Visual Information Shapes the Dynamics of Corticobasal Ganglia Pathways during Response Selection and Inhibition

Sara Jahfari[1], Lourens Waldorp[2], K. Richard Ridderinkhof[2], and H. Steven Scholte[2]

Abstract

■ Action selection often requires the transformation of visual information into motor plans. Preventing premature responses may entail the suppression of visual input and/or of prepared muscle activity. This study examined how the quality of visual information affects frontobasal ganglia (BG) routes associated with response selection and inhibition. Human fMRI data were collected from a stop task with visually degraded or intact face stimuli. During go trials, degraded spatial frequency information reduced the speed of information accumulation and response cautiousness. Effective connectivity analysis of the fMRI data showed action selection to emerge through the classic direct and indirect BG pathways, with inputs deriving form both prefrontal and visual regions. When stimuli were degraded, visual and prefrontal regions processing the stimulus information increased connectivity strengths toward BG, whereas regions evaluating visual scene content or response strategies reduced connectivity toward BG. Response inhibition during stop trials recruited the indirect and hyperdirect BG pathways, with input from visual and prefrontal regions. Importantly, when stimuli were nondegraded and processed fast, the optimal stop model contained additional connections from prefrontal to visual cortex. Individual differences analysis revealed that stronger prefrontal-to-visual connectivity covaried with faster inhibition times. Therefore, prefrontal-to-visual cortex connections appear to suppress the fast flow of visual input for the go task, such that the inhibition process can finish before the selection process. These results indicate response selection and inhibition within the BG to emerge through the interplay of top–down adjustments from prefrontal and bottom–up input from sensory cortex. ■

INTRODUCTION

Many of our day-to-day decisions rely on the processing and incorporation of visual information into action plans (e.g., "What signs should I follow home?" or "Should fog affect my driving speed or cautiousness?"). Classic neurophysiological theories describe the goal-oriented pFC to direct response selection or suppression in concert with the response-gating basal ganglia (BG; e.g., Mink, 1996). However, despite the importance of pFC–BG interactions in capturing most cognitive and motor domains of action control, very little is known about how these frameworks integrate interactions with sensory systems to influence the selection or cancelation of planned actions.

Mechanistically, response selection implicates the striatum in receiving modality-independent stimulus evaluations from the dorsolateral pFC (DLPFC; Philiastides, Auksztulewicz, Heekeren, & Blankenburg, 2011), whereas the pre-SMA monitors and evaluates response strategies (Shenhav, Botvinick, & Cohen, 2013; Alexander & Brown, 2011; Hikosaka & Isoda, 2010). The striatum then weighs the incoming information to facilitate response selection through a direct pallidal-thalamico pathway, whereas an indirect pathway via the subthalamic nucleus (STN) allows integration of additional information by slowing or stopping the motor output (Schmidt, Leventhal, Mallet, Chen, & Berke, 2013; Cavanagh et al., 2011; Nambu, 2009; Frank, 2006). When a planned response needs to be withdrawn, the right inferior frontal gyrus (rIFG) and pre-SMA are thought to influence the indirect and hyperdirect pathways (Herz et al., 2014; Jahfari et al., 2011, 2012; Aron & Poldrack, 2006). In contrast to the direct or indirect pathways (with pFC to striatum projections), the hyperdirect pathway involves pFC to STN projections and is thought to function as a fast "brake" on all responses. This study aims to specify how sensory, pFC, and BG regions collaborate to enable action selection or suppression.

One popular framework to study the neural indices of both response selection and inhibition has been the stop signal task, where the go process is initiated by a choice stimulus (often an arrow) and incidental stop signals designate that a planned response has to be withdrawn. Recently, we showed how variance in the quality of visual information (relevant for the go task) differentially,

[1]VU University Amsterdam, [2]University of Amsterdam

Journal of Cognitive Neuroscience 27:7, pp. 1344–1359
doi:10.1162/jocn_a_00792

but robustly, influences both go and stop processes in the stop signal task (Jahfari, Ridderinkhof, & Scholte, 2013). Specifically, when the perceptual quality of the choice stimulus was reduced (i.e., choice discrimination was difficult), go trial RTs and error rates increased, whereas the speed of information accumulation and the level of cautiousness were reduced. In contrast, longer stop signal RTs (SSRTs) were observed when response selection was progressing fast and accurately. The importance of perceptual processes during response selection and inhibition is further underlined through recent experiments (Chiu & Egner, 2015; Verbruggen, Stevens, & Chambers, 2014), computational modeling work (Salinas & Stanford, 2013; Boucher, Palmeri, Logan, & Schall, 2007), and theoretical frameworks (Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014). However, despite the importance of incoming sensory information during action control, it remains largely unknown how perceptual regions that first identify and process the perceptual input interact with classic pFC–BG frameworks of action control. Connectivity-based neuroimaging analysis may present one approach to the systematic evaluation of the functioning and architecture of these pFC–BG networks, with or without the incorporation of visual regions that first process the incoming information for the task.

In the current study, fMRI data were recorded from a stop signal task where participants discriminated gender in faces containing low, high, or all spatial frequencies (SFs; Fiorentini, Maffei, & Sandini, 1983; Figure 1A, B). Occasionally, face stimuli were followed by stop signals instructing participants to withhold planned responses. From a general perspective, we can distinguish, within the visual system, a ventral and dorsal pathway. The dorsal pathway responds faster and has a preference for more lower SFs, and the ventral system responds slower and has a preference for higher SFs (Goodale & Milner, 1992). In addition to this preference, the choice for the removal of SF, in contrast to other degrading methods, was based on a recent behavioral study where we found robust and distinct effects of SF on both action selection and inhibition (Jahfari et al., 2013).

A hierarchical Bayesian version of the drift diffusion model (DDM; Figure 1C) was used to model the RT data of correct and incorrect go responses (Wiecki, Sofer, & Frank, 2013; Ratcliff & McKoon, 2008). This approach allowed us to establish how the degradation of visual information affects the ease of evidence accumulation ("drift") and/or strategic adjustments in the amount of evidence required before a decision threshold is reached ("boundary"). The slower the progression of evidence accumulation, the greater the cost of acquiring additional information from visual regions. We therefore expected the selective removal of SFs to affect both drift rate and decision thresholds (Jahfari et al., 2013; Drugowitsch, Moreno-Bote, Churchland, Shadlen, & Pouget, 2012).

Effective connectivity analysis was formalized by modeling fMRI data with a recently developed method called

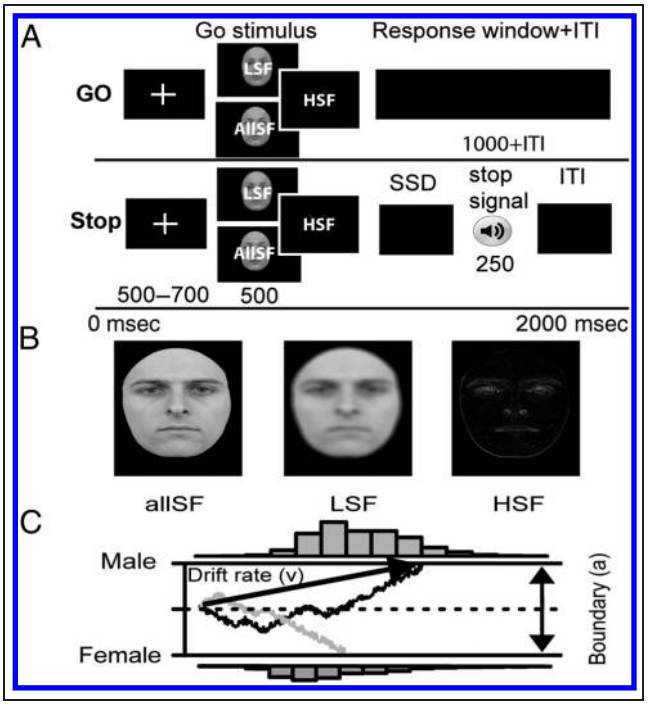

Figure 1. Outline of the task, stimuli, and DDM. (A) Each trial started with the presentation of a fixation cross, followed by a male or female face, indicating a left or right response. During stop trials, a tone was played at some SSD after the presentation of the go stimulus. The tone instructed participants to suppress the indicated response (ITI = intertrial interval, allSF = full SF spectrum, LSF = low SF, HSF = high SF). (B) To study the dynamics of information relay, the presented faces contained low, high, or all SFs. Note that the prints underestimate the contrasts used in the actual experiment, especially for HSF faces. (C) Simplistic outline of the DDM. The amount of information required "boundary," the ease of information accumulation "drift rate," and nondecision time were varied to find the best model for representing the observed RT distributions (lowest deviance information criterion).

ancestral graphs (Waldorp, Christoffels, & van de Ven, 2011). This method combines model comparison techniques (Burnham & Anderson, 2004) with effective connectivity analysis to identify which model best supports the observed pattern of activity across trials (Jahfari et al., 2011, 2012). The selection of visual regions that might interact with pFC or BG was based on a recent review, reporting dense connections to link the visual ventral pathway with the striatum while the dorsal pathway projects into DLPFC (Kravitz, Saleem, Baker, Ungerleider, & Mishkin, 2013). Therefore, potentially important sensory regions include regions that respond to the processing of faces (fusiform face area [FFA]; Kanwisher, McDermott, & Chun, 1997) or represent saliency or scene content (lateral occipital cortex [LOC]; Altmann, Deubelius, & Kourtzi, 2004) in the ventral stream or integrate sensorimotor information (parietal operculum [PO]; Tosoni, Galati, Romani, & Corbetta, 2008) in the dorsal stream.

