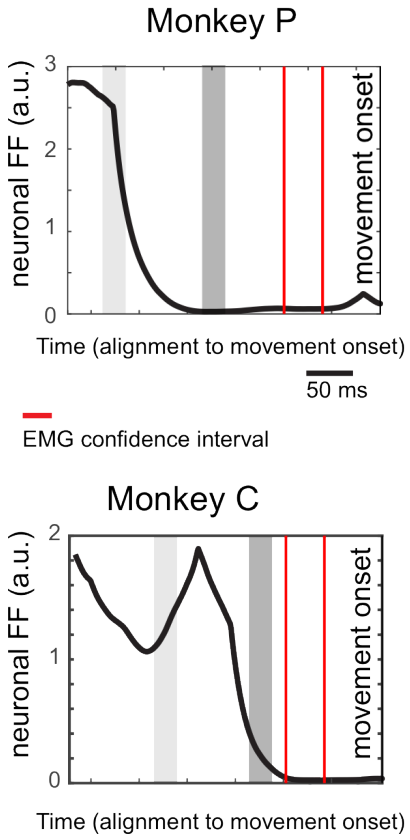

Figure S10. Reduction of variance of the HPA no-stop activity before movement onset.

The neuronal variability (Fano Factor; FF) is computed every millisecond, as the variance of the HPA activity across trials grouped in deciles (for both movement directions) and divided by the mean value. Gray regions represent epochs (40 ms interval) where values of high (light gray) and low neuronal variability (dark gray) have been measured and used for data in Figure 5. The temporal window indicated by the two vertical redlines is the EMG onset confidence interval.

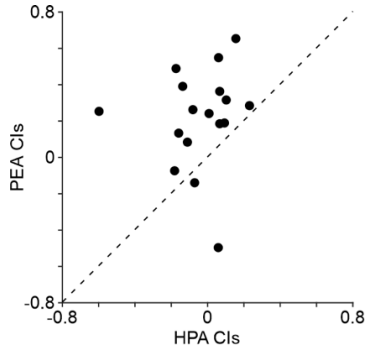

Figure S11: Across sessions the threshold mechanism is apparent only on PEA.

Scatterplot of the Contrast Indexes between no-stop and correct-stop trials in HPA and PEA axes. (see main text for details; n= 17; 2 directions for most of the sessions; 1 direction for MP05).

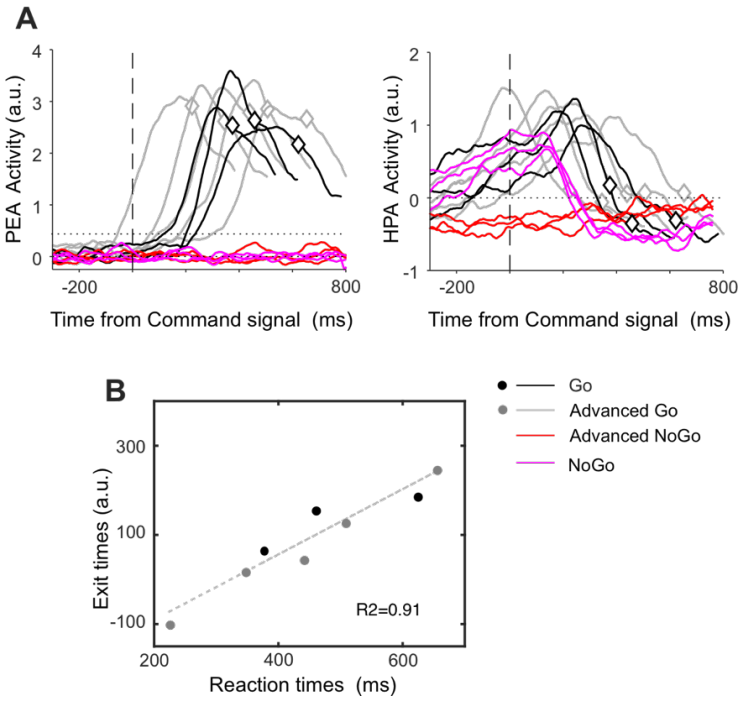

Figure S12: Results of the Context informed Go-NoGo task

A, PEA and HPA dynamics for the different trials. **B**, relationship between exit times (crossing between the dotted line in PEA [0.5 (a.u.)] and activities in Go and Advanced Go trials and corresponding RTs (dark dots: Go trials; gray dots: Advanced Go trials). The empty diamond corresponds to Movement onset (M_on).

