Behavioral/Cognitive

Common and Unique Inhibitory Control Signatures of Action-Stopping and Attentional Capture Suggest That Actions Are Stopped in Two Stages

Joshua R. Tatz,[1,2,3*] Cheol Soh,[1,3*] and Jan R. Wessel[1,2,3]

[1]Department of Psychological and Brain Sciences, University of Iowa, Iowa City, Iowa 52245, [2]Department of Neurology, University of Iowa Hospitals and Clinics, Iowa City, Iowa 52242, and [3]Cognitive Control Collaborative, University of Iowa, Iowa City, Iowa 52245

The ability to stop an already initiated action is paramount to adaptive behavior. Much scientific debate in the field of human action-stopping currently focuses on two interrelated questions. (1) Which cognitive and neural processes uniquely underpin the implementation of inhibitory control when actions are stopped after explicit stop signals, and which processes are instead commonly evoked by all salient signals, even those that do not require stopping? (2) Why do purported (neuro) physiological signatures of inhibition occur at two different latencies after stop signals? Here, we address both questions via two preregistered experiments that combined measurements of corticospinal excitability, EMG, and whole-scalp EEG. Adult human subjects performed a stop signal task that also contained "ignore" signals: equally salient signals that did not require stopping but rather completion of the Go response. We found that both stop- and ignore signals produced equal amounts of early-latency inhibition of corticospinal excitability and EMG, which took place ~150 ms following either signal. Multivariate pattern analysis of the whole-scalp EEG data further corroborated that this early processing stage was shared between stop- and ignore signals, as neural activity following the two signals could not be decoded from each other until a later time period. In this later period, unique activity related to stop signals emerged at frontocentral scalp sites, reflecting an increased stop signal P3. These findings suggest a two-step model of action-stopping, according to which an initial, universal inhibitory response to the saliency of the stop signal is followed by a slower process that is unique to outright stopping.

Key words: **attentional capture; EEG; EMG; inhibitory control; motor inhibition; stop signal task**

Significance Statement

Humans often have to stop their ongoing actions when indicated by environmental stimuli (stop signals). Successful action-stopping requires both the ability to detect these salient stop signals and to subsequently inhibit ongoing motor programs. Because of this tight entanglement of attentional control and motor inhibition, identifying unique neurophysiological signatures of action-stopping is difficult. Indeed, we report that recently proposed early-latency signatures of motor inhibition during action-stopping are also found after salient signals that do not require stopping. However, using multivariate pattern analysis of scalp-recorded neural data, we also identified subsequent neural activity that uniquely distinguished action-stopping from saliency detection. These results suggest that actions are stopped in two stages: the first common to all salient events and the second unique to action-stopping.

Introduction

Humans are able to use inhibitory motor control to rapidly stop already initiated actions. Neuroscientific research on action-stopping, typically using the stop signal task (SST) (Logan et al., 1984; Verbruggen et al., 2019), has suggested that a three-pronged neuroanatomical network, consisting of right inferior frontal cortex, presupplementary motor area, and subthalamic nucleus, implements an inhibitory control mechanism that purportedly underlies this ability (Aron et al., 2007). However, recent years have seen substantial controversy regarding the purported role of this network in uniquely instantiating inhibitory control during action-stopping. For example, events that result

Received May 26, 2021; revised Aug. 3, 2021; accepted Aug. 26, 2021.

Author contributions: J.R.T. and J.R.W. designed research; J.R.T. and C.S. performed research; J.R.T., C.S., and J.R.W. analyzed data; J.R.T., C.S., and J.R.W. edited the paper; J.R.T., C.S., and J.R.W. wrote the paper; J.R.W. wrote the first draft of the paper.

This work was supported by National Science Foundation Career Grant 1752355 to J.R.W.; and National Institutes of Health Grant R01 NS102201 to J.R.W.

*J.R.T. and C.S. contributed equally to this work.

The authors declare no competing financial interests.

Correspondence should be addressed to Jan R. Wessel at Jan-wessel@uiowa.edu.

https://doi.org/10.1523/JNEUROSCI.1105-21.2021

in attentional capture (e.g., infrequent or otherwise salient stimuli) activate overlapping brain regions, even when they do not instruct action-stopping (e.g., Sharp et al., 2010; Chatham et al., 2012; Sebastian et al., 2016, 2017; for EEG, see van de Laar et al., 2014; Sánchez-Carmona et al., 2019). Complicating this debate is the fact that the attentional detection of the infrequent stop signals in the SST is paramount to the success of stopping itself (Matzke et al., 2013; Elchlepp et al., 2016). Hence, since action-stopping involves both attentional detection and inhibitory control in close temporal succession, the clear-cut interpretation of neural activity during action-stopping is difficult. This and the fact that action-stopping in tasks like the SST depends strongly on the relative timing of stop- and go-related processes (rather than mere differences in their amplitude) (Logan et al., 1984) suggest that the identification of processes that uniquely index action-stopping necessitates methods with precise temporal resolution.

However, work with such methods has yielded a similarly controversial picture. Early work using EEG has suggested that the stop signal P3 event-related potential indexes inhibitory control during stopping, as its properties align with the horse-race model of action-stopping and stopping behavior itself (de Jong et al., 1990; Kok et al., 2004; Wessel and Aron, 2015). By contrast, more recent work has claimed that this activity occurs too late after the stop signal to contribute to stopping success and that earlier neural activity must hence reflect the critical processes (e.g., Huster et al., 2020; Jana et al., 2020; Skippen et al., 2020). Indeed, measurements of corticospinal excitability (CSE) using concurrent EMG and transcranial magnetic stimulation (TMS) have shown that the motor system is already broadly suppressed at ~150 ms after visual stop signals (i.e., before purported neural signatures of inhibition, e.g., the P3, emerge) (e.g., Majid et al., 2013; Wessel et al., 2013; Jana et al., 2020). "Broadly" here refers to CSE suppression that is found even at task-unrelated effectors (compare Wessel and Aron, 2017). Similar evidence comes from movement-related EMG in the responding effector itself: on successful stop trials that contained subthreshold activity of the to-be-stopped muscle, that activity begins to decline around the same latency (Raud and Huster, 2017; Raud et al., 2020). However, parallel work on the processing of salient events outside of stop signal situations suggests that the interpretation of these purported early-latency inhibitory signatures is affected by the same problem as the neuroimaging literature: salient events lead to similar suppression of CSE even when stopping is not required (Novembre et al., 2018, 2019; Dutra et al., 2018; Iacullo et al., 2020).

The current study aimed to resolve these contradictions and provide a definitive answer to the question of which neurophysiological signatures are unique to action-stopping and which are common to any type of salient event.

Across two preregistered experiments, we measured all three abovementioned neurophysiological candidate signatures of inhibitory control during action-stopping: the nonselective reduction of CSE after stop signals in task-unrelated muscles, the concurrent suppression of subthreshold EMG in the stopped muscle itself, and event-related whole-scalp EEG. Participants performed a stimulus-selective SST (Bissett and Logan, 2014), which includes a type of trial on which the go signal is followed not by a stop signal, but by an equally salient signal that does not instruct stopping and has to be ignored. We hypothesized that such "ignore" signals would trigger the same early-latency reduction of EMG and CSE as stop signals. We furthermore hypothesized that multivariate pattern analysis (MVPA) of whole-scalp EEG would reveal no decodable differences between stop- and ignore signals during this initial period. Instead, we hypothesized that the earliest signature of uniquely stop-related neural activity would be the emergence of frontocentral stop signal P3.

Experiment 1

Methods

Hypothesis. The preregistered hypothesis for Experiment 1 was that salient/infrequent events that do not instruct stopping (IGNORE signals) would induce short-latency CSE suppression, and that that degree of that initial CSE suppression would not differ significantly from that observed following STOP signals. The preregistration document can be found at https://osf.io/p6c3r.

Participants. Thirty-seven healthy, right-handed adults (19 female, mean age = 22.46 years, SD = 4.30 years) participated in the TMS version of the experiment after completing screening questions (Rossi et al., 2011) and providing written informed consent. The participants were recruited via an email sent out to the University of Iowa community, an online system providing research experience to students in psychology classes, or among personnel in psychology laboratories. Participants were compensated at an hourly rate of $15 or received course credit. The study was approved by the University of Iowa's Institutional Review Board (#201711750).

The target sample size ($n = 27$) was determined by an *a priori* analysis computed in G*Power (Faul et al., 2009) for repeated-measures ANOVAs for one group and three measurements. Based on a recent study that showed nonselective CSE suppression following infrequent events (Iacullo et al., 2020), we sought 80% power to identify a medium effect ($f = 0.25$) using a predefined α level of 0.05 (other G*power parameters were left to default). Two additional participants were collected to replace participants with insufficient trial counts, and 8 additional participants were collected to replace datasets in which the stop signal delay (SSD) exceeded 500 ms in at least one block of trials (for details, compare Materials, design, and procedure).

Materials, design, and procedure. The behavioral task was presented using Psychtoolbox (Brainard, 1997) and MATLAB 2015b (TheMathWorks) scripts on a Linux desktop computer running Ubuntu. Figure 1 (top) depicts the behavioral task. Each trial began with a white fixation cross on a black background. After 300 ms, the fixation cross was replaced by a centrally displayed white arrow pointing to the left or right. This represented the GO signal. Participants responded to the direction of the GO signal by pressing the corresponding left or right foot pedal on a Kinesis Savant Elite 2 response device. On two-thirds of trials, termed GO trials, the arrow remained white for 1 s. On one-sixth of trials, the white arrow turned magenta after an adaptive delay (see below). On another one-sixth of trials, the arrow turned from white to cyan after a matched delay. These signals represented the STOP and IGNORE signals, respectively (the color assignment was counterbalanced across subjects). Participants were instructed to try to cancel their response to the STOP signal and to continue responding in the case of an IGNORE signal. Participants were instructed to respond according to the direction of the GO signal and had a maximum deadline of 1 s to do so. If no response was made, the text "Too slow!" was displayed centrally in red for 1 s. A stop was considered successful if, following the STOP signal, no response was registered within the 1