We hypothesized that perceptual information, during response selection, might be conveyed either through

the ventral pathway into the direct and indirect striatal pathways and/or through dorsal projections into pFC. Next, we examined how the ease of stimulus identification affects the communication between cortex and BG in the optimal connectivity model for go trials. Previous work has shown adaptations in pFC, BG, and sensory regions when response strategies are adjusted (Egner & Hirsch, 2005) or when information accumulation is slowed (Bogacz, 2007; Palmer, Huk, & Shadlen, 2005). Therefore, modulations in communication strength could reflect the ease of action selection and/or adapted response strategies to optimize behavior.

An important final set of analysis examined the interplay between pFC–BG and perceptual systems during stop trials. Performance in the stop signal task is often described as an independent stochastic race, which asserts that initiation and inhibition processes are independent stochastic accumulation processes and compete for the first finishing time toward a certain bound or threshold (Band, van der Molen, & Logan, 2003; Logan, 1994; Logan & Cowan, 1984). In previous work, using ancestral graphs, we showed that the indirect and hyperdirect pFC–BG pathways combined best explain the pattern of brain activity during stop trials (Jahfari et al., 2011, 2012). Within both these studies, we assumed pFC to suppress only the planned motor output through the BG, once a stop signal was sounded. However, when the buildup of evidence for the go task is fast, successful stopping might require pFC to suppress the motor output via BG and/or slow down the accumulation of visual go task information to prevent it from transgressing a response threshold (Salinas & Stanford, 2013). These alternatives were examined by systematically comparing connectivity networks where pFC interacts with only BG (Jahfari et al., 2011, 2012) or with visual regions.

METHODS

Participants

Twenty-three young right-handed adults (seven men; mean age = 21.6 years, range = 19–26 years) participated in this study. All participants had normal or corrected-to-normal vision and provided written consent before the scanning session, in accordance with the Declaration of Helsinki. The ethics committee of the University of Amsterdam approved the experiment, and all procedures were in accordance with relevant laws and institutional guidelines.

Stimuli

A total of 30 grayscale full-front pictures of unfamiliar faces posing with a neutral expression (15 male, 15 female) were selected from the Radboud Face Database (Langner et al., 2010). Faces had neither hair nor glasses and were trimmed to remove all external features (neck, hairline). Three versions of each picture were then made to manipulate the SF of information available in the pictures for gender categorization (male, female). In all low SF (LSF) pictures, high SFs (HSFs) were removed by convolving the image with a symmetric Gaussian low-pass filter with a size of 20 pixels. The HSFs (e.g., the edges) were isolated by applying a range filter that returned the difference between the maximum and minimum values (using dilation and erosion function) in a 3 × 3 pixel neighborhood. This resulted in a total of 90 pictures where each picture had one version containing (1) all information (allSF), (2) only global information (LSF), or (3) only the local edge information (HSF). Note that the HSF and LSF conditions are comparable to the original stimuli with the exception of the experimental manipulation of removing either the HSFs or LSFs.

Experimental Protocol

A modified stop signal task (Figure 1A), with faces as go stimuli, was used to examine how perceptual information affects strategic responding during response initiation and withdrawal. To examine how different types of sensory information affect response strategies, go stimuli contained either all SFs (allSF) or they contained only LSFs or only HSFs (Figure 1B). On each trial, a white fixation cross was displayed on a black computer screen followed by a male or female face stimulus, indicating a left or right response with the index or middle finger of the right hand. Each trial started with a jitter interval of 500 to 700 msec (steps of 50 msec), during which a white fixation cross was presented in the center of the screen. A face stimulus was then presented for a period of 500 msec. On 30% of the trials, the go stimulus was followed by a high tone (stop signal). The stop signal delay (SSD) between the go stimulus and the stop signal was adjusted separately for each stimulus category (allSF, LSF, HSF) according to standard staircase methods to ensure convergence to p(inhibit) of .5. For example, if a stop signal was presented on an LSF trial and the participant responded ("failed stop"), then the SSD for the LSF staircase was reduced by 50 msec on the subsequent LSF stop trial; if the participant did not respond (i.e., "successful stop"), then SSD was increased by 50 msec. Initial SSD was set to 250 msec for all conditions. Instructions emphasized that participants should do their best to respond as quickly as possible while also doing their best to stop the response when an auditory stop signal occurred.

Each trial had a fixed duration of 2000 msec. In addition to the jitter (or interstimulus interval) used at the onset of each trial, all trials were further separated by a null trial where only the fixation cross was presented for another 2000 msec. All stimuli were presented on a back-projection screen that was viewed via a mirror system attached to the MRI head coil. Faces stimuli had a diameter of 4.2° and 5.3°. Before the MRI session,

participants performed a practice block of 60 trials to familiarize them with the task. In the MRI scanner, participants subsequently performed a total of 600 trials (go: 140 allSF, 140 LSF, 140 HSF; stop: 60 allSF, 60 LSF, 60 HSF) in four blocks of 150 trials each. The order of mapping rules for the index and middle fingers in gender categorization varied across participants and was reversed after each block.

Behavioral Analysis

The percentage choice errors and median RTs were calculated separately for each frequency condition for go and failed stop trials. For successful stop trials, SSRT was estimated separately for each frequency condition using the so-called "integration method" (Logan & Cowan, 1984). Repeated-measures ANOVAs were used to test how the available SF information (allSF, LSF, HSF) affects performance on go and stop trials. We could not obtain reliable SSRT estimates for two participants because of very long (average) SSD values during HSF trials. Therefore, these participants were excluded from all analysis. Because of the limited number of stimuli (30 faces), each face was presented six or seven times per conditions. The possibility of adaptation over time was excluded, as no significant differences were found across the four experimental blocks for either correct RT, $F(3, 186) < 1$, or percentage choice errors, $F(3, 186) < 1$.

Hierarchical Bayesian DDM

On the basis of go trial RT distributions of both correct and error responses, the formal Ratcliff DDM specifies two choice decisions as a noisy process, where information is accumulated over time (Figure 1C). This model can disentangle the speed or ease of evidence accumulation "drift rate" (v), the variability of evidence accumulation (η), the amount of evidence required for a decision boundary (a), the starting point of evidence accumulation (z), the variability of this starting point (S_z), the nondecision time (T_{er}), and nondecision time variability (S_t; Ratcliff & McKoon, 2008). To analyze the go RT data with the DDM, we used a recently developed hierarchical Bayesian estimation of DDM parameters, allowing for simultaneous estimation of individual and group parameters. Fits to individual participants are constrained by the group distribution but can vary from this distribution to the extent that their data are sufficiently diagnostic (Wiecki et al., 2013). To gain a deeper insight into how SF information affects choice RT (male or female), eight different models were investigated where three DDM parameters of interest were either fixed or varied across the three visual conditions: drift rate (v), boundary separation (a), and nondecision time (T_{er}). For each model, there were 100,000 samples generated from the posteriors, where the first 20,000 (burn-in) samples were discarded. Of the remaining 80,000 samples, every fifth sample was saved, resulting in a trace of 16,000 samples. The best model to describe the data across the three conditions was selected on the basis of deviance information criterion (Spiegelhalter, Best, Carlin, & der Linde, 2002), reflecting the best trade-off between fit quality and model complexity. To evaluate model performance, posterior predictives generated by the winning model were plotted on top of the observed correct and incorrect RT distributions for each participant. Figure 2 represents an example of three representative participants.

Image Acquisition

For registration purposes, a three-dimensional T1 scan was acquired before the functional runs (repetition time [TR] = 8.312 msec; echo time [TE] = 3.83 msec; flip angle [FA] = 8°; 1 mm slice thickness; 0 mm slice spacing; field of view [FOV] = 240 × 220 × 188). The fMRI data, collected during the stop task, were acquired in a single scanning session with four runs on a 3-T scanner (Philips Achieva TX, Andover, MA) using a 32-channel head coil. Each scanning run acquired 320 functional T2*-weighted echo-planar images (TR = 2000 msec; TE = 27.6 msec; FA = 76.1°; 3 mm slice thickness; 0.3 mm slice spacing; FOV = 240 × 121.8 × 240; 80 × 80 matrix; 37 slices, ascending slice order). Finally, a localizer task with faces, houses, objects, and scrambled scenes was used to identify FFA and LOC regions on an individual level (317 T2* weighted echo-planar images; TR = 1500 msec; TE = 27.6 msec; FA = 70°;

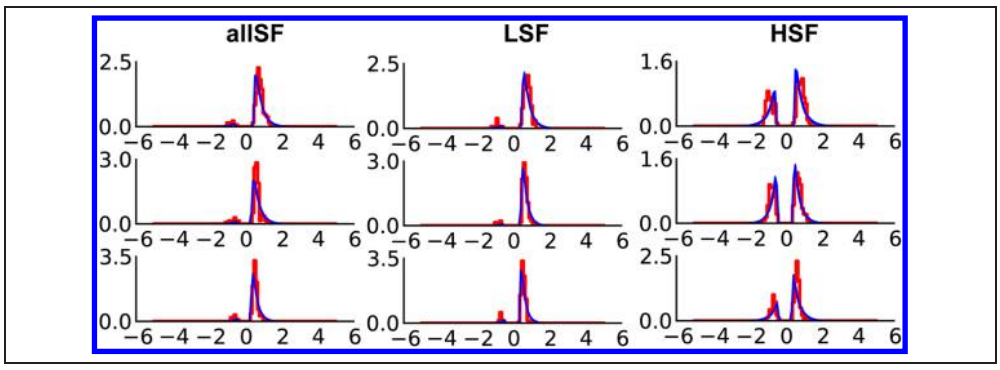

Figure 2. Examples of the posterior predictive (red lines) from the optimal Bayesian hierarchical DDM, on top of the observed normalized RT distributions (blue) for three representative participants (from top to bottom). Errors have been mirrored along the x axis to display correct and incorrect RT distributions in one plot.