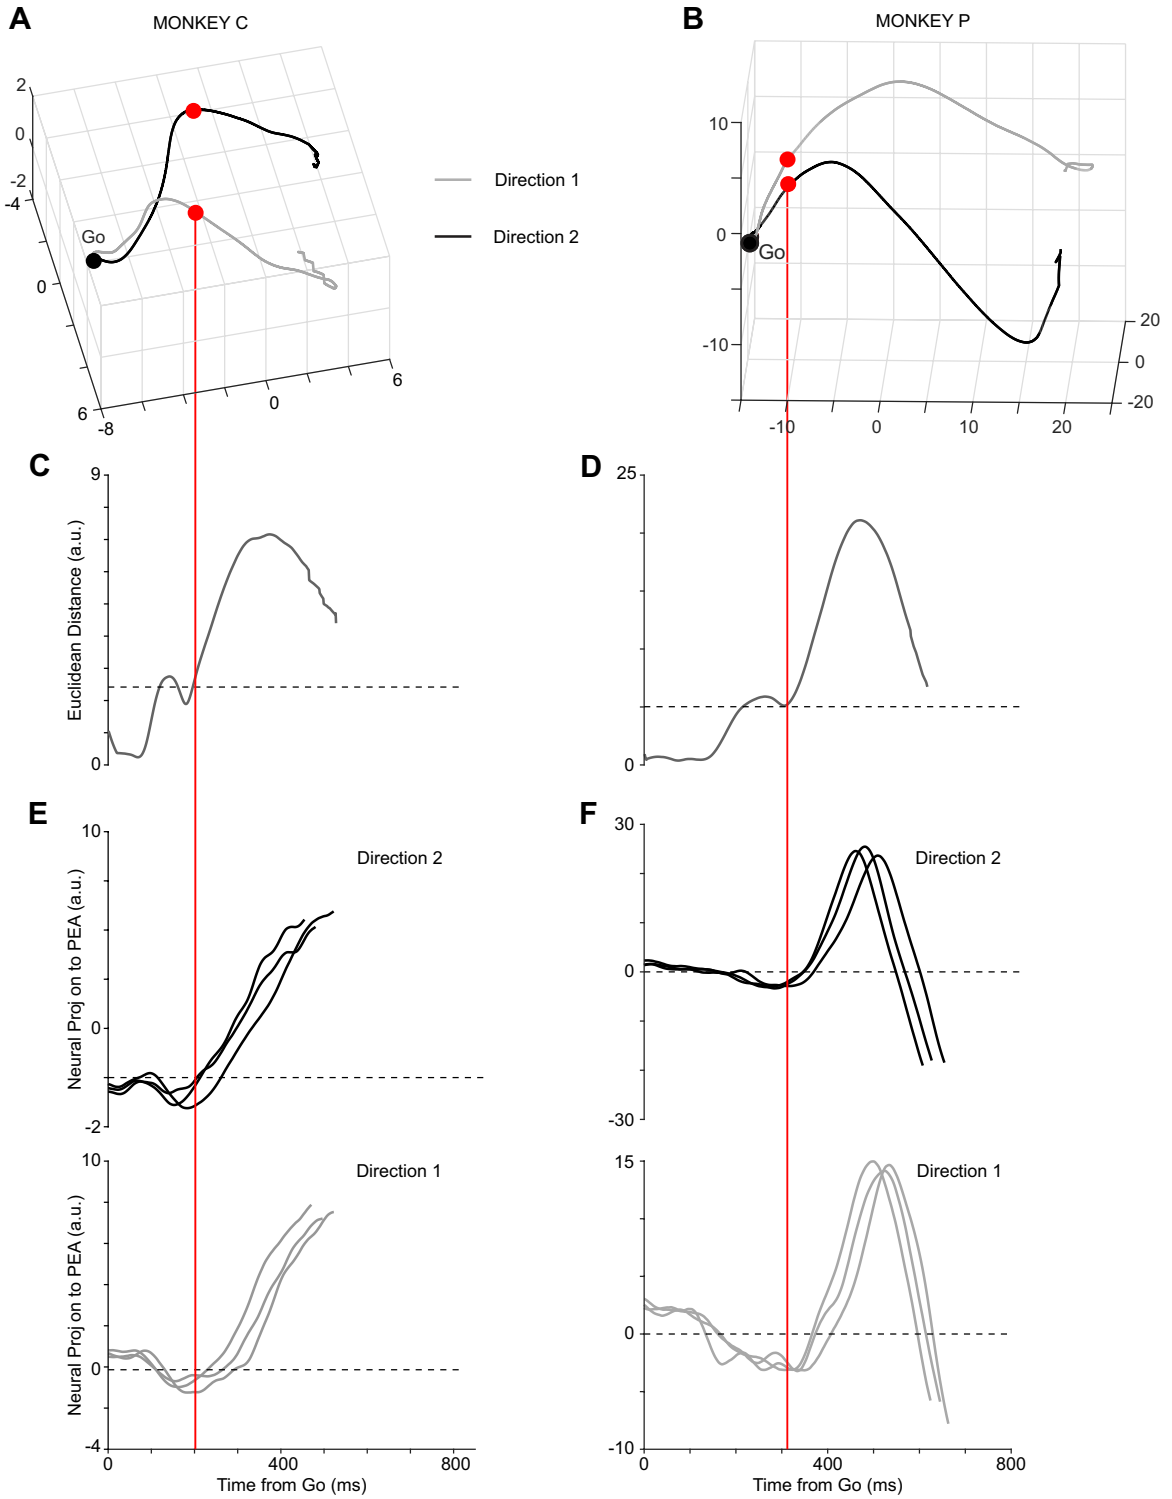

Figure S13. Directional information is reflected in neuronal trajectories

A and **B**: average neuronal trajectories of no-stop trials for the two movement directions in the 3D subspace of the first three PCs. **C** and **D**: Euclidean distances between trajectories for different movement directions: at latency of about 200 ms (300 ms) from the Go signal for monkey C (P) the two trajectories strongly diverged (the horizontal line represents a threshold computed as the mean ± 3SD of the Euclidean distance in the epoch from -50 to +50 ms around the Go signal). **E** and **F**: neuronal activities in the PEA corresponding to the 4th-, 5th and 6th decile of RT, separately for each

direction.