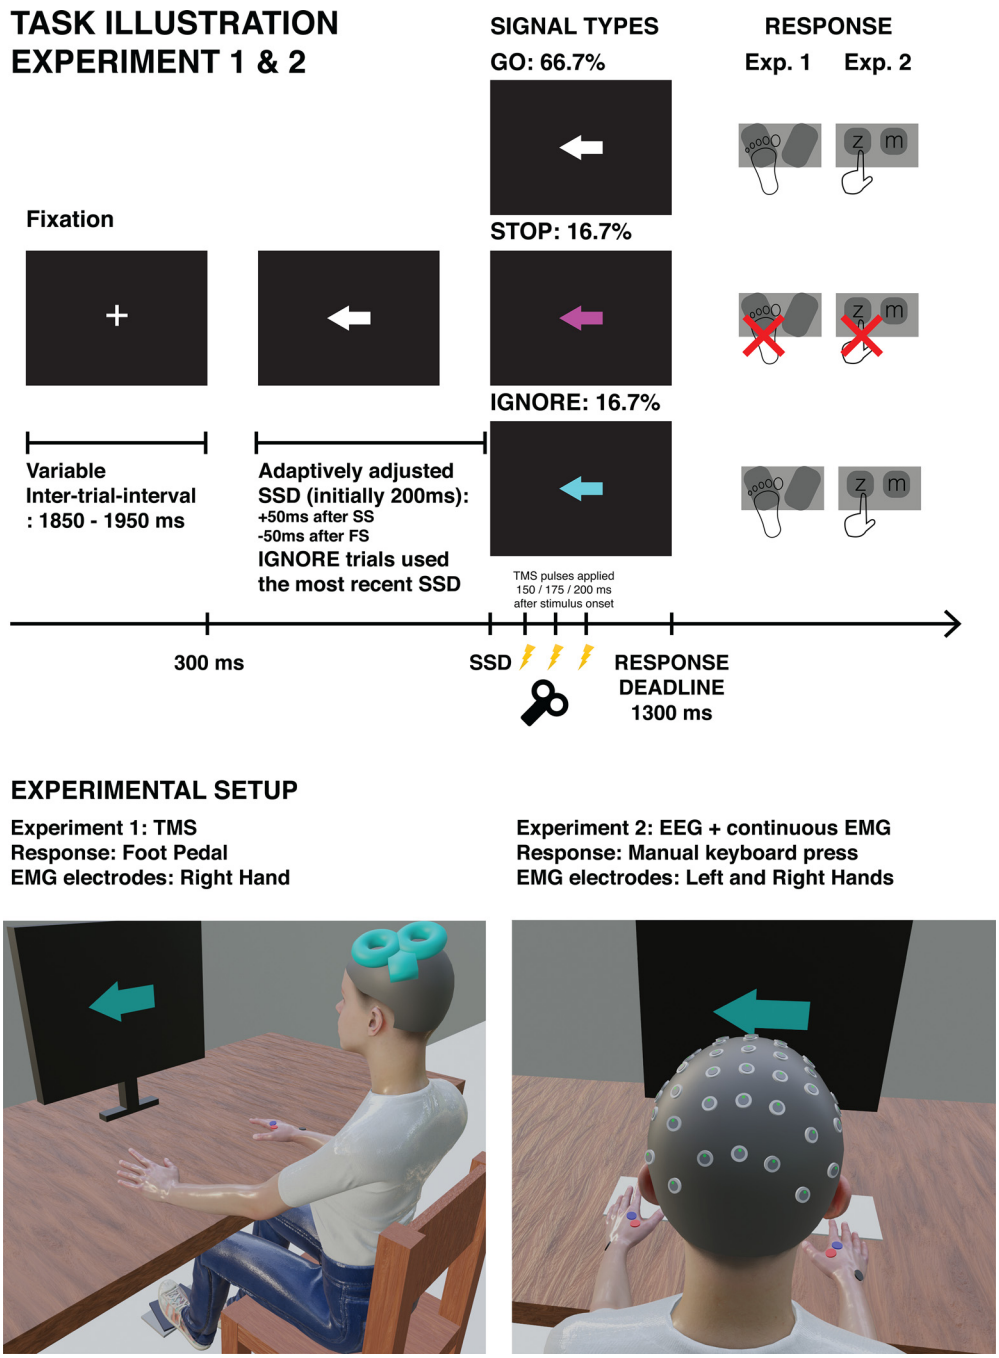

Figure 1. Task diagram. Top, Time course of the behavioral task. Following presentation of the GO stimulus, the arrow changed to magenta or cyan on a minority of trials (16.7% for each color). These signaled to the participant to try to cancel their response (STOP signal) or to continue with the indicated response (IGNORE signal). Bottom left, The setup for Experiment 1, which involved recording MEPs from the right hand while the feet were used to respond. Bottom right, The setup for Experiment 2, in which we recorded EEG from the scalp and EMG from both hands involved in the task.

s response deadline. No immediate feedback was provided following STOP signals, nor was immediate feedback provided following GO and IGNORE signals (except on misses, i.e., when no response was made on either of these trial types within 1 s). The delay of the color change, known as the SSD in the case of the STOP signal, was initialized at 200 ms and from there adjusted adaptively based on stopping success. If the participant was successful in withholding a response following the STOP signal, the SSD was increased by 50 ms for subsequent trials. If the participant made a response after presentation of the STOP signal (but within the response deadline), the SSD decreased by 50 ms with the constraint that the SSD could not fall below 50 ms. The delay for the IGNORE signal was inherited from the most recent STOP signal. Immediately following the 1 s response deadline, the arrow was again replaced with the fixation cross. The cross was presented for 1.7 s plus a variable intertrial interval ranging from 150 to 250 ms (sampled from a uniform distribution in 25 ms increments).

Before the main experiment, each participant completed a practice block of 36 trials. No TMS was delivered during the practice block, and these data were not included in the analyses. The main experiment featured a 3 (Signal Type: GO, IGNORE,

STOP) × 3 (Stimulation Time: 150, 175, 200 ms) design with both Signal type and Stimulation Time as within-subject factors. A random permutation function was used to achieve a pseudorandomized trial order with respect to Signal type and Stimulation Time. The constraints were that the first three trials were necessarily GO trials and that no two Infrequent trials (IGNORE or STOP) were presented consecutively. The arrow direction was randomly determined for each trial (sampling from a Bernoulli distribution). Participants completed 1080 trials in total. Single-pulse TMS was delivered on 840 of those trials, with 180 pulses corresponding to STOP trials, 180 to IGNORE trials, and 240 to GO trials. TMS was delivered 150, 175, or 200 ms after stimulus onset in equal numbers for each trial type. Thus, TMS was delivered on each STOP and IGNORE trial. Another 240 trials were dedicated to obtaining active baseline MEP measurements, during which TMS was delivered before the onset of the GO signal on GO trials. The remaining 240 trials (without TMS) were all GO trials that preceded the active baseline GO trials. These trials did not have TMS to allow adequate time for the TMS stimulator to recharge between pulses. Active baseline and no TMS trials were selected randomly from the list of all consecutive GO trials (with the constraint that equal numbers of 150, 175, and 200 ms stimulation times were preserved for regular GO trials). Because of an initial programming error, the first participant received 672 GO trials and only 24 active baseline trials (along with the planned 180 STOP trials and 180 IGNORE trials that everyone received). As this participants' data were usable and representative in all other respects, their data were not excluded from the analyses.

The 1080 trials were partitioned into 10 blocks of 108 trials each. In between blocks, participants were given as much rest time as they desired. In addition, the reaction time (RT), miss rate, and direction errors were displayed to the participant. If necessary, participants were given instructions based on these metrics as well as p(inhibit) and SSD (which were hidden in the decimal places of the other metrics to inform the experimenter without informing the subject). As typical for the SST, the experimenter's instructions aimed to keep the participant around p(inhibit) = 0.5 by encouraging the participant to either respond more quickly or to be more successful at stopping depending on these indicators. Blocks in which p(inhibit) was 0 or 1 or in which the miss (no response following the GO signal) rate was > 25% were excluded from all analyses. This was the case for two blocks in total, from two different participants.

As mentioned above, 8 subjects showed blocks in which the mean SSD became longer than 500 ms (this was the case for one block in 7 subjects and 2 blocks in 1 subject). While the likely reason for such a pattern is that these participants tried to "wait" for the stop signal in those blocks (contrary to the instruction), another potential reason could be that the original code did not register responses made before the stop signal. While the data from Experiment 2 indicated that such responses were exceedingly rare (only a single trial in one single participant showed a suprathreshold EMG burst before the stop signal) and while the short SSD of 205 ms in Experiment 1 made such occurrences unlikely in general, we still chose to replace these 8 participants from the original sample with 8 new participants. Notably, a comparison between the pattern of results from the original 27 participants and the new sample with the eight replacements indicated no substantial difference in the CSE results between the sets of subjects.

Motor-evoked potentials (MEPs). We measured MEPs as an index of CSE. As we were specifically interested in the "global," nonselective CSE suppression that is typical for reactive inhibition in the SST (Badry et al., 2009), we measured CSE from a motor effector that was not involved in the task itself, as is common practice. We recorded EMG from the right first-dorsal interosseous muscle and stimulated the corresponding region of contralateral M1 using TMS. Participants were instructed to keep both hands relaxed and pronated on the desk while data were collected and they were responding to the task stimuli with their feet.

TMS was delivered via a MagStim 200-2 system (MagStim) with a 70 mm figure-of-eight coil. Hotspotting was used to identify the correct location and intensity for each participant. We initially positioned the coil 5 cm left of and 2 cm anterior to the vertex and incrementally adjusted the location and intensity to determine the participant's resting motor threshold. Resting motor threshold was the minimum intensity needed to produce MEPs >0.1 mV in 5 of 10 consecutive probes (Rossini et al., 1994). For the main experiment, the stimulus intensity was increased to 115% of resting motor threshold. The mean experimental intensity was 60.9% of maximum stimulator output.

EMG was recorded using adhesive electrodes (H124SG, Covidien) placed on the belly and tendon of the right first-dorsal interosseous. A ground electrode was placed on the distal end of the ulna. EMG electrodes were passed through a Grass P511 amplifier (Grass Products) and sampled using a CED Micro 1401-3 sampler (Cambridge Electronic Design) and CED Signal software. The EMG was sampled at a rate of 1000 Hz with online 30 Hz high-pass, 1000 Hz low-pass, and 60 Hz notch filters. For GO-, STOP-, and IGNORE signal trials, EMG sweeps were triggered 90 ms before each TMS pulse, and EMG was recorded for 1 s. In all participants, the active baseline pulse was delivered before the start of GO trials (i.e., during the initial fixation period). In most participants, the EMG sweep triggered 10 ms after commencement of the trial and the TMS pulse was delivered at 100 ms. However, an initial programming error omitted this 90 ms delay on the baseline trials in the first 10 participants. Importantly, active baseline MEPs did not differ in amplitude between these 10 participants and the remaining 18 (see additional details in the subsequent paragraph). Thus, no distinction is made in the ensuing analyses.