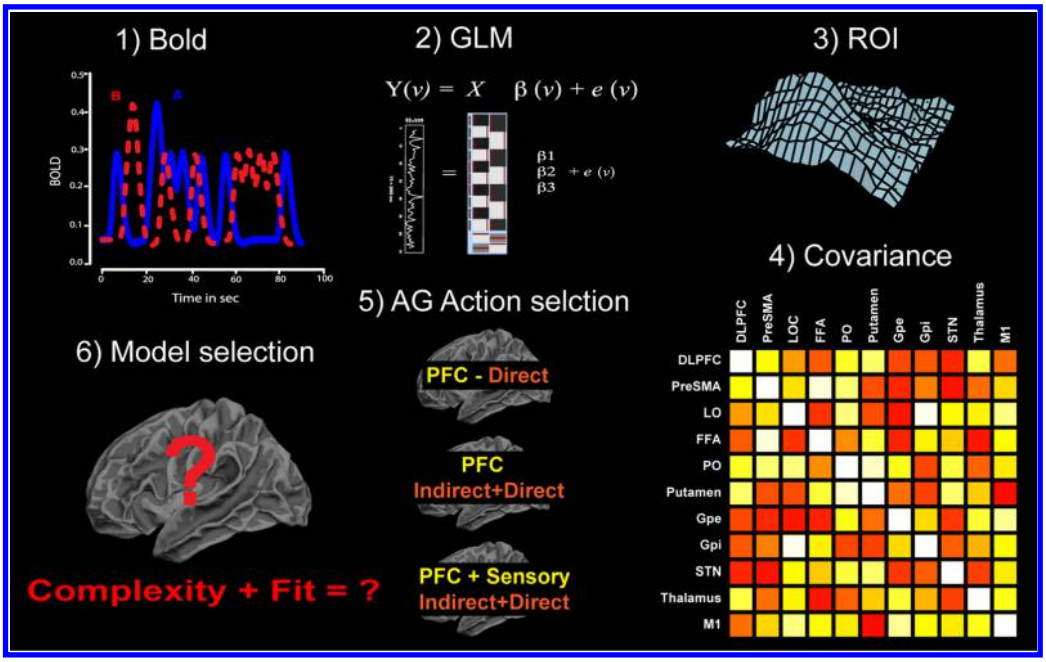

Figure 3. Graphical representation of all ancestral graphs connectivity steps. (1) Event related BOLD measurements are used as inputs to (2) the general linear model. (3) For go trials, 11 structural ROIs (DLPFC, pre-SMA, STN, putamen, GPi, GPe, thalamus, FFA, PO, LOC, left M1) determine the number of nodes (variables) in the ancestral graph and are used to determine the amplitudes for each trial separately for each region, such that the covariance matrix for these regions can be determined based on trial data (4). The data covariance matrix for each condition is then compared to all defined ancestral graph models, and each model obtains an BIC score (5). The model with the lowest BIC values represents the group network best and is selected (6). Please see Methods section for the selected ROIs during stop trials.

2.5 mm slice thickness; 0.25 mm slice spacing; FOV = 240 × 79.5 × 240; 96 × 96 matrix; 29 slices, ascending slice order).

Preprocessing

Analysis was performed using FEAT (fMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). The first six volumes were discarded to allow for T1 equilibrium effects. Preprocessing steps included motion correction, high-pass filtering in the temporal domain ($\sigma = 100$), and prewhitening (Woolrich, Ripley, Brady, & Smith, 2001). Spatial smoothing (FWHM = 5) was performed for the general fMRI analysis, but not on the ancestral graphs single-trial images. All functional data sets were individually registered into 3-D space using the participant's individual high-resolution anatomical images. The individual 3-D representation was then used to normalize the functional data into Montreal Neurological Institute (MNI) space by linear and nonlinear scaling.

General fMRI Analysis

The statistical evaluation was performed using the general linear model. The design matrix was generated with a hemodynamic response function and its first derivative with local autocorrelation correction (Woolrich et al., 2001). To replicate previous fMRI results using a stop signal task, we computed the following contrasts: (1) Successful Stop > Go and (2) Go > Successful Stop. Higher-level analysis was performed using FLAME Stage 1 and Stage 2 with automatic outlier detection (Beckmann, Jenkinson, & Smith, 2003). For the whole-brain analysis Z (Gaussianized T/F) statistic images were thresholded using clusters determined by $z > 2.3$ and $p = .05$ (GRFT).

Ancestral Graphs Connectivity

Ancestral graphs infer functional or effective connectivity by taking into account the distribution of BOLD activation per ROI over trials per participant and so is not dependent on the low temporal resolution of the time series in fMRI (Figure 3). A graphical model reflects the joint distribution of several neuronal systems with the assumption that for each individual the set of active regions is the same. The joint distribution (graphical model) of two nodes is estimated from the replications of condition-specific trials and not from the time series (Rissman, Gazzaley, & D'Esposito, 2004). With this method, we can infer three types of connections: (i) effective connectivity (directed connection →), (ii) functional connectivity (undirected connection –), and (iii) unobserved systems (bidirected connection ↔). Directed connections are regression parameters in the usual sense (denoted by β), and undirected connections are partial covariances (unscaled partial correlations; denoted by λ). The

bidirected connections refer to the covariance of the residuals from the regressions (denoted by ω). These three types of connection can be identified by modeling the covariance matrix (denoted by Σ) as

$$\Sigma = B^{-1}\begin{pmatrix}\Lambda^{-1} & 0\\ 0 & \Omega\end{pmatrix}(B^{-1})$$

where B contains the regression coefficients, Λ contains the partial covariances, and Ω contains the covariances between residuals. A random effects model is used to combine models across participants to then compare different models over the whole group using Bayes information criterion (BIC). The graph with the lowest BIC value will be selected. Individual (participant) fits are obtained by using an adjusted goodness-of-fit test, indicating whether the model explains the data well enough. For a full description of the ancestral graphs method, fit measures, and comparisons with other methods like structural equation modeling or dynamic causal modeling, please see citation (Waldorp et al., 2011).

With ancestral graphs, directionality (or the influence of one region to another) is inferred by comparing the joint probability distribution from a set of colliders (of graphs with different directionality or no directionality), which give rise to different conditional independencies (please see Waldorp et al., 2011, for an elaborate description). Effective connectivity (or influence) is defined as a regression, whereas functional connectivity with no direction is defined as a partial correlation. In general, the absolute regression (or correlation) value is an indication of connection strength, whereas the sign (+/−) or type (directed, undirected, or missing region) of the connection is informative for the relationship. For example, when there is a regression between variables A and B (A → B), a positive sign (i.e., positive slope) indicates that higher values in A predict (or relate to) higher values in B. However, when the regression sign is negative (anticorrelation), it means that higher values in A predict lower values in B. Although, ancestral graphs allow for directional inferences (A influences B), the explicit relationship between sign and inhibitory or excitatory connections is unknown. For the current study, inferences about the role of connections were based on the strength of connections or relationships with behavior.

ROI Selection, Single-trial Parameter Extraction, and Model Definition

The potential set of anatomical templates to define as ROIs was based on work reporting the involvement of these regions in decision-making (Heekeren, Marrett, & Ungerleider, 2008) and processes that involve the inhibition or initiation of a planned responses (Jahfari et al., 2011, 2012; Aron, 2011). We also included visual regions from the ventral stream (FFA and LOC) and dorsal stream (PO) that were either known to connect with pFC or BG (Kravitz et al., 2013) or reported to interact with top–down goals (Egner & Hirsch, 2005). The LOC is not directly involved in face categorization but plays an important role in transforming and segmenting visual input for subsequent processing. Because the dorsal pathway is known to project into DLPFC (Kravitz et al., 2013), PO (important for the transformation of visual input) was included in the model space as a potential link between the dorsal pathway and pFC. With the exception of the DLPFC and STN templates, all anatomical ROI templates (and coordinates) were defined from cortical and subcortical atlases available in FSL. The DLPFC template was obtained from a recent study, linking especially the posterior part to action execution (Cieslik et al., 2013). The STN template was derived from a previous study using ultrahigh 7 T scanning (Forstmann et al., 2012).

For go trials, this approach resulted in the selection of 11 key regions including the posterior DLPFC [center of gravity (cog): (−)37, 37, 27], pre-SMA [cog: (−)9, 25, 50], putamen [cog: (L) −25, 1, 1; (R) 26, 2, 1], globus pallidus externa (GPe) [cog: (−)21, −4, 0], globus pallidus interna (GPi) [cog: (−)17, −6, −4], thalamus [cog: (L) −10, −19, 7; (R) 11, −18, 7], STN [cog: (L) −10, −13, −7; (R) 10, −13, −7], primary motor cortex (M1) [cog: −18, −26, 61], PO [cog: (L) −50, −18, 17; (R) 51, −14, 18], LOC [cog: (L) −41, −81, −7; (R) 41, −79, −7], and the FFA [cog: (L) −42, −48, −17; (R) 45, −45, −19]. With the exception of M1, all selected go trial ROIs were bilateral. A left-hemisphere ROI for M1 was used because all participants responded only with their right hand.

On the basis of the dominant right hemispheric involvement of both BG and pFC regions during response inhibition (e.g., Garavan, Ross, & Stein, 1999), all selected regions for the stop network were right hemispheric with the exception of visual regions (FFA and LOC). Specifically, to examine the pattern of coactivation within the brain during stop trials, a set of nine regions including the rIFG [cog: 51, 19, 17], pre-SMA [cog: 9, 25, 50], rSTN [cog: 10, −13, −7], rGPe [cog: 21, −4, 0], rGPi [cog: 17, −6, −4], rCaudate [cog: 13, 10, 10], rThalamus [cog: 11, −18, 7], LOC [cog go], and FFA [cog go] were selected.