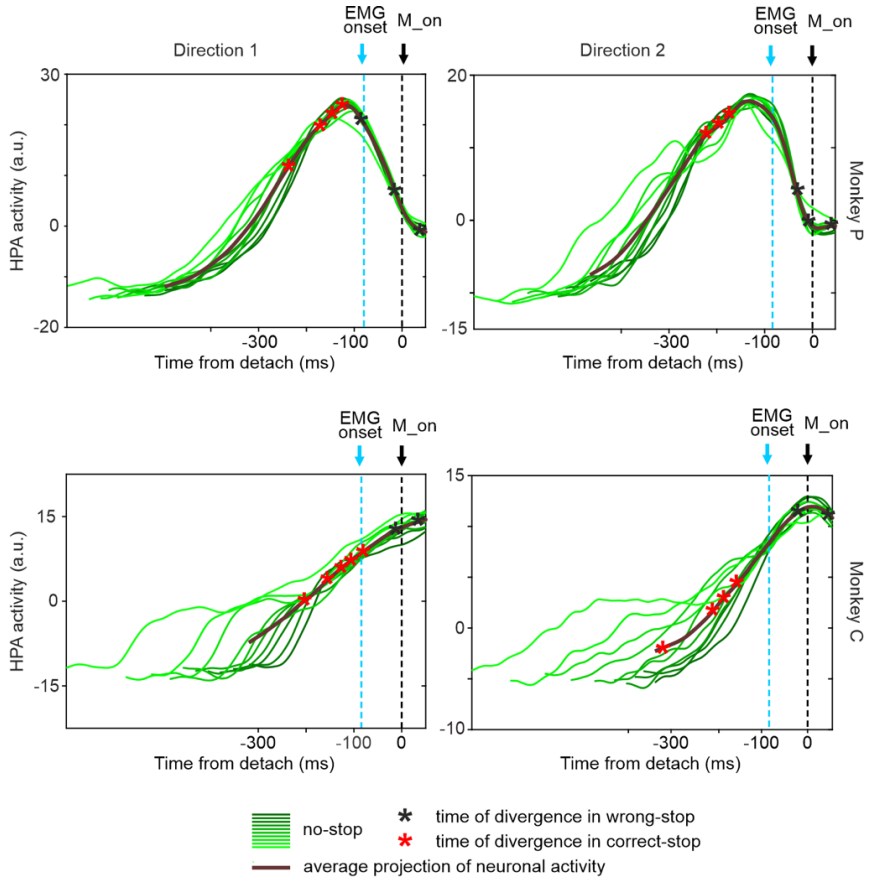

Figure S14. HPA no-stop activities grouped by RTs and aligned to movement onset.

Data from both animals. The solid line represents the average of the activities (also sorted by deciles and in grey levels) projected in the HPA. Red (correct-stop) and black (wrong-stop) dots highlight the level of activity observed at the time of divergence. For correct-stop trials this activity corresponds to the level of activity observed at the time of divergence between correct-stop and no-stop trials. The timing of the neuronal divergence is computed as in Figure 2B relying on single values of SSD for correct-stop and on the average SSD for wrong-stop trials. In correct-stop trials the modulation occurs before muscle activation; in wrong-stop trials the modulation occurs too late to halt the movement.