MEP amplitude was extracted semiautomatically from the EMG trace using ezTMS, a freely available TMS preprocessing tool developed for use in MATLAB (Hynd et al., 2021). MEP amplitude was defined as the difference between the maximum and minimum amplitude 10-50 ms after the TMS pulse. In addition to the automatic determination of MEP amplitude, each trial was visually inspected for accuracy without knowledge of the specific trial type. Trials in which the MEP amplitude was <0.01 mV or the root mean square (RMS) of the EMG trace was >0.01 mV were excluded from the MEP analyses. The first 80 ms of the EMG sweep was used to calculate the RMS power (i.e., for all GO, STOP, IGNORE, and most active baseline trials), whereas for the active baseline trials in the first 10 subjects, the baseline RMS-power was calculated for 81-160 ms after the TMS pulse. These RMS-EMG values (which were only used to identify baseline trials for exclusion for excessive noise) did not differ between the first 10 participants and the remaining 18 participants.

Mean amplitude MEPs were computed for each condition for each participant (as a function of Signal type and Stimulation Time) as well as separately for successful and failed STOP trials at each Stimulation Time. These values were then normalized by dividing them by the mean active baseline MEPs. For each condition, the average number of trials were as follows: 203 (active baseline), 76 (GO/150), 74 (GO/175), 78 (GO/200), 51 (IGNORE/150),

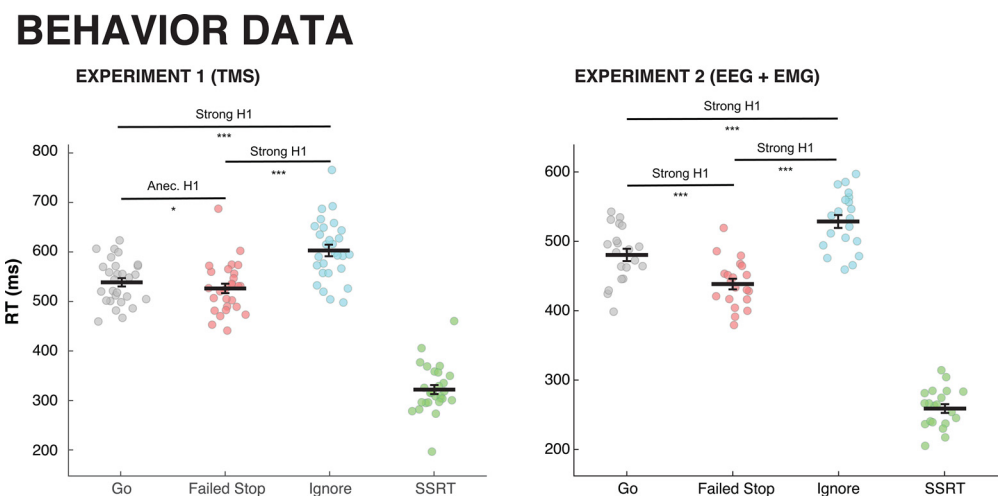

Figure 2. Behavioral RT results from Experiments 1 (left) and 2 (right). Horizontal black brackets represent comparisons made using paired *t* tests. We report both Frequentist (*p* values with Holm–Bonferroni correction) and Bayesian (Bayes factor) results. Horizontal black bars represent the group mean RT. Points represent individual participant mean RTs. Error bars indicate SEM. SSRT was estimated via the integration method (Verbruggen et al., 2019) and not statistically compared with the other, directly observed RTs. ***$p < 0.001$; *$p < 0.05$. Strong H1: BF > 10; Anec. H1: BF > 1, H1: alternative.

52 (IGNORE/175), 45 (IGNORE/200), 27 (failed STOP/150), 28 (failed STOP/175), 25 (failed STOP/200), 26 (successful STOP/150), 24 (successful STOP/175), and 25 (successful STOP/200). Data were excluded from analyses if, after combining successful and failed STOPs, any condition had <20 MEPs (which was the case for 1 participant). Additionally, for analyses that distinguished between successful and failed STOP trials, data were excluded if there were <10 MEPs in any condition (which was the case for 1 different participant).

Data analyses. The behavioral RT data were analyzed with repeated-measures ANOVAs (with Huynh–Feldt correction in instances where the assumption of sphericity was violated) and Bayesian equivalents with noninformative priors in JASP software (Love et al., 2019). Planned and *post hoc* comparisons were made using paired *t* tests with Holm–Bonferroni correction and Bayesian *t* tests. Bayesian analyses were provided in addition to the frequentist statistics because some of our *a priori* hypotheses explicitly stipulated the presence of a null effect. The MEP data were analyzed principally by comparing each Signal type at each Stimulation Time.

Behavioral RT data for mean GO RT, mean IGNORE RT, and mean failed STOP RT were analyzed with a one-way repeated-measures ANOVA. Errors were excluded from the analyses and were generally low for both GO and IGNORE trials (2% and 7%, respectively). Mean STOP signal reaction time (SSRT) was estimated according to the integration method with go-omission replacement, as detailed by Verbruggen et al. (2019). We also examined mean *p*(inhibit) and the difference in failed STOP and GO RT to ensure that the SST portion of our task functioned properly. Finally, we used the prescription of Bissett and Logan (2014) developed for this exact type of experimental paradigm to evaluate the purported behavioral strategy that participants used on IGNORE and STOP trials. This prescription is based on the comparison between RTs on failed STOP versus GO trials as well as between IGNORE and GO trials. Participants that show IGNORE RT > GO RT and failed STOP RT < GO RT are categorized as using a "Stop then discriminate" strategy. Participants that show IGNORE RT < GO RT and failed STOP RT < GO RT are categorized as "Independent Discriminate then Stop." Participants that show IGNORE RT > GO RT and failed STOP RT > GO RT are categorized as "Dependent Discriminate then Stop."

We completed a variety of analyses on the MEP data. First, and most critically for our prediction that MEP suppression would not differ for STOP and IGNORE signals at 150 ms, we compared mean MEPs for GO, STOP, and IGNORE signals at each Stimulation Time. Second, we likewise compared successful STOP, failed STOP, and IGNORE signals at each Stimulation Time. In both cases, we report the results of frequentist paired *t* tests (with Holm–Bonferroni α correction) alongside their Bayesian counterparts. Because of our null prediction, it should be noted that, when reporting Bayes factor (BF), we report BF_{01} (evidence supporting the null hypothesis) or BF_{10} (evidence supporting the alternative hypothesis), depending on which way the evidence pointed.

Data availability. The data files, as well as scripts for running the task and analyzing the data can be found on the Open Science Framework at https://osf.io/b72mp/.

Results

Behavior. The behavioral RT results for Experiment 1 are presented in Figure 2 (left). The comparison for GO, IGNORE, and failed STOP RTs revealed a strong effect ($BF_{10} = 1.12 \times 10^{17}$, $F_{(1.87,50.46)} = 118.58$, $p < 0.001$, $\eta_p^2 = 0.82$). *Post hoc* comparisons indicated that participants responded faster to GO signals (mean: 540 ms, SD = 44.7) than to IGNORE signals (mean: 605 ms, SD = 62.0; $BF_{10} = 1.41 \times 10^{8}$, $t_{(27)} = 12.03$, $p_{holm} < 0.001$, $d = 2.27$). Likewise, participants responded faster during failed STOP trials (mean: 528, SD = 50.7) than during IGNORE signal trials ($BF_{10} = 8.64 \times 10^{12}$, $t_{(27)} = 14.34$, $p < 0.001$, $d = 2.71$). RT was also faster on failed STOP trials than on GO trials ($BF_{10} = 2.23$, $t_{(27)} = 2.31$, $p_{Holm} = 0.025$, $d = 0.44$), and the mean *p*(inhibit) was 0.508 (SD = 0.02). Thus, the assumptions of the race model held (e.g., Verbruggen et al., 2019), and the task was effective at maintaining a *p*(inhibit) at ~0.5. The mean SSRT was estimated to be 323 ms (SD = 48.4), which is in the typical range for foot pedal responses (e.g., Tabu et al., 2012), and mean SSD was 205 ms (SD = 58). Miss rates were 1.33% in the GO and 5.54% in the IGNORE condition, respectively. According to the classification of Bissett and Logan (2014), 19 participants used a "Stop then discriminate" strategy according to their RT pattern, with the remaining 9 participants using "Dependent discriminate then stop," and none using "Independent discriminate then stop."

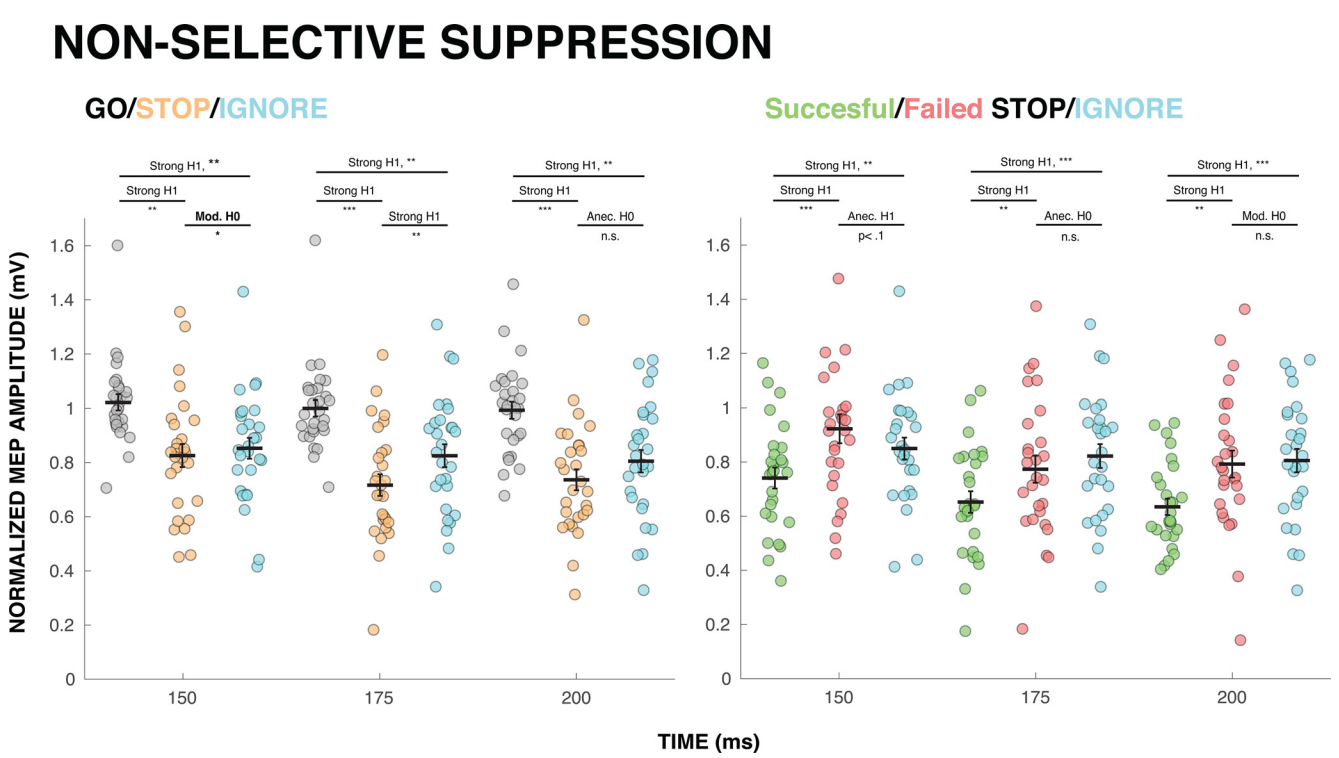

Figure 3. Mean normalized MEPs for each Signal type at each TMS Stimulation Time. Horizontal black brackets represent comparisons made using paired *t* tests. We report both Frequentist (*p* values with Holm–Bonferroni correction) and Bayesian (Bayes factor) results. Horizontal black bars represent the group mean MEP. Points represent individual participant's mean normalized MEPs. Error bars indicate SEM. Left, Comparisons between Go, Stop (including both successful and failed stop trials), and Ignore trials. Right, Comparisons between Successful Stop, Failed Stop, and Ignore trials. ***$p < 0.001$; **$p < 0.01$; *$p < 0.05$; n.s., $p > 0.10$. Strong H1/H0: BF > 10; Mod. H1/H0: BF > 3; Anec. H1/H0: BF > 1; H1: alternative; H0: null hypothesis.