Ancestral Graphs Model Fit Procedure

For each ROI, we obtained the standardized predictor estimates (βs) of only the active voxels (F test) per participant, per trial/condition. Activity was averaged across active voxels in each ROI to obtain a single parameter (averaged normalized β estimate) for each trial. During correct go trials, the average number of parameters per ROI for responses with the right-hand index finger was 64.5 (SD = 4.7) for allSF, 62.3 (SD = 4.7) for LSF, and 42.4 (SD = 5.8) for HSF faces. When participants responded with their right-hand middle finger, this average was 63.6 (SD = 5.3) for allSF, 61.3 (SD = 5.6) LSF, and 41.9 (SD = 7.9) HSF faces. During successful stop trials,

the average number of parameters per ROI was 27.7 (SD = 4.7) for allSF stop trials, 25.9 (SD = 3.8) for LSF stop, and 28.0 (SD = 6.0) for HSF stop. The average number of trials is given, because the number of successful stops and correct go responses varied across participants and conditions. Error trials and misses were excluded from ancestral graphs analysis. Connectivity analysis was conducted in R-Cran (version 3.0.2), including the packages ggm (version 1.995-3), graph (version 1.40.0), and RBGL (version 1.38.0).

For go and stop trials, model evaluations were performed for each SF condition separately. Model selection was based on lowest BIC (and later evaluated through fit) and not restricted to identical models across differing levels of sensory evidence for either go or stop trials.

On go trials, definition of the winning model proceeded in three phases. First, we determined that a combined direct–indirect pathway with top–down projections from pFC into the putamen results in the lowest BIC, with good fits across participants. Next, we added projections from parietal and temporal nodes into the BG, pFC, or both and discovered that the addition of these raw visual nodes improves BIC in all conditions. Finally, we explored the relationship between visual and pFC regions. Functional connectivity between PO, DLPFC, and pre-SMA improved BIC values for the winning model in all conditions. Because both PO and pFC influence the putamen, ancestral graphs cannot infer directionality from the possibilities between PO and pFC. However, replacement of the PO-to-putamen connection by (1) PO-to-DLPFC, (2) PO-to-pre-SMA, or (3) PO-to-pFC all increased BIC values substantially, with a minimum increase of 1375 points. Therefore, the PO-to-putamen influence combined with functional connectivity between PO and pFC (see Model 5) was most representative, within the possibilities of ancestral graphs. For two participants, the model with the lowest BIC did not fit (using a goodness-of-fit test, see Waldorp et al., 2011) the observed data. Therefore, these two participants were excluded from connectivity strength evaluations during go trials, as the poor fit made parameter evaluations unjust.

For stop trials, we first tested a previously published hyperdirect–indirect model with top–down projections from both the rIFG and pre-SMA into the rSTN (hyperdirect pathway) and rCaudate (indirect pathway). During this first step, visual regions were already included in the model space but contained no connections toward BG or pFC. We next examined how visual information is integrated into the BG, during stop trials, by extending the traditional hyperdirect–indirect model with visual (i.e., FFA and LOC) to BG (i.e., the rSTN or rCaudate) connections. Selecting the optimal model from this step, we proceeded by evaluating visual (FFA, LOC) to pFC (rIFG, and pre-SMA) and, finally, pFC to visual projections. The optimal model to describe the pattern of activity in the selected ROIs, across trials, per condition, was selected on the basis of lowest BIC.

Model selection was based on BIC values obtained from a random effects model that combined models across participants and, as such, resulted in a pooled BIC per model. This approach allowed us to select the most representative model across the whole group. For both go and stop trials, these results were supported by individual BIC values across participants (please see Table 5 for go models and Table 6 for stop models).

RESULTS

Behavioral Performance

Before examining how visual information interacts with connectivity networks for response initiation or inhibition, we first established how the manipulation of SF information affects correct and incorrect RTs or the efficiency to withdraw a planned response.

Table 1 gives an overview of the behavioral data. As expected, increased RTs, $F(2, 40) = 42.6$, $p < .001$ (Figure 4A), and percentage choice error, $F(2, 40) = 255.9$, $p < .001$ (Figure 4B), indicated that gender categorization in faces becomes more difficult with the removal of SF. A hierarchical DDM where the speed of information accumulation "drift rate" (v), the amount of evidence required before reaching a decision "boundary" (a), and nondecision time (T_{er}) were allowed to vary across SF conditions, best predicted the observed RT on go trials (Table 2). Inspection of model parameters showed that the accumulation of evidence progressed more slowly with the removal of SF information, $F(2, 40) = 212.3$, $p < .001$ (Figure 4D). The speed of information accumulation was lowest when only the detailed HSF information was presented. Notably, the slower accumulation of evidence for especially HSF faces motivated participants

Table 1. Behavioral Data Overview

	allSF	*LSF*	*HSF*
Go			
Median RT (msec)	718.2 (144.8)	745.5 (135.8)	798.7 (146.3)
Choice errors (%)	5.5 (4.4)	8.3 (5.4)	36.5 (8.0)
Stop Respond			
Median RT (msec)	705.2 (154.7)	716.8 (139.2)	754.3 (160.7)
Stop Inhibit			
SSD (msec)	512.4 (213.0)	555.6 (218.0)	596.8 (202.7)
P_resp	0.54 (0.06)	0.56 (0.06)	0.52 (0.08)
SSRT (msec)	223.7 (94.7)	217.5 (102.8)	212.1 (94.7)

Values in parentheses are *SD*s. P_resp = probability of responding given the stop signal.

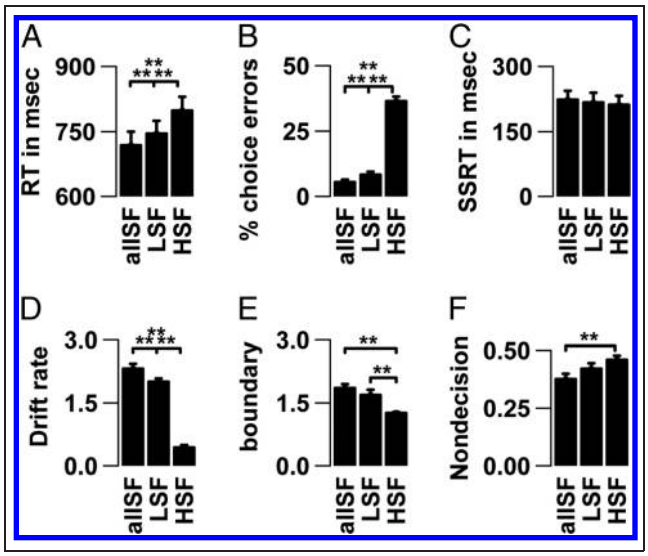

Figure 4. Across conditions, response strategies were adjusted when gender categorization was more difficult. Both median correct RT (A) and the percentage choice errors (B) increased with the removal of SF, whereas the efficiency to stop was unaffected (C). Formal analysis of the correct and incorrect RTs indicated that the ease of information accumulation is slowed when SF information is removed (D), while at the same time participants require less information before making a decision to respond (E). Removal of SF further affected nondecision time related to visual encoding and muscle activation (F) [error bars reflect *SEM*; ** = p(post hoc pairwise comparisons) < .01; sidak was used to adjust for multiple comparisons].

to change response strategies by reducing the amount of evidence required before reaching a decision, $F(2, 40) = 17.4$, $p < .001$ (Figure 4E). There was also a main effect of SF information on nondecision time, such that nondecision time was higher for HSF pictures relative to allSF pictures, $F(2, 40) = 6.9$, $p < .01$ (Figure 4F). This can be seen as a sanity check because it should indeed be harder to encode diminished visual input. Therefore, the sparse information in HSF pictures might affect not only the decision process but also the very early encoding of stimulus information or muscle initiation.

During stop signal trials, removal of SF information prolonged RTs on failed stop trials, $F(2, 40) = 6.8$, $p < .01$, and the SSD that was adjusted separately for each frequency condition, $F(2, 40) = 12.0$, $p < .01$. In contrast to our previous observations (Jahfari et al., 2013), the time needed to successfully stop a response was not influenced by the type of SF information presented, $F(2, 40) < 1$ (Figure 4C). This might be related to the fact that participants were lying horizontally in a noisy MRI scanner (Hommel, Fischer, Colzato, van den Wildenberg, & Cellini, 2011) or longer trial lengths chosen for the fMRI design (Jahfari et al., 2011).

Regional Brain Activation during Stop and Go Trials

A conventional set of contrasts used in the stop signal task was computed to evaluate the reliability of the fMRI data set. In line with previous findings, contrasting successful stop with go trials activated the IFG, pre-SMA, and inferior parietal lobe (Table 3). After cluster thresholding, no significant clusters remained in the STN. Contrasting go > successful stop resulted in higher activity levels in the left primary motor cortex, cerebellum, and striatum.

A Sensory Informed Cortico-BG Network for Action Selection

How do visual, pFC, and BG regions communicate to implement a selected motor plan? Using ancestral graphs, random effects analysis across the whole group indicated that a network comprising both the direct and indirect pathway best described the pattern of activity during all go trials, irrespective of SF information (Tables 4 and 5). Importantly, the optimal model (with the lowest BIC) contained effective connectivity projections from both pFC and visual areas into the BG (Figure 5A). Note that the DLPFC, pre-SMA, PO, and FFA all showed an information flow into the putamen, whereas LOC transmitted information into the STN. Functional connectivity (with no direction specified) was found between the pre-SMA, DLPFC, and PO and supported reported structural dorsal pathway projections into pFC (Kravitz et al., 2013). No relationship was found between visual regions in the ventral pathway (i.e., FFA and LOC) and pFC. Note that the inclusion of visual areas increases the complexity of response selection networks but improves the BIC criterion substantially for all go trials (Table 4). This improvement is remarkable because BIC favors simple models and penalizes for all increases in complexity.

