NeuroImage 119 (2015) 286–295

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Anticipating conflict: Neural correlates of a Bayesian belief and its motor consequence

Sien Hu [a,*], Jaime S. Ide [b], Sheng Zhang [a], Chiang-shan R. Li [a,c,d,**]

[a] Department of Psychiatry, Yale University, New Haven, CT 06519, USA
[b] Department of Biomedical Engineering, Stony Brook University, Stony Brook, NY 11794, USA
[c] Department of Neurobiology, Yale University, New Haven, CT 06520, USA
[d] Interdepartmental Neuroscience Program, Yale University, New Haven, CT 06520, USA

ARTICLE INFO

Article history:
Received 10 April 2015
Accepted 10 June 2015
Available online 18 June 2015

Keywords:
Top–down
Imaging
Stop signal task
Inhibitory control
Bayesian

ABSTRACT

Previous studies have examined the neural correlates of proactive control using a variety of behavioral paradigms; however, the neural network relating the control process to its behavioral consequence remains unclear. Here, we applied a dynamic Bayesian model to a large fMRI data set of the stop signal task to address this issue. By estimating the probability of the stop signal – p(Stop) – trial by trial, we showed that higher p(Stop) is associated with prolonged go trial reaction time (RT), indicating proactive control of motor response. In modeling fMRI signals at trial and target onsets, we distinguished activities of proactive control, prediction error, and RT slowing. We showed that the anterior pre-supplementary motor area (pre-SMA) responds specifically to increased stop signal likelihood, and its activity is correlated with activations of the posterior pre-SMA and bilateral anterior insula during prolonged response times. This directional link is also supported by Granger causality analysis. Furthermore, proactive control, prediction error, and time-on-task are each mapped to distinct areas in the medial prefrontal cortex. Together, these findings dissect regional functions of the medial prefrontal cortex in cognitive control and provide system level evidence associating conflict anticipation with its motor consequence.

Introduction

The ability to proactively adjust our behavior is integral to survival. Studying the neural bases of proactive control advances our understanding of how decisions are made in a changing environment and why individuals are engaged in impulsive behavior.

Proactive control has been studied in the laboratory with a variety of behavioral paradigms (Brass and Haggard, 2007; Horga et al., 2011; Kuhn et al., 2009). Frontal and parietal cortices respond to cued attention allocation (Luks et al., 2007) and preparatory control of a switch in response (Rushworth et al., 2001). When participants withheld movements while waiting to detect a target, activation of the superior medial prefrontal cortex (MPFC) and inferior parietal lobule supports proactive control (Jaffard et al., 2008). The importance of proactive control is demonstrated in a computational model of saccadic eye movement (Lo et al., 2009) and may be generalized to other systems (Ballanger, 2009).

In the stop signal task (SST), increased stop signal probability bolsters proactive control, evidenced by delayed activity in the primary motor cortex (Jahfari et al., 2010). Varying the occurrence of go trials prior to a stop trial, Vink et al. (2005) showed MPFC, caudate and left insula increasing activation to stop likelihood. Chikazoe et al. (2009) used a SST with two types of go trials, 'go-certain' and 'go-uncertain'. Motor responses may require interruption in 'go-uncertain' trials; thus elicited activations are thought to reflect proactive inhibitory control. A recent study of choice SST showed activation of the superior medial frontal and inferior frontal cortices when participants are informed, compared to uninformed, as to which effector to use (Smittenaar et al., 2013). Furthermore, in reaction time tasks proactive control is frequently followed by prolonged response times, and studies have also described medial prefrontal activities in association with time-on-task (Carp et al., 2010; Grinband et al., 2011). Together, these studies highlighted an important role of the MPFC in proactive control but it remains unclear whether distinct regions of the MPFC mediate conflict anticipation and behavioral outcome or whether these activities are related. Another important issue concerns the confound of stimulus prediction error, which is known to drive MPFC activation (Glascher et al., 2010; Ide et al., 2013; Nee et al., 2011; So and Stuphorn, 2012). As pointed out earlier, because the expectation of the stop signal is not realized during go signal onset, a violation of this expectation or prediction error