Nonselective CSE suppression. We first aimed to replicate the established findings of nonselective CSE suppression in task-unrelated muscles following both STOP (Badry et al., 2009) and IGNORE signals (Iacullo et al., 2020). As can be seen in Figure 3 (left), global MEP suppression was evident for both IGNORE and STOP signals at every Stimulation Time relative to GO signals (STOP/150: $BF_{10} = 47.81$, $t_{(27)} = 3.84$, $p_{Holm} = 0.005$, $d = 0.73$; IGNORE/150: $BF_{10} = 31.41$, $t_{(27)} = 3.66$, $p_{Holm} = 0.004$, $d = 0.69$; STOP/175: $BF_{10} = 10{,}629.48$, $t_{(27)} = 6.07$, $p_{Holm} < 0.001$, $d = 1.15$; IGNORE/175: $BF_{10} = 31.30$, $t_{(27)} = 3.66$, $p_{Holm} = 0.005$, $d = 0.69$; STOP/200: $BF_{10} = 10{,}634.29$, $t_{(27)} = 6.07$, $p_{Holm} < 0.001$, $d = 1.15$; IGNORE/200: $BF_{10} = 33.56$, $t_{(27)} = 3.69$, $p_{Holm} = 0.008$, $d = 0.70$).

More importantly for the present study, we then directly compared the MEPs for STOP and IGNORE signals to evaluate our prediction that the MEPs would be equivalent, at least at the earliest Stimulation Time of 150 ms. Critically, there was moderate evidence ($BF_{01} = 3.24$, $t_{(27)} = 0.98$, $p_{Holm} = 0.338$, $d = 0.18$) for a null effect at that time point. By contrast, strong evidence for a difference between STOP and IGNORE signals emerged at 175 ms ($BF_{10} = 14.83$, $t_{(27)} = 3.32$, $p_{Holm} = 0.008$, $d = 0.63$), with STOP signals exhibiting greater global MEP suppression. Although there was also numerically greater MEP suppression for STOP signals at the 200 ms Stimulation Time, there was anecdotal evidence to suggest no difference ($BF_{01} = 1.40$, $t_{(27)} = 1.70$, $p_{Holm} = 0.101$, $d = 0.32$). Restricting the analysis to only participants who used the "Stop-then-discriminate" strategy (i.e., excluding the 9 participants who used "Dependent discriminate then stop" according to the Bissett and Logan, 2014 method) did not change the pattern of results. Importantly, these results confirm our hypothesis that both STOP and IGNORE signals produce nonselective CSE suppression and that this effect does not initially differ.

When STOP trials were divided by stopping success, further interesting patterns emerged. First, when distinguishing between successful and failed STOP trials (Fig. 3, right), MEPs were clearly reduced for successful STOP trials at all Stimulation Times, again replicating prior work (Badry et al., 2009; Wessel et al., 2013; Jana et al., 2020). Strong evidence for this difference was present at all stimulation times (150 ms: $BF_{10} = 860.96$, $t_{(26)} = 5.08$, $p_{Holm} < 0.001$, $d = 0.98$; 175 ms: $BF_{10} = 7.02$, $t_{(26)} = 2.98$, $p_{Holm} = 0.006$, $d = 0.57$; 200 ms: $BF_{10} = 10.17$, $t_{(26)} = 3.16$, $p_{Holm} = 0.004$, $d = 0.61$).

We also found greater MEP suppression for successful STOP relative to IGNORE trials at all stimulation times, with Strong evidence present at all stimulation times (150 ms: $BF_{10} = 13.00$, $t_{(26)} = 3.27$, $p_{Holm} = 0.003$, $d = 0.63$; 175 ms: $BF_{10} = 475.73$, $t_{(26)} = 4.83$, $p_{Holm} < 0.001$, $d = 0.93$; 200 ms: $BF_{10} = 73.73$, $t_{(26)} = 4.04$, $p_{Holm} < 0.001$, $d = 0.78$). For failed STOP relative to IGNORE trials, there was Anecdotal evidence of greater MEP suppression on IGNORE trials at 150 ms ($BF_{10} = 1.75$, $t_{(26)} = 2.26$, $p_{Holm} = 0.098$, $d = 0.43$). However, there was Anecdotal evidence for no difference between failed STOP and the IGNORE trials at 175 ms ($BF_{01} = 2.75$, $t_{(26)} = 1.13$, $p_{Holm} = 0.269$, $d = 0.22$) and Moderate evidence for no difference at 200 ms ($BF_{01} = 4.69$, $t_{(26)} = 0.31$, $p_{Holm} = 0.762$, $d = 0.06$).

In summary, these results suggest that IGNORE and STOP trials featured equivalent amounts of nonselective CSE suppression at the early postsignal latency (here, 150 ms), with successful STOP trials representing the part of the STOP trial CSE distribution that happened to contain more suppression compared with failed STOP trials, and IGNORE trials showing amounts of CSE suppression that is between that occurring on successful and failed STOP trials. This suggests that this early activity contributes to the success of stopping, but is not unique to STOP trials, as STOP trials overall showed equal CSE suppression compared with IGNORE trials. Moreover, these results suggest that, in case

of a STOP signal, this initial suppression is sustained further into the postsignal period, as significant differences to IGNORE trials emerged at 175 ms.

Experiment 2

Hypothesis. The hypotheses for this follow-up experiment were registered after the results of Experiment 1 were known. The preregistration document can be found at https://osf.io/2krwy.

Based on the findings in Experiment 1, we made two predictions:

1. The suppression of the overt EMG that is observable on some successful STOP trials (Raud and Huster, 2017; Jana et al., 2020; Raud et al., 2020) is not unique to STOP trials but will instead also occur to the same degree on IGNORE trials.
2. MVPA of whole-scalp EEG will not show any significant differences between STOP and IGNORE signals in the early postsignal time period. However, in line with the CSE results from Experiment 1, we predicted that STOP and IGNORE signals would become decodable from one another at the earliest at 175 ms.

Participants. Twenty healthy, right-handed adults (16 female, mean age: 22.6 years, SD = 2.7 years) participated in the EEG experiment after providing written informed consent. The participants were recruited via an email sent out to the University of Iowa community. Participants were compensated at an hourly rate of $15 or received course credit. The study was approved by the University of Iowa's Institutional Review Board (#201511709).

Task. The task was identical to Experiment 1, with the following exceptions: (1) Responses were made using both index fingers ('z' for left and 'm' for right on a QWERTY keyboard). (2) The total number of trials was reduced to 720 trials that were separated into 10 blocks (66.7% GO, 16.7% STOP, and 16.7% IGNORE trials). While the manual responses are a departure from Experiment 1 (where foot responses were necessary to enable CSE recordings from the hand muscle as a task-unrelated effector), partial EMG (prEMG) research in the SST has hitherto exclusively been performed using manual responses (compare Raud and Huster, 2017). Hence, to most closely replicate past work, participants were made to respond with their hands in Experiment 2.

EEG and EMG recording. EEG was acquired using a Brain Product ActiChamp 63-channel system. The ground and reference electrodes were placed at AFz and Pz, respectively. EMG was recorded from each hand using two electrodes placed on the belly and tendon of each first-dorsal interosseous muscle. Ground electrodes were placed on the distal end of each hand's ulna. The EMG electrodes were connected to the EEG system using two auxiliary channels (via BIP2AUX adaptor cables). Thus, there were 65 total recording channels. All channels were sampled at a rate of 2500 Hz. To minimize task-unrelated EMG activity, the participants were instructed to only move their index fingers when necessary and to otherwise keep both hands relaxed and pronated on the desk.

EEG and EEG preprocessing. Custom MATLAB code was used to preprocess the EEG and EMG data. The EMG data analysis was adapted from Raud and Huster (2017). Both EEG and EMG data were bandpass filtered (EEG: 0.5-50 Hz; EMG: 2-200 Hz) and downsampled to a rate of 500 Hz. The continuous EEG data were visually inspected to identify nonstereotypical artifacts. Segments with nonstereotypical EEG artifacts were removed from both the EEG and EMG data. The EEG data were rereferenced to the common average and submitted to an infomax independent component algorithm (Makeig et al., 1996) as implemented in the EEGLAB toolbox (Delorme and Makeig, 2004). The resulting independent components were visually inspected to remove components that captured stereotypical artifacts (i.e., blinks and saccades). Both the EEG and EMG data were epoched relative to the onsets of GO, STOP, and IGNORE stimuli (−200 to 1000 ms). For GO and IGNORE trials, all trials with incorrect or absent responses were excluded from all analyses. The EMG data were then converted to RMS power using a sliding window of ±5 sample points and baseline-corrected by dividing the entire epoch by the mean of the 200 ms prestimulus EMG. The resulting data were then standardized for each hand across all types of trials and sample points using the *z*-transform.

Behavior analysis. Behavior data were analyzed identically to the behavior data in Experiment 1.