The LOC-to-STN connectivity was unexpected but is informative about the observed coactivation patterns in the recorded fMRI data during response selection. To evaluate whether this connection is indirect or caused by a third (missing) region that was not included in our model space, we redefined the connection from LOC to

Table 2. Model Selection with Hierarchical DDM

Model	*Free to Vary*	*DIC*
1	v, a, T_{er}	1225.7
2	v, T_{er}	1871.1
3	v, a	2664.7
4	a, T_{er}	2875.3
5	v	3057.0
6	T_{er}	3156.1
7	a	4480.2
8	Fix all	4525.3

Lower deviance information criterion values indicate a better balance between fit and complexity. v = drift rate; a = boundary; T_{er} = nondecision time.

Table 3. Location of Increased Activation for the Comparison of Go and Stop Trials

Anatomical Area	*Cluster Size (mm^2)*	*MNI Coordinates (mm)*			*Maximum Effect Size*
		x	*y*	*z*	
Successful Stop > Go					
IFG	4111	52	18	−4	5.79
IFG	913	−48	18	−6	4.2
Medial frontal cortex/pre-SMA	814	2	30	30	4.2
Inferior parietal lobe	6630	62	−38	8	5.79
Inferior parietal lobe	5272	−42	−20	−2	5.44
Go > Successful Stop					
Primary motor cortex	4955	48	21	8	7.48
Striatum	1617	6	32	40	4.09
Cerebellum	2058	−61	−47	25	5.12

MNI coordinates are those of the max gravity point for each cluster. Cluster thresholding with $z > 2.3$ and $p < .05$, whole brain-corrected.

STN as a missing region connection (bidirected arrow) in the full model. In all conditions, this step increased BIC values and suggested that a directed relationship is more representative (BIC–missing regions model: go allSF [L = 13,970, R = 13,705], GLSF [L = 13,298, R = 12,881], GHSF [L = 10,151, R = 10,675]). The degree to which this connection represents a direct or indirect anatomical projection is a matter for future empirical research.

A last improvement in BIC was observed when we modeled the STN and putamen to project directly into the thalamus, instead of first projecting into the GPi. One explanation for this finding might be the close proximity of these regions to the thalamus within the BG. Crucially, the number of participants where the model fitted the observed data decreased substantially when we removed the indirect pathway with both globus pallidus nodes (i.e., the number of participants where the model fits for left [L] and right [R] go responses: allSF [L = 0, R = 0], GLSF [L = 0, R = 0], GHSF [L = 9, R = 12]).

Finally, to confirm that differing error rates in behavior do not confound the optimal model shown in Figure 5A, we repeated all model selection steps using a random sampler, which equates the number of trials across SF conditions (this step was done 25 times to assure stability of selection). Consistent with the model fits across all trials, the optimal model was identical to Model 5 (Figure 5A). That is, across all 25 equal sample draws, BIC values for Model 5 were lower when compared to Models 1–4, with a minimum difference of 187 points.

Overall, these effective connectivity results imply that response selection is based on the communication of

Table 4. Model Fits for Go Trial the Direct and Indirect Pathway and the Extension with Sensory Regions during Response Selection

		Go allSF				*Go LSF*				*Go HSF*			
		Left Key		*Right Key*		*Left Key*		*Right Key*		*Left Key*		*Right Key*	
Model	*Specification*	*BIC*	*n*	*BIC*	*n*	*BIC*	*n*	*BIC*	*n*	*BIC*	*n*	*BIC*	*n*
1	pFC + BG direct	14053	17	13969	16	13348	18	13242	18	10175	21	10567	21
2	pFC + BG indirect	13857	21	13822	21	13206	21	12969	21	10149	21	10701	21
3	pFC + BG direct&indirect	13729	21	13713	21	12079	21	12854	21	10009	21	10554	21
4	pFC + BG direct&indirect + Sensory	12933	19	12818	18	12340	19	11984	20	9502	21	9996	21
5	**pFC + BG direct&indirect (no GPi) + Sensory**	**12271**	**19**	**12194**	**19**	**11703**	**21**	**11341**	**21**	**9150**	**21**	**9581**	**21**

Lower BIC values indicate a better balance between fit and complexity of the estimated model connections. *n* indicates the number of participants (total = 21) where the defined model actually fitted the pattern of activity in the a priori anatomically defined ROIs. For clarity, the winning model (Model 5) is printed in **bold**.

Table 5. Averaged Individual BIC Scores for Go Trials

		Go allSF				*Go LSF*				*Go HSF*			
		Left Key		*Right Key*		*Left Key*		*Right Key*		*Left Key*		*Right Key*	
Model	*Specification*	*BIC*	*SD*	*BIC*	*SD*	*BIC*	*SD*	*BIC*	*SD*	*BIC*	*SD*	*BIC*	*SD*
1	pFC + BG direct	753	(114)	749	(133)	719	(131)	714	(133)	560	(80)	578	(91)
2	pFC + BG indirect	743	(115)	741	(126)	712	(136)	700	(122)	558	(84)	584	(100)
3	pFC + BG direct&indirect	737	(119)	736	(128)	705	(134)	694	(123)	551	(85)	577	(99)
4	pFC + BG direct&indirect + Sensory	699	(113)	693	(125)	670	(127)	652	(121)	526	(80)	550	(96)
5	**pFC + BG direct&indirect (no GPi) + Sensory**	**667**	**(104)**	**663**	**(111)**	**639**	**(115)**	**622**	**(104)**	**510**	**(72)**	**530**	**(94)**

For clarity, the winning model (Model 5) is printed in **bold**.

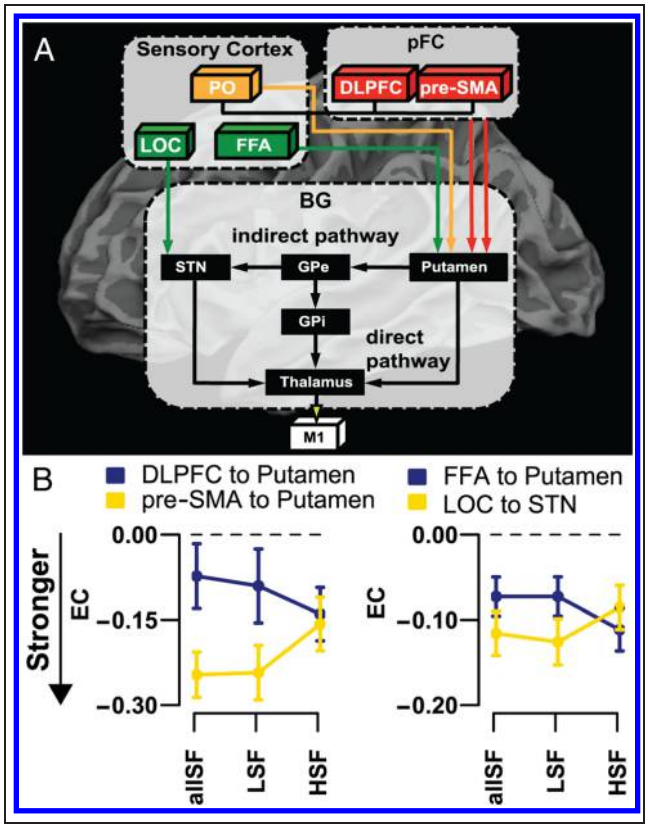

Figure 5. Sensory informed fronto-BG interactions in response selection. (A) Graphical representation of the most representative effective connectivity network on go trials. A cortico-BG model, with top–down projections from both the pFC and sensory areas into the BG, was identified to best represent the pattern of activity during all go trials. Directed arrows represent effective connectivity between two regions; undirected lines represent functional connectivity. (B) The strength of cortico-BG projections was adjusted dynamically when gender identification was more difficult. Across SF conditions (*x* axis), analysis of effective connectivity (EC) strength indicated a significant interaction between the two pFC regions projecting into BG (left) and two sensory regions projecting into BG (right). Note that in ancestral graphs effective connectivity is defined as a regression. Therefore, greater deviations from zero (dotted lines) indicate stronger connectivity strengths (error bars reflect *SEM*).

both raw stimulus information (sensory to BG) and abstract response plans (pFC to BG), through the direct and indirect striatal pathways. With the exception of the LOC-to-STN connection, the connections found between the dorsal and ventral stream with pFC or BG merge well with the known structural connections. However, the structure of the graph in the dorsal or ventral pathways was not specific to LSF or HSF information. Nevertheless, the specific relationship between SF and the dorsal/ventral route might prevail from a more time-sensitive approach.

Perceptually Motivated Cortico-BG Dynamics in Response Selection

Next, we examined how the cortex-to-BG regression strength (effective connectivity) is modulated when the visual input for response selection is degraded. Results showed that changes in categorization difficulty are accompanied by dynamic adjustments in the strength of communication from both pFC and visual cortex into BG. Within the ancestral graphs, method effective connectivity is defined as a regression from Region A to Region B. Therefore, stronger deviations from zero indicate stronger connection strengths. Two separate repeated-measures ANOVAs with the factors Region and SF information were defined to examine the influence of pFC (region defined as: DLPFC and pre-SMA) or visual cortex (region defined as: FFA and LOC) on BG. A significant interaction indicated that projections from the DLPFC into the putamen become stronger with the removal of SF information whereas projections from the pre-SMA into the putamen are weakened, $F(2, 36) = 4.4$, $p < .05$ (Figure 5B, left). Similarly, when inspecting visual-to-BG projections, a significant interaction showed increased FFA-to-putamen connectivity when face identification was more difficult while the LOC-to-STN connectivity was weakened, $F(2, 36) = 3.3$, $p < .05$ (Figure 5B, right). The ease of gender identification did not affect PO-to-BG connection strengths, $F(2, 36) < 1$.