* Correspondence to: S. Hu, Connecticut Mental Health Center, S108, 34 Park Street, New Haven, CT 06519, USA.
** Correspondence to: C.S. Li, Connecticut Mental Health Center, S112, 34 Park Street, New Haven, CT 06519, USA.
E-mail addresses: sien.hu@yale.edu (S. Hu), chiang-shan.li@yale.edu (C.S. Li).

http://dx.doi.org/10.1016/j.neuroimage.2015.06.032
1053-8119/

occurs at the same time, presenting a confound to proactive control (Zandbelt et al., 2013).

The current study aimed to address these issues. We used a Bayesian model to compute the likelihood of stop signal – p(Stop) – trial by trial in the SST and established a correlation between p(Stop) and reaction time (RT) – a sequential effect – for individual subjects. We modeled the fMRI signals at trial onset to characterize activations to p(Stop), and at go signal onset to characterize activations to prediction error and RT slowing. With exclusive masking we identified neural responses specific to each component process of proactive control. We showed that stop signal anticipation, stimulus prediction error, and RT slowing ("time-on-task") are mapped to distinct areas in the MPFC. Importantly, neural activities specific to stop signal anticipation are both correlated and Granger causally related to activities specific to RT slowing, supporting a directional link between these two processes.

Methods

Participants and behavioral task

One hundred fourteen healthy adults (64 females; 30.7 ± 11.0 years of age) participated in this study. All participants signed a written consent after they were given a detailed explanation of the study in accordance with a protocol approved by the Yale Human Investigation Committee. We were able to include a large number of participants in the current study by combining data sets both from studies exclusively of healthy individuals (Hendrick et al., 2010; Hu et al., 2012; Ide and Li, 2011; Zhang and Li, 2012; Zhang et al., 2012) and from the healthy cohort of clinical studies (Bednarski et al., 2012; Hendrick et al., 2012; Li et al., 2009b; Yan and Li, 2009), all conducted under the same behavioral task and in the identical 3 T scanner.

Participants performed a stop signal task or SST (Hu and Li, 2012; Li et al., 2009a), in which go and stop trials were randomly intermixed in presentation with an inter-trial-interval of 2 seconds (s). A fixation dot appeared on screen to signal the beginning of each trial. After a fore-period varying from 1 s to 5 s (uniform distribution), the dot became a circle – the "go" signal – prompting participants to quickly press a button. The circle disappeared at button press or after 1 s if the participant failed to respond. In approximately one quarter of trials, the circle was followed by a 'cross' – the stop signal – prompting participants to withhold button press. The trial terminated at button press or after 1 s if the participant successfully inhibited the response. The time between the go and stop signals, the stop signal delay (SSD), started at 200 ms and varied from one stop trial to the next according to a staircase procedure, increasing and decreasing by 67 ms each after a successful and failed stop trial (Levitt, 1971). The 67 ms step reflects four screen frames (monitor refreshing rate = 60 Hz) during stimulus presentation, as used in almost all of our previous studies (Chao et al., 2012; Farr et al., 2014; Ide et al., 2013; Liao et al., 2014; Luo et al., 2013; Winkler et al., 2013; Zhang et al., 2014, in press). In an earlier work we have tried a smaller staircase (32 ms) in a pilot study, which yielded similar performance profiles (Li et al., 2005). With the staircase procedure we anticipated that participants would succeed in withholding the response half of the time. Participants were trained briefly on the task before imaging to ensure that they understood the task. They were instructed to quickly press the button when they saw the go signal while keeping in mind that a stop signal might come up in some trials. In the scanner, they completed four 10-minute sessions of the task, with approximately 100 trials in each session.