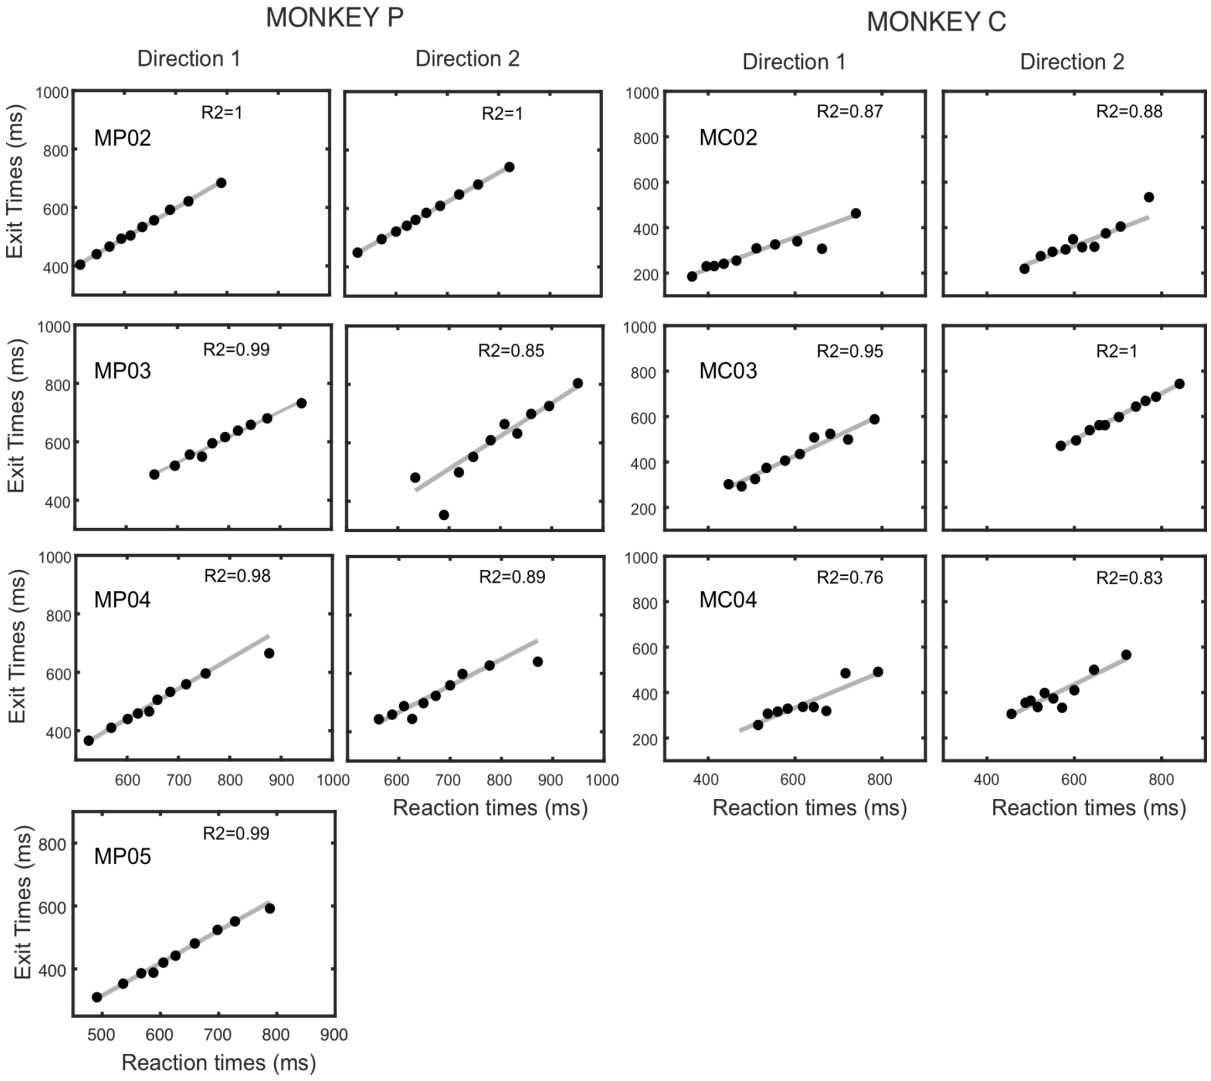

Figure S15: Relationship between exit times and corresponding RTs (as in Figure 6 main text) across sessions and directions.

Monkey P sessions: MP02, MP03, MP04, MP05; Monkey C sessions: MC02, MC03, MC04. In MP05 only one direction was available.

Session	Mean RT(SD)	Mean Wrong(SD)	Estimated SSRT	Mean SSD(SD)	P_{resp}
MC01*	516(124)	447(58)	191	290(11)	0.48
MC02	560(115)	522(90)	174	351(109)	0.40
MC03	637(110)	616(88)	175	447(111)	0.43
MC04	578(92)	534(72)	186	351(104)	0.39
MP01*	603(75)	574(63)	197	399(91)	0.51
MP02	651(94)	623(83)	184	446(89)	0.47
MP03	786(96)	757(63)	197	572(96)	0.45
MP04	667(100)	624(80)	177	456(101)	0.42
MP05	620(92)	584(94)	214	391(104)	0.5

Table S1. Behavioral results (related to Figure 1).

For each session the mean ± SD of the behavioral indexes is reported. All values, except the probability of response P_{resp}), are in milliseconds (ms). Exemplative sessions used in the main text are indicated by *.

Session	Initial pool of isolated units; n =	Units employed to run analysis (Euclidean, Trajectories, ecc.); n= [%]
MC01	91	46 [50.5]
MC02	78	51 [65.4]
MC03	59	37 [62.7]
MC04	69	54 [78.3]
MP01	113	93 [82.3]
MP02	131	93 [71]
MP03	106	64 [60.4]
MP04	82	45[54.9]
MP05	119	89 [74.8]

Table S2. Details of the isolated units and those selected for in depth analysis.

Supplementary Material and Methods

Subjects and apparatus

Two monkeys 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.