EMG validation. As a manipulation check, we examined whether the EMG data accurately reflected motor output by correlating the latency of each trial's GO trial EMG peak to GO RT on the same trial within each subject. If EMG and GO RT are meaningfully related as predicted, the mean individual-subject correlation should be significantly larger than zero. Pearson correlation coefficients (r) were calculated for each subject at the single-trial level. For those correlations, trials in which Cook's D exceeded 4/N were deemed statistical outliers and were not included in calculating the individual r's (on average, 6.2% of trials per subject). To ensure normality, the individual r's were transformed using Fisher's z transformation. The individual, z-transformed r values were then submitted to a one-sample t test against 0 on the group level. To enable their original interpretation as correlation coefficients, the averaged z-transformed r values were then transformed back to r (see Silver and Dunlap, 1987).

prEMG trial identification. Previous research has shown that residual EMG on successful STOP trials shows a marked reduction at ~150 ms following the STOP signal (Raud and Huster, 2017; Raud et al., 2020). We here hypothesized that the same would be true for IGNORE trials. This analysis was again adapted from Raud and Huster (2017). For each subject, we limited the time window of successful STOP EMG traces to 0-200 ms and averaged the time-constrained EMG traces across trials. Using the subject-specific EMG traces, a detection threshold was identified by finding the peak amplitude of the averaged EMG trace in the first 200 ms following the signal for each subject: [mean threshold (z score) = 0.3, SD = 0.39]. Within each subject, a trial was then denoted a prEMG trial if the individual trial peak in the same window exceeded the threshold. The same procedure was used to identify prEMG on IGNORE trials, except that all IGNORE trials with correct responses were used, as they all contained prEMG activity.

EMG analysis. Figure 4 (top left) depicts the grand-averaged EMG traces for successful STOP, failed STOP, and IGNORE trials, separately for the selected and unselected hands. Statistical analyses were performed on EMG data from the selected hand. Both successful STOP and IGNORE trials showed an initial pattern of ramping activity, which was at ~150 ms following either signal, interrupted by a sudden downturn. In STOP trials, the EMG trace subsequently returned toward zero, in line with the outright cancellation of the response. In IGNORE trials, the downturn was followed by a second ramping, in line with a reactivation of the response. Consistent with prior studies using only STOP trials, the variable of interest was the latency of the downturn after the initial ramping, which is

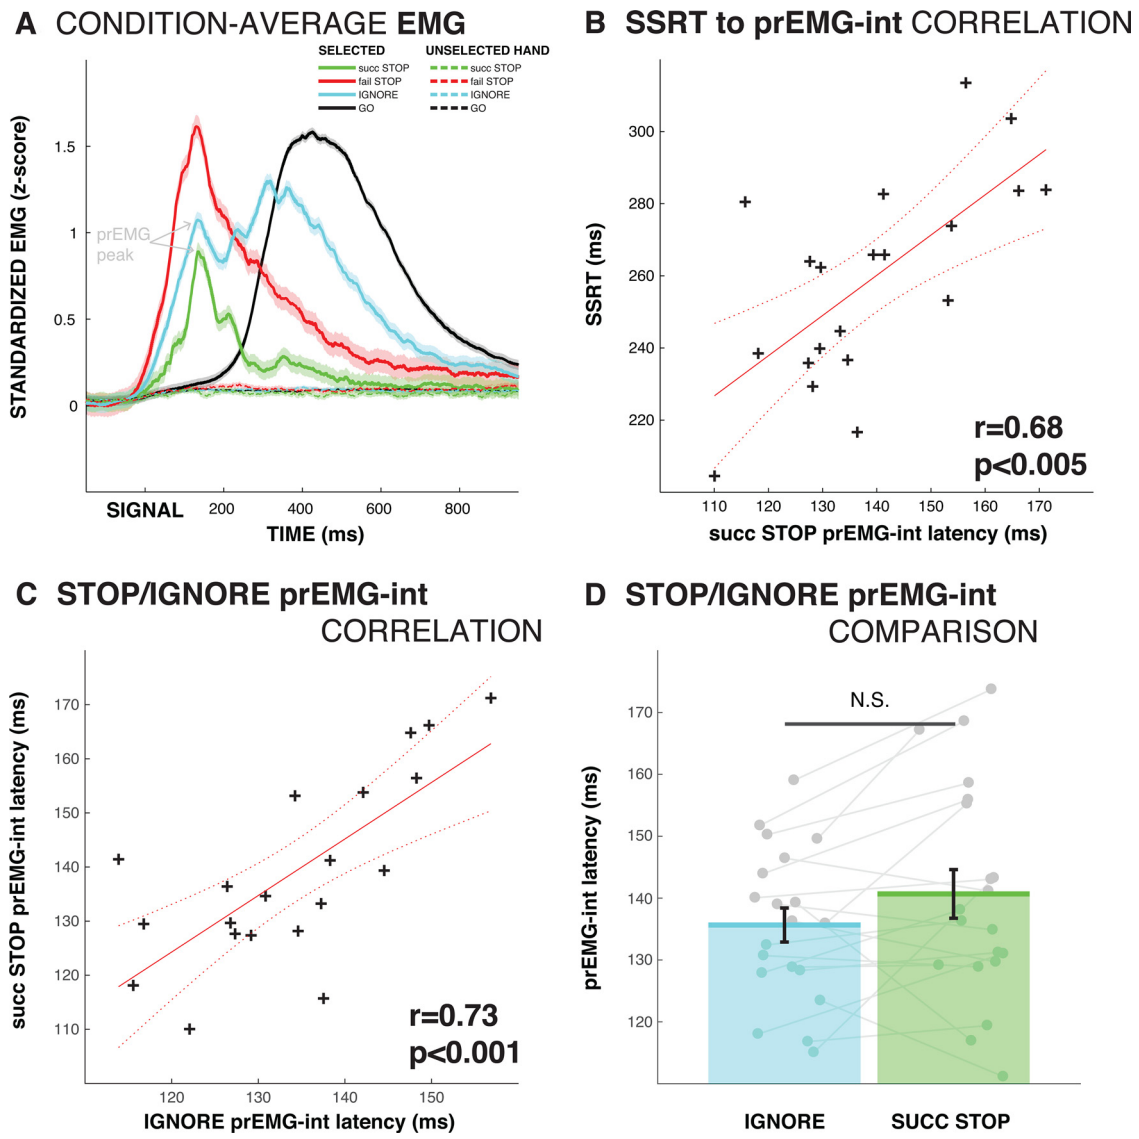

Figure 4. EMG data results. Top left, ***A***, EMG traces for all conditions. Arrows on the left indicate prEMG-int latency for the successful STOP and IGNORE condition. Top right, ***B***, Group-level correlation between successful STOP prEMG-int latency and SSRT. Bottom left, ***C***, Group-level correlation between successful STOP and IGNORE prEMG-int latency. Bottom right, ***D***, Successful STOP and IGNORE prEMG-int latency comparison.

taken as an indicator of early inhibitory activity (Raud and Huster, 2017; Jana et al., 2020; Raud et al., 2020). To this end, the peak of the EMG trace was detected within the first 200 ms after successful STOP and IGNORE signals and its latency averaged for each subject and condition separately. We refer to this latency, which was the DV of interest for the EMG analysis, as the prEMG-int (prEMG interruption) latency.

To test Hypothesis 1 for Experiment 2 (see Hypothesis, above), first, we correlated mean STOP trial prEMG-int latency with IGNORE trial prEMG-int latency at the group level to examine the extent to which both measures related to each other. Furthermore, we compared both latencies with a paired *t* test. Since we hypothesized that the same interruption that has been observed on STOP trials would also occur on IGNORE trials, we predicted that these latencies would be highly correlated across subjects and that their condition means would not differ. As with Experiment 1, we report BF alongside frequentist statistics for all EMG analyses in Experiment 2.

We also examined the relationships between prEMG and behavioral RT measures. First, we correlated STOP prEMG measures with mean SSRT to replicate prior work that showed a strong correlation between the two (e.g., Raud and Huster, 2017). For IGNORE trials, we used the same single-trial level approach to correlate prEMG measures with RT that was described for the GO trial analysis in the EMG validation section above. To provide an unbiased estimator of IGNORE RT, it was quantified from the onset of the IGNORE signal, not of the GO signal (although the results were qualitatively unchanged when the full RT was taken instead).

EEG MVPA. To test Hypothesis 2 for Experiment 2, we investigated at which time point an MVPA classifier applied to the whole-scalp EEG data recorded during the task would be able to distinguish between the conditions of interest above chance, with a particular focus on the first time point at which whole-scalp EEG showed above-chance decoding of STOP versus IGNORE trials.

As control analyses (and to illustrate the performance of the classifier approach), we also examined time points during which successful STOP, failed STOP, and IGNORE trials differed from matched GO trials. To do so, GO trials were split into fast and

slow GO trials by each subject's median GO RT for each hand. Fast GO trials were matched to failed STOP trials (which represent the fast part of the GO RT distribution) (see Logan et al., 1984) and slow GO trials were matched to successful STOP and IGNORE trials (since successful STOP trials represent the slower part of the RT distribution, and IGNORE RT was longer than regular GO RT; compare behavioral results for Experiments 1 and 2). This matching procedure ensured that decoding was not confounded by differences in RT. In addition, the number of left and right arrow trials were balanced for all decoding analyses to ensure nonbiased results.

MVPA analyses were performed using the ADAM toolbox (Fahrenfort et al., 2018). We used leave-one-out cross-validation in which a Linear Discriminant Analysis classifier was trained on all trials but one and tested on the remaining trial that was not part of the training set. This validation was implemented in the sample pointwise fashion, in which training and testing were done at each sample point following the GO/STOP/IGNORE signal. Classifier performance was measured as the area under the receiver operating characteristics curve (AUC) (Wickens, 2002), which quantified the total AUC when the cumulative true positive rate (probability of correct classification) was plotted against the false positive rate (probability of incorrect classification) for each decoding problem between the conditions of interest.

Time series of AUC values were obtained for each pair of conditions and each participant. To identify time points during which classification performance was significantly above chance, we compared each AUC value against chance-level performance (0.5) using one-sample *t* tests, which was corrected for multiple comparisons using cluster-based permutation testing (10,000 iterations, cluster-*p* value = 0.00001, individual sample point $\alpha = 0.0001$) (Maris and Oostenveld, 2007).

Event-related potential analysis. Since the frontocentral ERPs are commonly associated with stop signal performance (e.g., de Jong et al., 1990; Huster et al., 2020), we also quantified the condition ERPs at the frontocentral electrodes Cz and FCz. All epochs were baseline-corrected by subtracting the mean of the 200 ms prestimulus period. Differences between the conditions of interest (successful STOP vs IGNORE, failed STOP vs IGNORE, and successful STOP vs failed STOP) were tested for significance using sample pointwise paired-samples *t* tests, in which the resulting *p* values were corrected for multiple comparisons using the false discovery rate correction procedure (Benjamini and Hochberg, 1995).