Behavioral data analysis

A critical SSD was computed for each participant that represented the time delay required for the participant to successfully withhold the response in half of the stop trials, following a maximum likelihood procedure (Wetheril et al., 1966). Briefly, SSDs across trials were grouped into runs, with each run being defined as a monotonically increasing or decreasing series. We derived a mid-run estimate by taking the middle SSD (or average of the two middle SSDs when there was an even number of SSDs) of every second run. The critical SSD was computed by taking the mean of all mid-run SSDs. It was reported that, except for experiments with a small number of trials (<30), the mid-run measure was close to the maximum likelihood estimate of X50 (50% positive response, Wetheril et al., 1966). The stop signal reaction time (SSRT) was computed for each participant by subtracting the critical SSD from the median go trial reaction time (Logan et al., 1984).

The SSRT can also be computed from a critical SSD estimated from an "inhibitory function", by fitting the response rates at different SSDs to a sigmoid function. In order to have a robust estimate of correct response rates, the trial numbers at different SSDs should be the same or similar, which applies to experimental designs where the SSD's are blocked. In the current work, we used a staircase procedure with varying trial numbers across SSDs. In particular, the low number of trials at the lower and higher ends of SSD oftentimes resulted in a less than ideal fit. Therefore, we used a maximum likelihood procedure to estimate the critical SSD and SSRT, following many of our published studies and the literature.

A sequential effect was quantified by Pearson correlation between p(Stop) – the Bayesian estimate of the probability of a stop signal (see below) – and RT on go trials for individual subjects.

Trial by trial Bayesian estimate of the likelihood of a stop signal

We used a dynamic Bayesian model (Yu and Cohen, 2009) to estimate the prior belief of an impending stop signal on each trial, based on prior stimulus history. The model assumes that subjects believe that stop signal frequency r_k on trial k has probability α of being the same as r_{k-1}, and probability $(1-\alpha)$ of being re-sampled from a prior distribution $\pi(r_k)$. Subjects are also assumed to believe that trial k has probability r_k of being a stop trial, and probability $1-r_k$ of being a go trial. Based on these generative assumptions, subjects are assumed to use Bayesian inference to update their prior belief of seeing a stop signal on trial k, $p(r_k|s_{k-1})$ based on the prior on the last trial $p(r_{k-1}|s_{k-1})$ and last trial's true category ($s_k = 1$ for stop trial, $s_k = 0$ for go trial), where $s_k = \{s_1,\ldots, s_k\}$ is short-hand for all trials 1 through k. Specifically, given that the posterior distribution was $p(r_{k-1}|s_{k-1})$ on trial $k-1$, the prior distribution of stop signal in trial k is given by:

$$p(r_k|s_{k-1}) = \alpha p(r_{k-1}|s_{k-1}) + (1-\alpha)\pi(r_k),$$

where the prior distribution $\pi(r_k)$ is assumed to be a beta distribution with prior mean *pm*, and shape parameter *scale*, and the posterior distribution is computed from the prior distribution and the outcome according to Bayes' rule:

$$p(r_k|s_k) \propto P(s_k|r_k)p(r_k|s_{k-1}).$$

The Bayesian estimate of the probability of trial k being a stop trial, which we colloquially call p(Stop) in this paper, given the predictive distribution $p(r_k|s_{k-1})$ is expressed by:

$$\begin{aligned} P(s_k = 1|s_{k-1}) &= \int P(s_k = 1|r_k)P(r_k|s_{k-1})dr_k \\ &= \int r_k P(r_k|s_{k-1})dr_k = \langle r_k|s_{k-1}\rangle. \end{aligned}$$

In other words, the probability p(Stop) of a trial k being a stop trial is simply the mean of the predictive distribution $p(r_k|s_{k-1})$. The assumption that the predictive distribution is a mixture of the previous posterior distributions and a generic prior distribution is essentially equivalent to using a causal, exponential, linear filter to estimate the current rate of stop trials (Yu et al., 2009). In summary, for each subject, given a sequence of observed go/stop trials, and the three model parameters {α, *pm*, *scale*}, we estimated p(Stop) for each trial.