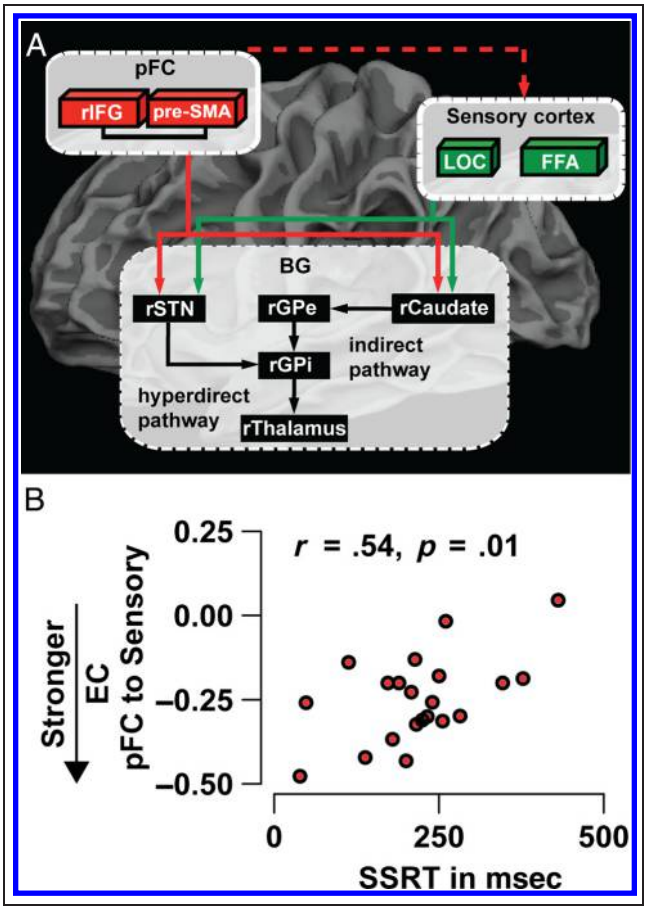

Figure 6. Sensory fronto-BG interactions during response inhibition. (A) Schematic of the winning hyperindirect connectivity model during successful stop trials. In all conditions, model representations (BIC values) improved when LOC and FFA were modeled to project directly into BG, as compared to the traditional hyperindirect model with no sensory input. Importantly, only when sensory information (for response selection) was not degraded (allSF faces), the optimal model contained additional effective connectivity projections (dotted red arrow) from pFC into sensory cortex to update or enforce the stop plan (i.e., pre-SMA to LOC and FFA, rIFG to LOC and FFA). Directed arrows represent effective connectivity. The black undirected line between the pre-SMA and rIFG represents functional connectivity. (B) When information for response selection was not degraded (allSF pictures), stronger projection strengths form pFC into sensory cortex were related to shorter inhibition times (more efficient stopping).

These interactions portray a dynamic interplay between regions involved with the transformation or representation of visual information (i.e., DLPFC and FFA) and regions involved with the evaluation of input (LOC) or response goals (pre-SMA). That is, regions involved with the representation of raw face information (FFA) and the processing of choice evidence (DLPFC) were found to increase connectivity toward the putamen, when gender identification became more difficult. Concurrently, regions involved with the evaluation of response strategies (pre-SMA) or the representation of scene saliency (LOC) were found to decrease connectivity strengths toward the BG.

To further specify the observed cortex–BG interactions, we explored the relationship between connectivity and behavioral indices. Although the interactions above seem consistent with the estimated changes in drift rate or response cautiousness, no direct relationship was found between hierarchical DDM parameters and cortico-BG connectivity strength (all ps $> .05$). Interestingly, however, activity levels within the DLPFC correlated positively with drift rates in all but the HSF condition—where drift rats were very low, r(allSF) $= .47$, $p < .05$, r(LSF) $= .44$, $p < .05$, r(HSF) $= .09$, $p > .05$. In contrast, weaker pre-SMA-to-striatum connections (closer to zero) were predictive for higher error rates during HSF trials ($r = .49$, $p < .05$), but not during allSF ($r = -.08$, $p > .05$) or LSF trials ($r = -.03$, $p > .05$). We return to the possible interpretation of these relationships in the discussion.

Response Inhibition: Visual Cortex, pFC, and BG Coordination

A final analysis examined the involvement of visual areas when a prepared response needs to be countermanded (i.e., the response plan in preparation needs to be withdrawn) and the visual input is no longer relevant to the task. Figure 6A shows a graphical schematic of the most representative model during successful stop trials. Replicating previous work, a right hemispheric pFC–BG model with the hyperdirect and indirect pathway showed a good fit to the pattern of brain activity during successful stop trials [BIC (n_fit): Stop allSF = 4679 (21), Stop LSF = 4495 (21), Stop HSF = 4771 (21)]. Importantly, in all conditions (i.e., allSF, LSF, HSF), model representations (BIC) improved when directed connections were added from LOC and FFA into, respectively, the rSTN (hyperdirect pathway) and rCaudate [indirect pathway; BIC (n_fit): Stop allSF = 4487 (21), Stop LSF = 4200 (21), Stop HSF = 4496 (21)]. Most notably, only when a stop signal followed allSF faces (i.e., nondegraded faces with high drift rate) the most representative model contained additional projections from pFC into visual cortex [BIC (n_fit): Stop allSF = 4428 (21), Stop LSF = 4222 (21), Stop HSF = 4512 (21); dotted red line Figure 6A, see Table 6 for individual BIC scores]. That is, only during

Table 6. Individual BIC Scores for Successful Stop Trials

	Stop allSF	*Stop LSF*	*Stop HSF*
Hyperindirect model (HIM)	290 (44)	280 (46)	295 (53)
HIM + sensory to BG	280 (44)	**265 (42)**	**280 (48)**
HIM + sensory to BG + pFC to sensory	**276 (45)**	265 (42)	280 (47)

Mean individual BIC scores with standard deviations within parentheses. Lower BIC values represent a more optimal model. With equal BIC scores, the model with the lowest number of parameters was evaluated as most representative. For clarity, the winning model, for each condition, is printed in **bold**.

successful allSF stop trials the most representative model contained four additional connections from the rIFG and pre-SMA into LOC and FFA. The optimal model (based on lowest BIC) during successful stop trials with degraded stimuli (either LSF or HSF) contained no projections form pFC into visual cortex.

The extra pFC projections into visual cortex (during allSF stop trials) might suppress the fast flow of perceptual information for the go task, such that the inhibition process can finish before the response selection process. To test this hypothesis, we examined the relationship between SSRT and the overall effective connectivity strength from pFC into visual cortex during successful allSF stop trials (overall = average connectivity for the four pFC–visual projections: (1) rIFG to LOC, (2) rIFG to FFA, (3) pre-SMA to LOC, and (4) pre-SMA to FFA). Results indicated a significant positive relationship between SSRT and the overall connectivity strength from pFC into visual cortex ($r = .54$, $p < .05$; Figure 6B). That is, faster inhibition times (lower SSRT) were related to stronger connectivity strengths from pFC into visual cortex. As can be seen in Figure 6B, two participants had SSRTs smaller than 100 msec. We included these participants because a boxplot analysis did not identify them as outliers. However, SSRTs around 50 msec are very low and hard to interpret. To exclude the possibility that unreliable estimates of SSRTs confound the relationship found, we repeated the correlation without these participants and again found a significant relationship between SSRT and pFC-to-visual regions connectivity strength ($r = .48$, $p < .05$).

These results specify the interaction between sensory and fronto-BG networks when a go plan is changed and a response needs to be withdrawn. Because a stop signal is always sounded after the stimulus presentation for the go task, the current study (with the slow fMRI resolution) cannot differentiate whether the flow of information from visual regions into BG is a residual of the response selection process or an essential part of the stop network. To assess the specificity of this sensory informed hyperdirect–indirect model to stop trials, we examined how well this model represents the pattern of activity during go trials (when no stop signal is sounded). In line with our previous findings, this stop model showed a poor fit for pure go trials (n_fit: Go allSF = 0; Go LSF = 0; Go HSF = 0) and was therefore specific to the process of response inhibition.

Note that, effective connectivity projections from pFC into visual regions were not observed during go trials and only improved stop trial model representations for those stimuli with the highest processing speed or drift rate (allSF faces). Therefore, the relationship between SSRT and pFC-to-visual connectivity indicates that, when information accumulation is easy and fast, pFC suppresses not only the release of a selected response through BG but additionally might intervene with the flow of information from visual cortex into BG so that the stop process can finish before the go process.

DISCUSSION

This study provides novel insights into how visual information is integrated into fronto-BG circuits to optimize action selection and indicates how the transmission of visual information is suppressed during response inhibition. On go trials, effective connectivity analysis showed that a direct–indirect cortico-BG network including both pFC and visual regions best represented the pattern of brain activity during action selection. Further inspection of model dynamics showed how the ease of stimulus identification is directly reflected in adjusted projections from both pFC and visual regions into BG. Bayesian parameter estimation using the DDM showed removal of SF to reflect the ease of information accumulation and the adjustment of response cautiousness in behavior.

During stop trials, a sensory informed hyperdirect–indirect pFC–BG network was most representative for the observed pattern of brain activity. Importantly, only when a stop signal followed easily processed nondegraded faces (i.e., high drift rate, low errors) did the optimal stop network contain four extra pFC-to-visual connections. This finding suggested that, when the accumulation of evidence toward a decision-threshold is fast, pFC might direct both BG and visual regions to suppress the motor output and possibly the buildup of visual input. In accord, stronger connectivity from pFC toward visual regions was related to more efficient inhibition times.