Monkeys were pair-housed with cage enrichment (toys, foraging devices). They were daily fed with standard primate chow supplemented with nuts, raisins, and fresh fruits if necessary. The monkeys received their daily water supply during the experiments.

All the surgeries were performed under sterile conditions and veterinary supervision. Antibiotics and analgesics were administered postoperatively. Anesthesia was induced with ketamine (Imalgene, 10 mg kg^{-1} i.m.) and medetomidine hydrochloride (Domitor, 0.04 mg kg^{-1} i.m. or s.c.) and maintained by inhalation isoflurane (0.5–4%) in oxygen (5 l min−1). Antibiotics were administered prophylactically during surgery and postoperatively for at least 1 week. Postoperative analgesics were given at least twice per day. Recordings started after full recovery from surgery (3/4 weeks).

Stimuli were presented on a 17-inch LCD monitor (800×600 resolution; refresh rate 75 Hz) equipped with a touch-screen (MicroTouch, USA; sampling rate 200 Hz). 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 analysis

The race model (10) 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 processes 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.

To validate the race model's independence assumption and compute the SSRT, the mean RT of wrong-stop trials must be shorter than the one of no-stop trials (see for example 14,3,5). Other tests employed to validate such assumption, require that the mean RT in wrong-stop trials should increase with increasing duration of the SSD and that the same mean RT in wrong-stop trials should be equal to those predicted from the race model for each SSD. This second test requires calculating the expected wrong-stop RT from the no-

stop trials distribution of RT. For each SSD, a threshold RT given by the sum SSD + SSRT is considered. The distribution of expected RT for wrong-stop trials is then assumed to be the same as the distribution of no-stop trials RT shorter than the above threshold RT. These predicted RTs are typically tested when the SSD does not dynamically change. Indeed, a reliable estimate of expected RTs requires many trials for each SSD (2,15). Nevertheless, we ran these tests of the independent race model also on our data, although SSD distribution was adjusted depending on the animal performance.

Since we used a LCT monitor and a touch screen with limited temporal resolution we assume that a systematic error could have been added to both the estimation of the SSRT and RT. However, the main goal of our work was not to provide precise evaluation of the behavioral parameters of the monkeys in the countermanding task but to describe the dynamics of the PMd population of neurons involved in the task management. In this respect we believe that a behavioral systematic error has little effect.

Neuronal analysis at the single unit level

We extracted units by employing the spike sorting toolbox KiloSort3 (16): Thresholds: [9 9] (thresholds for template-matching on spike detection); Lambda: 10 (bias factor of the individual spike amplitude towards the cluster mean); Area Under the Curve split: 0.9 (threshold for cluster splitting); Number of blocks: 5 (amount of blocks channels are divided into for estimating probe drift). The output was manually curated in Phy (v2.0; 17) to merge clusters mistakenly separated by the automated sorter.

We selected for in depth 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.

One purpose of the study was to confirm, as previously observed by our group (18), that single neurons in PMd were related to movement inhibition, i.e., if they were differently modulated when the movement was inhibited or when the movement was executed. To this aim we compared the neuronal activity in stop trials to the neuronal activity in the latency-matched no-stop trials (also defined as control trials). For correct-stop trials the latency-matched no-stop trials were those with RTs longer than SSD+SSRT for each specific SSD and session. Latency-matched no-stop trials and correct-stop trials are characterized by the same maturation level of movement planning before the Stop-signal, and their comparison allows one to evaluate what happens when the movement generation is suppressed.

To perform such comparison, the neuronal activity in the latency-matched no-stop trials for each SSD were aligned to the onset time of the hypothetical Stop signal (i.e., the SSD) working out the SDF. The same procedure was used to compare correct-stop and wrong-stop trials.

We then calculated, for each unit, the stop-related neural latency by subtracting the SDF of correct-stop trials from that of latency-matched no-stop trials considering the time (starting from the Stop Signal) when the differential SDF exceeded by 3.5 SD (or fall behind by -3.5 SD) the mean difference in baseline activity calculated from −400 ms to the time of Stop signal, and remained above (below) this threshold for at least 50 ms. We then obtained the time lag between the stop-related neuronal latency and the behavioral estimate of successful inhibition (i.e., the SSRT) to establish when the neuronal modulation occurred relative to the estimated SSRT A negative time lag indicated a stop-related neuronal latency that occurred before the end of the SSRT.