Results

Behavior. The RT results for the EEG experiment, which involved hand responses, paralleled those of Experiment 1, which was performed using the feet (with slightly faster RT values across the board; Fig. 2, right). There was again a strong effect of RT ($BF_{10} = 6.84 \times 10^{14}$, $F_{(1.73,32.80)} = 134.25$, $p < 0.001$, $\eta_p^2 = 0.88$). As with Experiment 1, RT was faster to the GO signal (mean = 480 ms, SD = 39.8) than to the IGNORE signal (mean 529 ms, SD = 41.7; $BF_{10} = 5.35 \times 10^5$, $t_{(19)} = 8.75$, $p_{Holm} < 0.001$, $d = 1.96$). RT was also faster to failed STOP trials (mean = 438, SD = 34.4 ms) than to IGNORE trials ($BF_{10} = 1.33 \times 10^9$, $t_{(19)} = 16.37$, $p_{Holm} < 0.001$, $d = 3.66$) and faster to failed STOP than to the GO trials ($BF_{10} = 6.57 \times 10^6$, $t_{(19)} = 7.62$, $p_{Holm} < 0.001$, $d = 1.70$). The mean *p*(inhibit) was 0.5 (SD = 0.016), the mean SSRT was 259 ms (SD = 28.3), and the mean SSD was 211 ms (SD = 56 ms). The difference in SSRT between the manual responses in Experiment 2 and the foot pedal responses in Experiment 1 was in the typical range for the SST across these effectors (Tabu et al., 2012). Miss rates were 0.43% in the GO and 1.14% in the IGNORE condition, respectively. Thus, the task was effective at producing approximately equal numbers of successful and failed stops, and the assumptions of the independent race model held (Verbruggen et al., 2019). According to the classification of Bissett and Logan (2014), 19 participants used a "Stop then discriminate" strategy according to their RT pattern, with the remaining one using "Independent discriminate then stop," and none using "Dependent discriminate then stop."

EMG validation. Single-trial correlation analysis revealed a strong positive correlation between GO trial EMG peak latency and GO RT (mean Pearson's $r = 0.85$, $BF_{10} = 1.30 \times 10^7$, $t_{(19)} = 11.21$, $p < 10^{-9}$, $d = 2.51$), validating the relationship between single-trial EMG and behavior.

EMG. Grand-averaged EMG traces showed that there was an early EMG peak (at ~140 ms) following both the STOP and IGNORE signals, which was followed by a sudden downturn (Fig. 4, top left). In line with our expectations, after this initial dip, the IGNORE EMG trace promptly began rebounding (at ~200 ms), whereas the successful STOP EMG trace became further suppressed.

Further in line with our hypothesis that both STOP and IGNORE trials would show a comparable early-latency suppression of EMG reflected in this downturn at ~140 ms (Hypothesis 1), successful STOP and IGNORE trial prEMG-int latencies showed a strong positive correlation (Pearson's $r = 0.73$, $BF_{10} = 141.26$, $p < 0.001$) and were not statistically different from one another, although evidence for this was inconclusive ($t_{(19)} = 1.88$, $p = 0.076$, $d = 0.42$, $BF_{10} = 1.01$). However, if anything, prEMG-int latency was numerically longer on successful STOP trials compared with IGNORE trials. Hence, nothing suggests that the prEMG reduction is unique to (or uniquely fast in) STOP compared with IGNORE trials.

In line with prior studies, successful STOP trial prEMG-int latency (mean latency = 139 ms, SD = 17.2) was significantly positively correlated with SSRT ($r = 0.68$, $BF_{10} = 40.79$, $p < 0.005$). Moreover, both IGNORE trial prEMG-int latency (mean Pearson's $r = 0.22$, $BF_{10} = 64.84$, $t_{(19)} = 4.17$, $p < 0.001$, $d = 0.93$) and its peak amplitude (mean Pearson's $r = -0.46$, $BF_{10} = 2.96 \times 10^6$, $t_{(19)} = -10.18$, $p < 0.001$, $d = 2.28$) were significantly correlated to IGNORE RT.

In summary, the EMG data from Experiment 2 show that both IGNORE signals and STOP signals are followed by low-latency reductions of the EMG. While this pattern on STOP trials is typically interpreted as a unique signature of outright action-stopping, our analyses indicate that this is a universal pattern that is common to all salient events. Successful STOP trials did not incur faster EMG suppression compared with IGNORE trials, and the latencies of the prEMG-int was highly correlated across subjects, suggesting that the same process is active on both types of trials.

MVPA. The results of the MVPA are presented in Figure 5, alongside the frontocentral ERPs. In line with our hypotheses and our findings from Experiment 1, IGNORE trials and successful STOP trials could not be successfully decoded from one another until 180 ms after signal onset (periods of above-chance decoding: 180-792 ms, 820-880 ms). IGNORE trials and failed STOP trials could not be decoded from one another until 428 ms following signal onset (significant period: 428-758 ms).

Notably, when the classifier was trained to distinguish between GO and other trials (failed STOP, successful STOP, and IGNORE), decoding performance was significantly above chance for almost the entirety of the epoch (Fig. 5). This suggests that

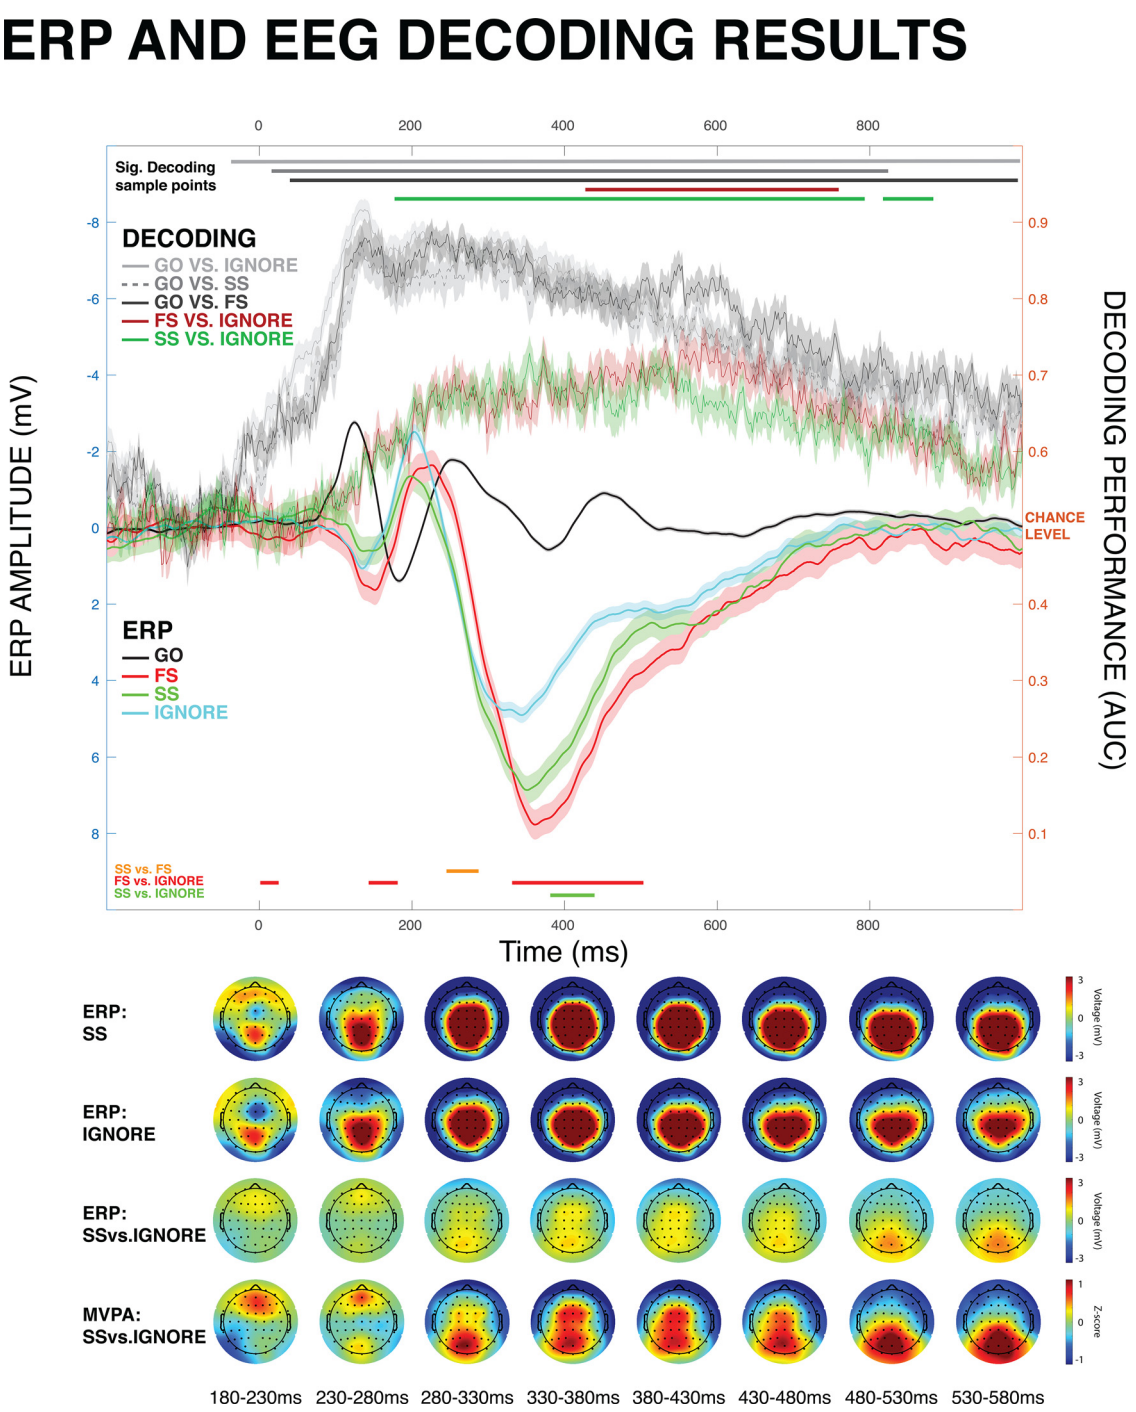

Figure 5. ERP and EEG decoding results. Top plot, Condition ERPs at frontocentral electrode sites FCz/Cz (left scale) and decoding performance for the contrasts of interest (right scale, 0.5 = chance-level decoding). Color bars above the figure represent time periods with significant above-chance whole-scalp MVPA decoding (color-coded identically to the legends). Color bars below the figure represent significant differences in the frontocentral ERP (orange represents SS vs FS; red represents FS vs IGNORE; green represents SS vs IGNORE). Bottom plot, Topographical distribution for SS, IGNORE, and SS-IGNORE ERP; and forward-estimated decoding estimates of SS versus IGNORE decoding.

our analysis methods were not too conservative to detect early differences in the contrasts of interest. One interesting anecdotal observation is that classifier performance peaked at ~140 ms for all combinations of STOP/IGNORE versus GO comparisons. This coincides with the time period during which all three trial types (successful/failed STOP, IGNORE) showed CSE suppression for STOP and IGNORE trials in Experiment 1.