Race model frameworks suggest that a fast and accurate go process (high drift rates, low boundaries, and low errors) should be accompanied by an even faster stop process (White et al., 2014). We found that, when a stop signal is sounded, visual information for the go plan is still forwarded into BG, possibly as a residue. Model-based effective connectivity results further showed that, when information is processed fast and accurately for the go process (allSF faces), the stop process involves projections from pFC into BG and visual regions. Compelling work has suggested a prominent role for pFC–BG frameworks in the suppression of planned motor outputs (Aron, 2011; Munakata et al., 2011). In parallel, a largely separate field has documented the importance of pFC and visual systems during conflict resolution (Danielmeier, Eichele, Forstmann, Tittgemeyer, & Ullsperger, 2011; Egner & Hirsch, 2005), motor preparation (Clementz et al., 2010), and sustained control (Miller & D'Esposito, 2005). This study indicates how pFC interacts with BG and visual regions to disrupt a fast go process.

Performance in the stop signal task is often described as an independent stochastic race between the go process (activated here by the face stimulus) and the stop process (beep sound presented at a variable delay; Logan et al., 2014; Middlebrooks & Schall, 2014). By questioning how the activated go process might be suppressed in neurophysiological data, Boucher et al. (2007) proposed an interactive race model. Here, the go process is described as a single diffusion toward threshold, whereas

Figure 7. Schematic of the interactive race model, showing how the buildup rate of a go plan is inhibited on a successful stop trial. Figure reproduced with permission from Verbruggen and Logan (2008).

an interactive stage of the stop process—that follows an independent stop stage—inhibits the growth of activation for the go response (Figure 7). The efficiency of stopping (SSRT) reflects the duration of the independent stage (the time where both stop and go progress independent of each other). During the independent stage, the stop process is initiated (after a variable SSD) but has not yet become active in suppressing the buildup of the go response. Thus, SSRT reflects the time between the initiation of the stop process (sounding of beep) and onset of the interactive stage, where active suppression of the go process begins. Responses are inhibited if the interactive stage prevents the go process from reaching threshold.

In a recent study, Salinas and Stanford (2013) suggested that successful cancelation of a response relies on (1) how soon the stop process can start interfering with the buildup rate of the go process and (2) the strength of change imposed by the active stop process. This study relates higher pFC-to-visual connectivity to shorter SSRT. Therefore, when the buildup rate of the go process is fast, the communication between pFC and visual cortex seems predictive for the onset of the interactive stage (active suppression). On-time interference with visual regions that carry information for the go process could be achieved by blocking the visual input (Palmeri, Schall, & Logan, 2015) or the allocation of attentional resources (Chiu & Egner, 2015), such that the buildup rate is slowed (Salinas & Stanford, 2013).

Consistent with our previous connectivity work, no relationship was found between SSRT and pFC-to-STN connections (Jahfari et al., 2011, 2012). The hyperdirect pathway is described to exert fast and powerful excitatory effects on the output nuclei of BG. Moreover, signal conduction within this pathway is thought to be faster than the direct or indirect pathway (Nambu, Tokuno, & Takada, 2002). Although our connectivity approach with fMRI is consistently able to detect direct connections from pFC-to-STN, the relationship between the hyper-indirect route and SSRT is better captured with more time-sensitive designs. In example, one study with scalp EEG recordings showed that simulation of the STN improves the efficiency of stopping (SSRT) and leads to increased IFG activity (Swann et al., 2011).

On go trials, our model-driven connectivity approach was motivated by classical accounts of action selection (Mink, 1996) and advances in the field of perceptual decision-making (Gold & Shadlen, 2007; Schall, 2001). In compliance, we show that a model best represents action selection when visual information is linked to both pFC and BG. Supporting theoretical frameworks, combining direct and indirect striatal pathways best represented the observed fMRI data. However, incorporating the information flow from visual cortex identified the optimal model. In accord with anatomical and neurophysiological evidence, visual information within the ventral stream projected into the BG whereas the dorsal stream was connected with both the BG and pFC (Kravitz et al., 2013). In both humans and primates, the ventral pathway (including FFA and LOC) has been reported to carry detailed stimulus information into the striatum, whereas the dorsal stream (including parietal cortex and PO) links to the striatum and further relates information into pFC (Yamamoto, Monosov, Yasuda, & Hikosaka, 2012; Saint-Cyr, Ungerleider, & Desimone, 1990). Our results comply and extend these reports by specifying how information is integrated into the BG to select an overt response.

At subcortical levels, go trial activity was best represented when both the facilitating direct pathway and the more deliberate indirect pathway originated from the putamen. Previous work has described the striatum to first constrain or weigh potential response options (Robbins & Brown, 1990). Consistent with such claims, our connectivity analysis identified the putamen to receive inputs from DLPFC, pre-SMA, PO, and FFA. Notably, LOC, implicated in the integration of scene information projected directly into the STN. Much is known about the relationship between the goal-oriented pFC and the response slowing or "braking" STN. This study implies the STN to also receive inputs from the ventral visual stream, associated with the evaluation of stimuli used for action selection. Anatomically, the STN has been reported to intervene with the pallidum (Carpenter, Carleton, Keller, & Conte, 1981), thalamus (Rico et al., 2010), and cerebral cortex (Degos, Deniau, Le Cam, Mailly, & Maurice, 2008). At the same time, the STN is known to receive inputs (among others) from both pFC and BG nodes (Nambu et al., 2002). Unfortunately, very little is known about how the STN is connected to the ventral pathway. Therefore, follow-up anatomical studies focusing on how the STN is connected to the ventral pathway are critical for a firm interpretation of the coactivation patterns found between LOC and STN.

After model selection, this study further examined how the ease of stimulus identification affects the strength of information relay between the cortex and BG. In behavior, prolonged RTs and increased error rates confirmed that gender identification in faces becomes more difficult with the selective removal of SF information. Formal RT analysis indicated that, when the ease of information accumulation is slowed, the acquisition of additional information comes at a greater cost. Thus, when gender identification was most difficult and processed slowly, participants reduced the criteria for the amount of evidence required to respond.

The ease of stimulus identification was also reflected in the strength of information relay between the cortex and BG. Specifically, we observed reduced connectivity between the pre-SMA and putamen when stimulus identification was most difficult, while concurrently the influence of DLPFC on the putamen was increased. Previous work has linked the communication between the pre-SMA and striatum to the evaluation and adjustment of response strategies (Jahfari et al., 2011, 2012; Forstmann et al., 2008). In this study, weaker pre-SMA-to-striatum connectivity was predictive for higher error rates only when gender discrimination was very difficult and information accumulation progressed slowly (HSF trials). In contrast, DLPFC activity correlated with the rate of information accumulation in all but the difficult HSF condition. Activity levels within the DLPFC have been repeatedly linked to evidence accumulation or the computation of a difference score between the desired and undesired response choice (Philiastides et al., 2011). One possible role for the increased DLPFC-to-striatum connectivity might be to relate all noisy choice information into the striatum, whereas reduced connection strengths from the pre-SMA allow for more impulsive or uncertain responses (higher error rates).

Importantly, comparable dynamics were observed when inspecting visual–BG projections from the visual ventral stream. With the removal of SF, FFA projections into the putamen became stronger, whereas LOC-to-STN projections were weakened. The increased FFA-to-putamen possibly represents the transmission of raw face information alongside the top–down input from DLPFC. A recent review described the ventral pathway (including LOC) to contain neural representations of object quality (Kravitz et al., 2013). The LOC-to-STN relationship might serve the communication of object quality based on SF, possibly to justify a swift acquisition of additional information at a low cost (with allSF faces).

Although connection strengths of both pFC regions and the ventral visual stream were adjusted dynamically with the removal of SF, we found no effects when inspecting projections from PO to striatum. One possible explanation might be the focus of this region on processing spatial and motion information. Because our stimuli did not entail the detection of movement, the function of this region in our experiment might be the integration of information with pFC or the presentation of all potential action intentions (i.e., right-hand index and middle fingers) to the striatum (Cisek & Kalaska, 2010). However, more research using both moving and still objects is required to shed light on this matter.

The current study is not without limitations. First, the race between male/female or go/stop is typically resolved within a framework of milliseconds and can only be partially captured with fMRI. The analysis of coactivation patterns across trials gives valuable insights into how key brain regions interact during selection or inhibition but contains no temporal information. Future studies using high temporal resolution techniques are key to further specify the role of cortico–cortico and cortico–subcortical interactions during voluntary selection or control. Second, ancestral graphs only incorporate linear interactions. As such, we cannot exclude the possibility of nonlinear effects. Previous work using dynamic causal modeling reported prediction errors within the striatum to modulate the relationship between pFC and sensory regions (Den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010). Because prediction errors influence choice certainty, future work should examine the possibility of nonlinear interactions between visual cortex, BG, and pFC when the quality of visual input is manipulated. Finally, the mechanisms of decision-making within the brain were only investigated through the separation of SF conditions across go and stop trials. This approach allowed us to investigate how SF modulates some (but not all) latent processes of decision-making such as response cautiousness, accumulation speed, and stopping efficiency. Other factors that could shape the neural process of decision-making include the motivation to perform the task, the relative emphasis on go or stop, or trial history (Gold & Shadlen, 2007). These factors should be systematically evaluated in future experiments focusing on cortico-BG networks of decision-making.