We also calculated a Contrast Index, by contrasting the average firing rate in the 50 ms following the onset of the stop-related neuronal latency in correct-stop trials to the corresponding 50 ms of average firing rate in the control (latency matched; see main text) no-stop trials, as follows:

$$\text{Contrast Index} = \frac{(control\ no{-}stop\ trials) - (correct{-}stop\ trials)}{(control\ no{-}stop\ trials) + (correct{-}stop\ trials)},$$

Using the same approach, a Contrast index was calculated for wrong-stop trials. Note that for wrong-stop trials, the latency-matched no-stop trials were those with RTs shorter than SSD+SSRT.

These Contrast Indexes were used to compare the effect of the Stop signal in inhibited (correct -stop) *vs* not inhibited (wrong-stop) trials. Contrast Indexes were positive (negative) when neuronal activity was higher (lower) in control no-stop trials compared to the corresponding correct-stop and wrong-stop trials. In all neurons with no significant stop-related neuronal latency, the Contrast Index was set equal to 0.

Electromyographic recordings and analysis

Electromyographic (EMG) signal was obtained by needle electrodes (monopolar derivation) inserted into the target muscles with the monkey calm. Muscles recorded were: pectoralis, deltoid (anterior, middle, posterior), infraspinatus, biceps, triceps (Figure S6A).

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.

To establish how the muscles recorded participated in arm movements generation and inhibition, we compared their patterns of activation for movements toward the two possible opposite directions (left and

right). For each muscle we compared all trials with planned movements towards one direction vs trials towards the other direction aligned to movement onset in 10 ms, not overlapping steps (Wilcoxon rank-sum test, p was corrected to 0.0001 for multiple comparisons). From this analysis we obtained a vector of zeros (no difference) and ones (significant difference) to be multiplied by the differential obtained by contrasting the EMG activity of the two directions. Thus, only significant differences had values higher than 0.

To study the muscular modulations that characterize movement inhibition, we selected those conditions with at least 5 correct-stop trials per single SSD. We first established whether in correct-stop trials EMG activity increased respect to the 300 ms preceding the Stop signal (at least 50 ms after stop signal presentation above 2.5 SD of the EMG activity in the 300 ms preceding the stop signal). Then we estimated the electromyographic equivalent of the average time lag between the neuronal modulation after the Stop signal presentation and the SSRT. To this end, we first calculated the differential EMG activity between correct-stop and latency matched no-stop trials and we searched for values above a level for at least 50 ms. The baseline level corresponded to the mean + 3 SD of the differential EMG activity in the 400 ms preceding Stop signal presentation). Once the criteria were respected, for each muscle we considered as electromyographic latency the difference between the SSRT and the first time point above the baseline level.

We also calculated the latency of the divergence in EMG activities between wrong-stop and correct-stop trials. To perform this analysis, we considered all SSDs with at least 6 trials of each type. We calculated as baseline the mean difference between wrong-stop and correct-stop trials in the 400 ms around the time of the Go signal presentation. We then considered as latency of the divergence the first time point after the Stop signal presentation in which the difference between the wrong-stop and the correct-stop trials was above the mean ± 2 SD and continued to stay above this level in the following 75 ms.

To establish the latency of the difference in EMG activation between wrong-stop trials and no-stop trials with similar RT, we subtracted EMG activity of no-stop from the one of wrong-stop trials aligned to the Movement onset. The latency was the first time point when the difference surpassed for at least 150 ms the limit defined by the average difference in the -300 to -100 ms before Movement onset of at least 2 SD.

Moving threshold regression analysis between neuronal projections and RTs.

We performed regression analyses to have an account of the temporal evolution of neuronal activity after the Go signal and to establish the relationship between this evolution and the RT.