ERP. Both successful and failed STOP trials showed significantly stronger positive voltage deflection compared with IGNORE trials during latencies corresponding to the P3 (successful: 388-436 ms; failed: 336-506 ms). No significant differences were found in the N2 time range (Fig. 5). Indeed, if anything, the N2 amplitude was larger for IGNORE trials compared with both STOP trial types. There were earlier ERP differences between failed STOP trials and IGNORE trials (2-22 ms and 148-182 ms), but not for successful STOP and IGNORE trials. During both periods, failed STOP trials showed increased positivity relative to IGNORE trials, perhaps indexing motor processes. Finally, the successful versus failed STOP ERP comparison showed significant differences

around SSRT (250-288 ms; SSRT = 259 ms) reflecting the fact that, in line with prior literature, successful STOP trials showed an earlier onset of the P3 compared with failed STOP trials (Kok et al., 2004; Wessel and Aron, 2015). These findings show that late-latency frontocentral ERPs (viz., the frontocentral P3) uniquely reflect processes related to outright stopping.

Discussion

In this study, we tackled two of the most prominent debates in the recent literature on inhibitory control and action-stopping: (1) Which neurophysiological processes reflect the attentional detection of an infrequent STOP signal and which index the actual implementation of inhibitory control? (2) Why do purported signatures of inhibition occur at two different latencies following STOP signals. Our results support the following view: Early signs of motor inhibition following STOP signals, including EMG suppression and CSE reduction, are not unique to action-stopping, and instead occur after any salient event. Processes unique to action-stopping only emerge after this initial inhibitory activity, signified by the emergence of the frontocentral stop signal P3.

Our results have several notable features. First, in scalp-EEG, our results show that the earliest activity that distinguishes STOP from IGNORE trials occurs at ~180 ms following STOP signals. However, the underlying ERP activity at the frontocentral scalp sites that provided the strongest decoding power indicates that this initial difference was because of an increased N2 ERP in the IGNORE condition (i.e., not to activity that is unique to stopping). (In principle, it is also possible that a separate frontocentral ERP in the same time range accounted for the difference in MVPA decoding, though ICA-based EEG studies have shown no evidence that frontocentral activity in that time range includes an additional ERP.) The first signature at which significant STOP-vs-IGNORE decoding was because of increased frontocentral STOP trial activity was the frontocentral P3. This supports the proposal that this ERP reflects a process that is unique to action-stopping (de Jong et al., 1990; Kok et al., 2004). Second, regarding CSE suppression, our results confirmed that such effects can be observed not just after STOP signals, but also after salient non-stop events (Wessel and Aron, 2013; Dutra et al., 2018; Iacullo et al., 2020). The current results substantially add to these findings by demonstrating that STOP and IGNORE signals do not differ in the degree to which they induce this early-latency CSE suppression (here, at 150 ms after signal). However, there was a clear divergence in this activity at 175 ms, when STOP trials showed additional suppression. This suggests that, on STOP trials, additional processing may sustain the initial CSE suppression that is shared by STOP and IGNORE signals. Moreover, while early CSE suppression did not differ between STOP and IGNORE trials, there was a significant difference between successful and failed STOP trials already at ~150 ms. This suggests that, although early CSE suppression is not unique to action-stopping, it does contribute to its success. It appears that successful STOP trials at least partially reflect the part of the STOP trial distribution in which the initial saliency-related CSE suppression was stochastically stronger. Hence, successful action-stopping likely results from a combination of early-latency inhibitory processes that are triggered by the saliency of the infrequent STOP signal and slower processes that are unique to STOP trials, reflected in later EEG signatures (Diesburg and Wessel, 2021). Finally, regarding EMG suppression on successful STOP trials, a signature of much recent interest (Raud and Huster, 2017; Jana et al., 2020; Raud et al., 2020), the current results are the first demonstration that this signature is not uniquely indicative of inhibitory control during action-stopping, but again common to all salient events, just like the CSE suppression. This is further in line with the assertion that the rapid invocation of inhibitory control is a stereotypic consequence of stimulus-driven attentional capture (Wessel and Aron, 2017).

More broadly, the current results are compatible with the view that action-stopping involves inhibitory control deployment in two stages. The first stage consists of an automatic invocation of inhibitory control that is common to all salient events. This is followed by a second stage, which is unique to action-stopping. As such, we propose that the neural cascade underlying human action-stopping may correspond to a recently proposed rodent model that postulates such a two-stage stopping process on the level of the basal ganglia (the pause-then-cancel model of Schmidt and Berke, 2017; Schmidt et al., 2013). This model proposes that action-stopping is achieved by a combination of two processes: an early-latency "pause" that consists of a broad, rapid inhibition of motor activity via the subthalamic nucleus, followed by a slower "cancel" process, which removes the ongoing striatal invigoration of motor activity. These two processes operate at different latencies, with the "pause" process buying time for the "cancel" process to shut off motor invigoration and ultimately enable successful action-stopping (see also Wiecki and Frank, 2013). Adapting this framework to the human domain, we propose that the early signatures that are common to STOP and IGNORE signals (early CSE/EMG suppression and the corresponding early-latency EEG activity) reflect the "pause" stage. We further propose that later signatures (e.g., the frontocentral P3) reflect the "cancel" stage.

Such an adaptation of such a two-stage pause-then-cancel model is in line with a large body of both recent and classic work in humans (Diesburg and Wessel, 2021). Specifically, we hypothesize that the initial "pause" phase is implemented via the hyperdirect pathway from inferior frontal cortex to subthalamic nucleus, which has recently been empirically described in humans (Chen et al., 2020) and has been proposed to be crucial to action-stopping (Aron et al., 2007). However, we propose that this "pause"-related hyperdirect pathway activity is not specific to action-stopping, but follows any type of salient event. This perspective unites the proposal that right inferior frontal cortex is a stimulus-driven attentional "circuit-breaker" (Corbetta and Shulman, 2002) with the proposal that it is an inhibitory node (Aron et al., 2007, 2014), and explains why it is active even when salient events do not require action-stopping (Sharp et al., 2010; Chatham et al., 2012). Furthermore, we hypothesize that, in situations that explicitly require outright action-stopping, a second "cancel" phase is implemented by structures underlying the frontocentral P3 (e.g., the pre-SMA) (Enriquez-Geppert et al., 2010; Swann et al., 2011). More broadly, we propose that this second phase entails a specific retuning of active motor plans, geared toward the exact behavioral requirements posed by the task. In the SST, this would entail an arrest of any still-active invigoration of the motor program ("cancel") (Schmidt and Berke, 2017). However, for salient signals that do not require action-stopping, this second phase could consist of other motor adaptations. This would explain why the frontocentral P3 after salient events relates to elongations in RT or subtle changes in motor force following salient events (Novembre et al., 2018, 2019; Wessel and Huber, 2019; Waller et al., 2019). As such, we propose a broader adaptation of the SST-specific "pause-then-cancel" model of Schmidt

and Berke (2017): a pause-then-*retune* model, where "retune" means "cancel" in the specific case of the SST, but can mean more subtle adjustments for other salient events. Moreover, although we operationalized saliency as infrequency in the present work, aspects, such as stimulus properties (e.g., visual contrast, auditory loudness) or behavioral relevance, may also make events salient (in the broader sense of Corbetta and Shulman, 2002), and we would expect such events to also prompt the Pause phase. While this model is still largely speculative, it does provide a framework to explain the current set of results, unite them with classic findings, and provide a translation between rodent and human work.

One potential point of criticism of the current study could be that since IGNORE signals were presented within the context of an SST, participants may have been biased toward using inhibitory control after IGNORE signals. Indeed, behavioral work using the current task indicates that participants may use a so-called "stop-then-discriminate" strategy when facing salient stimuli while they anticipate potential stop signals (Bissett and Logan, 2014). Notably, the vast majority of our participants used this "stop-then-discriminate" strategy. In line with our neuroscientific theory outlined above, we propose that the concept of the "stop-then-discriminate" strategy is a way of framing the "pause" in the pause-then-cancel model (or the related "Hold-your-horses" concept of Frank; compare Frank, 2006; Frank et al., 2007) in behavioral terms. According to our view (for a more detailed theoretical account, compare Diesburg and Wessel, 2021), any salient event obligatorily invokes the "pause" phase, whereas only stop signals that were successfully identified as such engage the subsequent "cancel" phase. This maps onto the "stop-then-discriminate" strategy described by Bissett and Logan (2014). In contrast, the Bissett and Logan (2014) "independent discriminate then stop" strategy may describe subjects that, either strategically or presdispositionally, minimize the influence of the "pause" phase, relying primarily on the "cancel" process to achieve stopping (which should, in principle, be less effective in effecting a stop, but in turn reduce the effect of IGNORE signals on GO trial RT). However, only one subject across both experiments used this strategy, which is hence unlikely to be the default. Regardless, to address the question whether proactive control confounded our results, we note that previous work has already shown that even outside of stop signal contexts, salient signals produce nonselective CSE suppression (Iacullo et al., 2020), frontocentral EEG activity (Courchesne et al., 1975; Wessel and Huber, 2019), slowing of motor behavior (Waller et al., 2019), changes in isometric motor activity (Novembre et al., 2018, 2019), and activity in right inferior frontal cortex, as mentioned (Sharp et al., 2010; Chatham et al., 2012). Hence, the observed IGNORE trial changes in CSE, EMG, and EEG in the current study are not contingent on proactive control. Instead, however, the current study shows that when STOP and IGNORE signals occur under conditions of equal proactive control, both produce the same degree of early-latency inhibition. Moreover, if our participants would have first performed an IGNORE-but-not-STOP task and then a separate STOP-but-not-IGNORE task, this would have invalidated the MVPA, precisely because if only STOP signals were presented in the presence of proactive control, this would have induced systematic changes to the EEG activity both before and after STOP signals (e.g., Kenemans, 2015; Elchlepp et al., 2016; Soh et al., 2021). Hence, we believe that the current approach presents an ideal comparison between IGNORE and STOP trials under conditions of equal proactive control.

In conclusion, we aimed to find unique neurophysiological signatures of action-stopping by comparing STOP and IGNORE trials across methodologies. We found no unique signatures of action-stopping before the emergence of the stop signal P3. Particularly, IGNORE signals produced both early EMG and CSE suppression on par with actual STOP signals, suggesting that these early-latency signatures do not uniquely signify action-stopping. As such, we believe that these results warrant a reinterpretation of past work in the domain of inhibitory control, perhaps along the lines of recently proposed two-stage models of action-stopping, with the prediction that only the second stage of such models is actually unique to outright action-stopping.

References

Aron AR, Robbins TW, Poldrack RA (2014) Inhibition and the right inferior frontal cortex: one decade on. Trends Cogn Sci 18:177–185.

Aron AR, Durston S, Eagle DM, Logan GD, Stinear CM, Stuphorn V (2007) Converging evidence for a fronto-basal-ganglia network for inhibitory control of action and cognition. J Neurosci 27:11860–11864.