In summary, using a model-driven effective connectivity approach, we specify how information from visual regions is used or modulated by classic fronto-BG circuits to optimize either response selection or inhibition. By changing the quality of stimulus information, we additionally show how the relay of information between the cortex and BG is adjusted dynamically when response selection becomes more uncertain. Finally, we argue that when visual information for the go task is processed fast, pFC updates action plans by inhibiting the motor output via BG and direct modulation of the visual buildup rate so that response selection can be turned into a full stop (response inhibition).

Acknowledgments

This work was supported by Mozaiek grant 017.005.107 (S. J.) from the Netherlands Organization for Scientific Research (NWO).

Reprint requests should be sent to Sara Jahfari, Department of Cognitive Psychology, VU University, Van der boechorststraat 1, 1081 BT Amsterdam, The Netherlands, or via e-mail: sara.jahfari@gmail.com.

REFERENCES

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience, 14,* 1338–1444.

Altmann, C. F., Deubelius, A., & Kourtzi, Z. (2004). Shape saliency modulates contextual processing in the human lateral occipital complex. *Journal of Cognitive Neuroscience, 16,* 794–804.

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry, 69,* e55–e68.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *The Journal of Neuroscience, 26,* 2424–2433.

Band, G. P., van der Molen, M. W., & Logan, G. D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologica, 112,* 105–142.

Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2003). General multilevel linear modeling for group analysis in fMRI. *Neuroimage, 20,* 1052–1063.

Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences, 11,* 118–125.

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review, 114,* 376–397.

Burnham, K. P., & Anderson, D. R. (2004). Multimodel inference understanding AIC and BIC in model selection. *Sociological Methods & Research, 33,* 261–304.

Carpenter, M. B., Carleton, S. C., Keller, J. T., & Conte, P. (1981). Connections of the subthalamic nucleus in the monkey. *Brain Research, 224,* 1–29.

Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., et al. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience, 14,* 1462–1467.

Chiu, Y.-C., & Egner, T. (2015). Inhibition-induced forgetting: When more control leads to less memory. *Psychological Science, 26,* 27–38.

Cieslik, E. C., Zilles, K., Caspers, S., Roski, C., Kellermann, T. S., Jakobs, O., et al. (2013). Is there "one" DLPFC in cognitive action control? Evidence for heterogeneity from co-activation-based parcellation. *Cerebral Cortex, 23,* 2677–2689.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience, 33,* 269–298.

Clementz, B. A., Gao, Y., McDowell, J. E., Moratti, S., Keedy, S. K., & Sweeney, J. A. (2010). Top–down control of visual sensory processing during an ocular motor response inhibition task. *Psychophysiology, 47,* 1011–1018.

Danielmeier, C., Eichele, T., Forstmann, B. U., Tittgemeyer, M., & Ullsperger, M. (2011). Posterior medial frontal cortex activity predicts post-error adaptations in task-related visual and motor areas. *The Journal of Neuroscience, 31,* 1780–1789.

Degos, B., Deniau, J.-M., Le Cam, J., Mailly, P., & Maurice, N. (2008). Evidence for a direct subthalamo-cortical loop circuit in the rat. *The European Journal of Neuroscience, 27,* 2599–2610.

Den Ouden, H. E. M., Daunizeau, J., Roiser, J., Friston, K. J., & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *The Journal of Neuroscience, 30,* 3210–3219.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience, 32,* 3612–3628.

Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience, 8,* 1784–1790.

Fiorentini, A., Maffei, L., & Sandini, G. (1983). The role of high spatial frequencies in face perception. *Perception, 12,* 195–201.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences, 105,* 17538–17542.

Forstmann, B. U., Keuken, M. C., Jahfari, S., Bazin, P. L., Neumann, J., Schäfer, A., et al. (2012). Cortico-subthalamic white matter tract strength predicts interindividual efficacy in stopping a motor response. *Neuroimage, 60,* 370–375.

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks, 19,* 1120–1136.

Garavan, H., Ross, T. J., & Stein, E. A. (1999). Right hemispheric dominance of inhibitory control: An event-related functional MRI study. *Proceedings of the National Academy of Sciences, 96,* 8301–8306.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–574.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences, 15,* 20–25.

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience, 9,* 467–479.

Herz, D. M., Christensen, M. S., Bruggemann, N., Hulme, O., Ridderinkhof, K. R., Madsen, K. H., et al. (2014). Motivational tuning of fronto-subthalamic connectivity facilitates control of action impulses. *The Journal of Neuroscience, 34,* 3210–3217.

Hikosaka, O., & Isoda, M. (2010). Switching from automatic to controlled behavior: Cortico-basal ganglia mechanisms. *Trends in Cognitive Sciences, 14,* 154–161.

Hommel, B., Fischer, R., Colzato, L. S., van den Wildenberg, W. P., & Cellini, C. (2011). The effect of fMRI (noise) on cognitive control. *Journal of Experimental Psychology: Human Perception and Performance, 38,* 290–301.

Jahfari, S., Ridderinkhof, K. R., & Scholte, H. S. (2013). Spatial frequency information modulates response inhibition and decision-making processes. *PLoS ONE, 8,* e76467.

Jahfari, S., Verbruggen, F., Frank, M. J., Waldorp, L., Colzato, L., Ridderinkhof, K. R., et al. (2012). How preparation changes the need for top–down control of the basal ganglia when inhibiting premature actions. *The Journal of Neuroscience, 32,* 10870–10878.

Jahfari, S., Waldorp, L., van den Wildenberg, W. P., Scholte, H. S., Ridderinkhof, K. R., & Forstmann, B. U. (2011). Effective connectivity reveals important roles for both the hyperdirect (fronto-subthalamic) and the indirect (fronto-striatal-pallidal) fronto-basal ganglia pathways during response inhibition. *The Journal of Neuroscience, 31,* 6891–6899.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience, 17,* 4302–4311.

Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral visual pathway: An expanded neural framework for the processing of object quality. *Trends in Cognitive Sciences, 17,* 26–49.

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H., Hawk, S. T., & Knippenberg, A. (2010). Presentation and validation

of the Radboud Faces Database. *Cognition and Emotion, 24,* 1377–1388.
Logan, G. D. (1994). On the ability to inhibit thought and action: A user's guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 189–239).
Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91,* 295–327.
Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review, 121,* 66–95.
Middlebrooks, P. G., & Schall, J. D. (2014). Response inhibition during perceptual decision making in humans and macaques. *Attention, Perception & Psychophysics, 76,* 353–366.
Miller, B. T., & D'Esposito, M. (2005). Searching for "the top" in top–down control. *Neuron, 48,* 535–538.
Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology, 50,* 381–425.
Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences, 15,* 453–459.
Nambu, A. (2009). Functions of direct, indirect and hyperdirect pathways. *Brain Nerve, 61,* 360–372.
Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal "hyperdirect" pathway. *Neuroscience Research, 43,* 111–117.
Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision, 5,* 376–404.
Palmeri, T. J., Schall, J. D., & Logan, G. D. (2015). Neurocognitive modeling of perceptual decision making. In *Oxford handbook of computational and mathematical psychology*. Oxford: Oxford University Press.
Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology, 21,* 980–983.
Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20,* 873–922.
Rico, A. J., Barroso-Chinea, P., Conte-Perales, L., Roda, E., Gómez-Bautista, V., Gendive, M., et al. (2010). A direct projection from the subthalamic nucleus to the ventral thalamus in monkeys. *Neurobiology of Disease, 39,* 381–392.
Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage, 23,* 752–763.
Robbins, T. W., & Brown, V. J. (1990). The role of the striatum in the mental chronometry of action: A theoretical review. *Reviews in the Neurosciences, 2,* 181–214.
Saint-Cyr, J. A., Ungerleider, L. G., & Desimone, R. (1990). Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. *Journal of Comparative Neurology, 298,* 129–156.
Salinas, E., & Stanford, T. R. (2013). The countermanding task revisited: Fast stimulus detection is a key determinant of psychophysical performance. *The Journal of Neuroscience, 33,* 5668–5685.
Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience, 2,* 33–42.
Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., & Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience, 16,* 1118–1124.
Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron, 79,* 217–240.
Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society, Series B, Statistical Methodology, 64,* 583–639.
Swann, N. C., Poizner, H., Houser, M., Gould, S., Greenhouse, I., Cai, W., et al. (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. *The Journal of Neuroscience, 31,* 5721–5729.
Tosoni, A., Galati, G., Romani, G. L., & Corbetta, M. (2008). Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. *Nature Neuroscience, 11,* 1446–1453.
Verbruggen, F., & Logan, G. D. (2008). Response inhibition in the stop signal paradigm. *Trends in Cognitive Sciences, 12,* 418–424.
Verbruggen, F., Stevens, T., & Chambers, C. D. (2014). Proactive and reactive stopping when distracted: An attentional account. *Journal of Experimental Psychology: Human Perception and Performance, 40,* 1295–1300.
Waldorp, L., Christoffels, I., & van de Ven, V. (2011). Effective connectivity of fMRI data using ancestral graph theory: Dealing with missing regions. *Neuroimage, 54,* 2695–2705.
White, C. N., Congdon, E., Mumford, J. A., Karlsgodt, K. H., Sabb, F. W., Freimer, N. B., et al. (2014). Decomposing decision components in the stop signal task: A model-based approach to individual differences in inhibitory control. *Journal of Cognitive Neuroscience, 26,* 1601–1614.
Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics, 7,* 14.
Woolrich, M. W., Ripley, B. D., Brady, M., & Smith, S. M. (2001). Temporal autocorrelation in univariate linear modeling of fMRI data. *Neuroimage, 14,* 1370–1386.
Yamamoto, S., Monosov, I. E., Yasuda, M., & Hikosaka, O. (2012). What and where information in the caudate tail guides saccades to visual objects. *The Journal of Neuroscience, 32,* 11005–11016.