These analyses were conducted separately on the neuronal projections in the PEA and HPA, and the corresponding RTs. Starting from the minimum neuronal activity we lifted a threshold activity up to the highest level (50 steps) of activity reached. For each step we extracted the times when the activity crossed these thresholds, and we regressed these times against the corresponding RTs. This was done only if we were

able to obtain at least 7 values of the crossing times. We then calculated the regression slope (Beta coefficient, see Figure S9), and we kept it if the regression analysis was significant ($p < 0.01$). When the projections showed an oscillation (see for example PEA activity in monkey P) we performed twice the procedure to account separately for two groups of crossing. For the initial part of the PEA (see Figure 4A and the relative description on the main text) we observed an almost flat activity or a decay of activity drawing a through that did not show any relationship with RTs.

The Beta coefficients observed were in line with previous works (see 19). We expected to find a broader distribution of slopes for activity showing a ramping like activity: indeed, these profiles usually start with regression slopes well below 1 that increases over time; differently it is possible to have stereotyped responses for which regression slopes are very close to 1 throughout most part of the trial.

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, separately for each direction and monkey, the steps for which the times at which the neuronal projections crossed the thresholds were significantly related to the RTs. We extracted, for each step, the crossing times of the deciles from 5^{th} to 6^{th}, and we then averaged the first ten of them to have an estimate of the time at which the relationship between neuronal projections and RTs started.

We then compared, for each direction and monkey, these time estimates between the PEA and the HPA (Wilcoxon rank-sum test, $P < 0.01$).

Supplementary References

1. 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.
2. Verbruggen F, Logan GD. 2009. Models of response inhibition in the stop-signal and stop-change paradigms. *Neurosci Biobehav Rev* 33:647-661.
3. Band GP, Van der Molen MW, Logan GD. 2003. Horse-race model simulations of the stop-signal procedure. *Acta Psychol* 112:105-142.
4. Verbruggen F, Chambers CD, Logan GD. 2013. Fictitious inhibitory differences: how skewness and slowing distort the estimation of stopping latencies. *Psychol Sci* 24:352-362.
5. Verbruggen F, et al. 2019. A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife* 8:e46323.
6. 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.
7. Buneo CA, Soechting JF, Flanders M. 1994. Muscle activation patterns for reaching: the representation of distance and time. *J Neurophysiol* 71:1546-1558.
8. d'Avella A, Portone A, Fernandez L, Lacquaniti F. 2006. Control of fast-reaching movements by muscle synergy combinations. *J Neurosci* 26:7791-7810.

9. Scangos K W, Stuphorn V (2010). Medial frontal cortex motivates but does not control movement initiation in the countermanding task. *J Neurosci*, 30:1968–1982.
10. Logan GD. 2015. The point of no return: A fundamental limit on the ability to control thought and action. *Q J Exp Psychol* 68:833-57.
11. Kalaska JF, Crammond DJ. 1995. Deciding not to GO: neuronal correlates of response selection in a GO/NOGO task in primate premotor and parietal cortex. *Cereb Cortex* 5:410–28.
12. Riehle A, Grammont F, MacKay WA. 2006. Cancellation of a planned movement in monkey motor cortex. *Neuroreport* 17:281-285.
13. Jerjian SJ, Sahani M, Kraskov A. 2020. Movement initiation and grasp representation in premotor and primary motor cortex mirror neurons. *eLife* 9:1–26. doi:10.7554/eLife.54139.
14. 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.
15. Hanes DP, Schall JD. 1996. Neuronal control of voluntary movement initiation. *Science* 274:427-430.
16. Pachitariu M, Steinmetz NA, Kadir SN, et al. 2016. Fast and accurate spike sorting of high channel count probes with kilosort. In *Advances in Neural Information Processing Systems*, pages 4448–4456.
17. Rossant C, Harris KD. 2013. Hardware-accelerated interactive data visualization for neuroscience in Python. *Front Neuroinform*.9;7:36.
18. 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.
19. Maimon G, Assad JA. 2006. A cognitive signal for the proactive timing of action in macaque LIP. *Nat Neurosci* 9: 948–955.