Badry R, Mima T, Aso T, Nakatsuka M, Abe M, Fathi D, Foly N, Nagiub H, Nagamine T, Fukuyama H (2009) Suppression of human cortico-motoneuronal excitability during the Stop signal task. Clin Neurophysiol 120:1717–1723.

Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc B 57:289–300.

Bissett PG, Logan GD (2014) Selective stopping? Maybe not. J Exp Psychol Gen 143:455–472.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433–436.

Chatham CH, Claus ED, Kim A, Curran T, Banich MT, Munakata Y (2012) Cognitive control reflects context monitoring, not motoric stopping, in response inhibition. PLoS One 7:e31546.

Chen W, de Hemptinne C, Miller AM, Leibbrand M, Little SJ, Lim DA, Larson PS, Starr PA (2020) Prefrontal-subthalamic hyperdirect pathway modulates movement inhibition in humans. Neuron 106:579–588.e573.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci 3:201–215.

Courchesne E, Hillyard SA, Galambos R (1975) Stimulus novelty, task relevance and the visual evoked potential in man. Electroencephalogr Clin Neurophysiol 39:131–143.

de Jong R, Coles MG, Logan GD, Gratton G (1990) In search of the point of no return: the control of response processes. J Exp Psychol Hum Percept Perform 16:164–182.

Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9–21.

Diesburg DA, Wessel JR (2021) The Pause-then-Cancel model of human action-stopping: theoretical considerations and empirical evidence. Neurosci Biobehav Rev 129:17–34.

Dutra IC, Waller DA, Wessel JR (2018) Perceptual surprise improves action stopping by nonselectively suppressing motor activity via a neural mechanism for motor inhibition. J Neurosci 38:1482–1492.

Elchlepp H, Lavric A, Chambers CD, Verbruggen F (2016) Proactive inhibitory control: a general biasing account. Cogn Psychol 86:27–61.

Enriquez-Geppert S, Konrad C, Pantev C, Huster RJ (2010) Conflict and inhibition differentially affect the N200/P300 complex in a combined go/nogo and stop signal task. Neuroimage 51:877–887.

Fahrenfort JJ, van Driel J, van Gaal S, Olivers CN (2018) From ERPs to MVPA using the Amsterdam Decoding and Modeling Toolbox (ADAM). Front Neurosci 12:368.

Faul F, Erdfelder E, Buchner A, Lang AG (2009) Statistical power analyses using G*Power 3.1: tests for correlation and regression analyses. Behav Res Methods 41:1149–1160.

Frank MJ (2006) Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. Neural Netw 19:1120–1136.

Frank MJ, Samanta J, Moustafa AA, Sherman SJ (2007) Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism. Science 318:1309–1312.

Huster RJ, Messel MS, Thunberg C, Raud L (2020) The P300 as marker of inhibitory control: fact or fiction? Cortex 132:334–348.

Hynd M, Soh C, Rangel BO, Wessel JR (2021) Paired-pulse TMS and scalp EEG reveal systematic relationship between inhibitory GABAa signaling

in M1 and frontocentral cortical activity during action stopping. J Neurophysiol 125:648–660.
Iacullo C, Diesburg DA, Wessel JR (2020) Nonselective inhibition of the motor system following unexpected and expected infrequent events. Exp Brain Res 238:2701–2710.
Jana S, Hannah R, Muralidharan V, Aron AR (2020) Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. eLife 9:e50371.
Kenemans JL (2015) Specific proactive and generic reactive inhibition. Neurosci Biobehav Rev 56:115–126.
Kok A, Ramautar JR, De Ruiter MB, Band GP, Ridderinkhof KR (2004) ERP components associated with successful and unsuccessful stopping in a stop signal task. Psychophysiology 41:9–20.
Logan GD, Cowan WB, Davis KA (1984) On the ability to inhibit simple and choice reaction time responses: a model and a method. J Exp Psychol Hum Percept Perform 10:276–291.
Love J, Selker R, Marsman M, Jamil T, Dropmann D, Verhagen J, Ly A, Gronau QF, Šmíra M, Epskamp S, Matzke D, Wild A, Knight P, Rouder JN, Morey RD, Wagenmakers EJ (2019) JASP: graphical Statistical Software for Common Statistical Designs. J Stat Soft 88.
Majid DS, Cai W, Corey-Bloom J, Aron AR (2013) Proactive selective response suppression is implemented via the basal ganglia. J Neurosci 33:13259–13269.
Makeig S, Bell AJ, Jung TP, Sejnowski TJ (1996) Independent component analysis of electroencephalographic data. Adv Neural Information Process Syst 145–151.
Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods 164:177–190.
Matzke D, Love J, Wiecki T, Brown S, Logan G, Wagenmakers EJ (2013) Release the BEESTS: Bayesian estimation of ex-Gaussian STop signal reaction time distributions. Front Psychol 4:918.
Novembre G, Pawar VM, Bufacchi RJ, Kilintari M, Srinivasan M, Rothwell JC, Haggard P, Iannetti GD (2018) Saliency detection as a reactive process: unexpected sensory events evoke corticomuscular coupling. J Neurosci 38:2385–2397.
Novembre G, Pawar VM, Kilintari M, Bufacchi RJ, Guo Y, Rothwell JC, Iannetti GD (2019) The effect of salient stimuli on neural oscillations, isometric force, and their coupling. Neuroimage 198:221–230.
Raud L, Huster RJ (2017) The temporal dynamics of response inhibition and their modulation by cognitive control. Brain Topogr 30:486–501.
Raud L, Westerhausen R, Dooley N, Huster RJ (2020) Differences in unity: the go/no-go and stop signal tasks rely on different mechanisms. Neuroimage 210:116582.
Rossi S, Hallett M, Rossini PM, Pascual-Leone A (2011) Screening questionnaire before TMS: an update. Clin Neurophysiol 122:1686.
Rossini PM, Barker AT, Berardelli A, Caramia MD, Caruso G, Cracco RQ, Dimitrijević MR, Hallett M, Katayama Y, Lücking CH, Maertens de Noordhout AL, Marsden CD, Murray NM, Rothwell JC, Swash M, Tomberg C (1994) Non-invasive electrical and magnetic stimulation of the brain, spinal cord and roots: basic principles and procedures for routine clinical application. Report of an IFCN committee. Electroencephalogr Clin Neurophysiol 91:79–92.
Sánchez-Carmona AJ, Santaniello G, Capilla A, Hinojosa JA, Albert J (2019) Oscillatory brain mechanisms supporting response cancellation in selective stopping strategies. Neuroimage 197:295–305.
Schmidt R, Berke JD (2017) A Pause-then-Cancel model of stopping: evidence from basal ganglia neurophysiology. Philos Trans R Soc Lond B Biol Sci 372:20160202.
Schmidt R, Leventhal DK, Mallet N, Chen F, Berke JD (2013) Canceling actions involves a race between basal ganglia pathways. Nat Neurosci 16:1118–1124.
Sebastian A, Jung P, Neuhoff J, Wibral M, Fox PT, Lieb K, Fries P, Eickhoff SB, Tüscher O, Mobascher A (2016) Dissociable attentional and inhibitory networks of dorsal and ventral areas of the right inferior frontal cortex: a combined task-specific and coordinate-based meta-analytic fMRI study. Brain Struct Funct 221:1635–1651.
Sebastian A, Rössler K, Wibral M, Mobascher A, Lieb K, Jung P, Tüscher O (2017) Neural architecture of selective stopping strategies: distinct brain activity patterns are associated with attentional capture but not with outright stopping. J Neurosci 37:9785–9794.
Sharp DJ, Bonnelle V, De Boissezon X, Beckmann CF, James SG, Patel MC, Mehta MA (2010) Distinct frontal systems for response inhibition, attentional capture, and error processing. Proc Natl Acad Sci USA 107:6106–6111.
Silver NC, Dunlap WP (1987) Averaging correlation coefficients: should Fisher's *z* transformation be used? J Appl Psychol 72:146–148.
Skippen P, Fulham WR, Michie PT, Matzke D, Heathcote A, Karayanidis F (2020) Reconsidering electrophysiological markers of response inhibition in light of trigger failures in the stop signal task. Psychophysiology 57: e13619.
Soh C, Hynd M, Rangel BO, Wessel JR (2021) Adjustments to proactive motor inhibition without effector-specific foreknowledge are reflected in a bilateral upregulation of sensorimotor β-burst rates. J Cogn Neurosci : 33:784–716.
Swann N, Poizner H, Houser M, Gould S, Greenhouse I, Cai W, Strunk J, George J, Aron AR (2011) Deep brain stimulation of the subthalamic nucleus alters the cortical profile of response inhibition in the beta frequency band: a scalp EEG study in Parkinson's disease. J Neurosci 31:5721–5729.
Tabu H, Mima T, Aso T, Takahashi R, Fukuyama H (2012) Common inhibitory prefrontal activation during inhibition of hand and foot responses. Neuroimage 59:3373–3378.
van de Laar MC, van den Wildenberg WP, van Boxtel GJ, van der Molen MW (2014) Development of response activation and inhibition in a selective stop signal task. Biol Psychol 102:54–67.
Verbruggen F, Aron AR, Band GP, Beste C, Bissett PG, Brockett AT, Brown JW, Chamberlain SR, Chambers CD, Colonius H, Colzato LS, Corneil BD, Coxon JP, Dupuis A, Eagle DM, Garavan H, Greenhouse I, Heathcote A, Huster RJ, Jahfari S, et al. (2019) A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop signal task. eLife 8:e46323.
Waller DA, Hazeltine E, Wessel JR (2019) Common neural processes during action-stopping and infrequent stimulus detection: the frontocentral P3 as an index of generic motor inhibition. Int J Psychophysiol 163:11–21.
Wessel JR, Aron AR (2013) Unexpected events induce motor slowing via a brain mechanism for action-stopping with global suppressive effects. J Neurosci 33:18481–18491.
Wessel JR, Aron AR (2015) It's not too late: the onset of the frontocentral P3 indexes successful response inhibition in the stop signal paradigm. Psychophysiology 52:472–480.
Wessel JR, Aron AR (2017) On the globality of motor suppression: unexpected events and their influence on behavior and cognition. Neuron 93:259–280.
Wessel JR, Huber DE (2019) Frontal cortex tracks surprise separately for different sensory modalities but engages a common inhibitory control mechanism. PLoS Comput Biol 15:e1006927.
Wessel JR, Reynoso HS, Aron AR (2013) Saccade suppression exerts global effects on the motor system. J Neurophysiol 110:883–890.
Wickens TD (2002) Elementary signal detection theory. Oxford: Oxford UP.
Wiecki TV, Frank MJ (2013) A computational model of inhibitory control in frontal cortex and basal ganglia. Psychol Rev 120:329–355.