◆ Human Brain Mapping 38:1333–1346 (2017) ◆

Filling the Void—Enriching the Feature Space of Successful Stopping

René J. Huster,[1,2]* Signe Schneider,[3] Christina F. Lavallee,[4] Stefanie Enriquez-Geppert,[5] and Christoph S. Herrmann[6,7]

[1]Department of Psychology, University of Oslo, Norway
[2]Psychology Clinical Neurosciences Center, University of New Mexico, Albuquerque, New Mexico, USA
[3]Department of Systems Nseuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany
[4]Center for Mind/Brain Sciences (CIMeC), University of Trento, Italy
[5]Department of Clinical and Developmental Neuropsychology, University of Groningen, The Netherlands
[6]Experimental Psychology Lab, Department of Psychology, Cluster of Excellence "Hearing4all", European Medical School, Carl von Ossietzky University, Oldenburg, Germany
[7]Research Center Neurosensory Science, Carl-von-Ossietzky University Oldenburg, Oldenburg, Germany

Abstract: The ability to inhibit behavior is crucial for adaptation in a fast changing environment and is commonly studied with the stop signal task. Current EEG research mainly focuses on the N200 and P300 ERPs and corresponding activity in the theta and delta frequency range, thereby leaving us with a limited understanding of the mechanisms of response inhibition. Here, 15 functional networks were estimated from time–frequency transformed EEG recorded during processing of a visual stop signal task. Cortical sources underlying these functional networks were reconstructed, and a total of 45 features, each representing spectrally and temporally coherent activity, were extracted to train a classifier to differentiate between go and stop trials. A classification accuracy of 85.55% for go and 83.85% for stop trials was achieved. Features capturing fronto-central delta- and theta activity, parieto-occipital alpha, fronto-central as well as right frontal beta activity were highly discriminating between trial-types. However, only a single network, comprising a feature defined by oscillatory activity below 12 Hz, was associated with a generator in the opercular region of the right inferior frontal cortex and showed the expected associations with behavioral inhibition performance. This study pioneers by providing a detailed ranking of neural features regarding their information content for stop and go differentiation at the single-trial level, and may

René J. Huster and Signe Schneider contributed equally to this work.

The authors declare no competing financial interests.
Contract grant sponsor: German Research Foundation; Contract grant number: HU 1729/2-1 (awarded to RJH and CSH)

*Correspondence to: Rene J. Huster, PhD, Department of Psychology, University of Oslo, Norway. E-mail: rene.huster@psykologi.uio.no

Received for publication 19 April 2016; Revised 30 September 2016; Accepted 25 October 2016.
DOI: 10.1002/hbm.23457
Published online 11 November 2016 in Wiley Online Library (wileyonlinelibrary.com).

further be the first to identify a scalp EEG marker of the inhibitory control network. This analysis allows for the characterization of the temporal dynamics of response inhibition by matching electrophysiological phenomena to cortical generators and behavioral inhibition performance. *Hum Brain Mapp 38:1333–1346, 2017.*

Key words: inhibition; stop signal task; N200; P300; delta; theta; prediction; classification

INTRODUCTION

The ability to stop behavior or cognitive processes is considered crucial for the adaptation to a fast changing environment [Aron, 2007, 2011]. The stop-signal task is a common method to study this ability: participants have to respond to a go signal yet have to withhold their response to a less often occurring stop signal that might follow a go signal. This task enables the computation of the stop-signal reaction time [SSRT; Band et al. 2003], which estimates the speed of the inhibitory process. Although the SSRT varies across paradigms and subject populations, it usually seems to fall within the range of 200–300 ms relative to stop stimulus onset. Then again, studies using transcranial magnetic stimulation of the motor cortex suggest the onset of inhibitory effects at about 150 ms post presentation of the stop stimulus [Badry et al., 2009; Coxon, 2006; van den Wildenberg et al., 2010], leaving us with a critical period of about 50–150 ms during which the exertion of successful inhibitory control is to be expected.

Electroencephalography (EEG) is a common tool used to investigate response inhibition. Here, the N200 and P300 event-related potentials (ERPs), both of which are more pronounced in successful stop as opposed to go trials, are probably the most extensively studied phenomena [e.g., Huster et al., 2013]. When subjected to time-frequency transforms, these ERPs roughly correspond to activity in the theta and delta frequency bands, respectively [e.g., Huster et al., 2013; Schmiedt-Fehr and Basar-Eroglu, 2011]. More recent work also investigated the relevance of activity in other frequency bands for behavioral stopping. Swann et al. [2009, 2011], utilizing intracranial recordings, reported increased beta band activity and coherence with successful stop trials in the right inferior frontal gyrus (IFG) and pre-supplementary motor areas (preSMA). Similarly, Alegre et al. [2008] found increases in beta-band activity with successful stopping over frontal electrode sites. At the same time, Krämer et al. [2011] and Lavallee et al. [2014] reported beta band activity with local maxima over the motor regions. Lavallee et al. [2014] found that this beta activity was differently modulated in bimanual as opposed to unimanual stopping. Activity in the same frequency band also correlated highly with response times in go-trials. Thus, it is unclear whether these different reports on beta band activity refer to the same neural region or system, or whether they reflect dissociable neural phenomena. This issue is further complicated by the fact that there might not always be a direct correspondence between intracranial and scalp EEG recordings, and in fact a direct mapping of IFG and preSMA activity to distinct EEG events is still illusive [Huster et al., 2013]. At last, a few studies also reported stop/no-go versus go differences in the alpha band [Bengson et al., 2012; Huster et al., 2014; Tallet et al., 2009], but it is unclear if these alpha-effects are at all associated with response inhibition, or whether they merely reflect the processing of the task stimuli and potential differences thereof.

Studies linking behavioral performance measures (such as SSRT and accuracy) to EEG phenomena would be helpful to expose potential EEG fingerprints of inhibitory control, yet current evidence is sparse. Some studies reported correlations of SSRTs with features associated with the P300 though. Recently, we found that activity of an independent component that also explained most of the variance of the P300, and correspondingly exhibited elevated activity in the delta frequency range, showed significant correlations with the SSRT at around 200 ms post presentation of the stop stimulus [Huster et al., 2014]. This finding was conceptually replicated by Wessel and Aron [2015], who found that the latency of the P300 onset, with an average of 225.3 ms, correlated positively with the SSRT. It still needs to be determined though whether any of these P3-related effects directly reflect inhibition itself or rather its outcome. The late occurrence of this effect rather suggests the latter, since myographic recordings at effector muscles indicate inhibitory influences as early as 150 ms stop stimulus presentation [e.g., van den Wildenberg et al., 2010]. Data from our lab further revealed the behavioral relevance of other neural processes such as parieto-occipital alpha depletion at about 150 ms post stimulus presentation, which correlated with individual SSRTs [Huster et al., 2014]. Given its spatial, temporal, and spectral features, this alpha-related effect points to the relevance of early visual processing for successful inhibition, yet seems as well rather unlikely to directly reflect inhibition proper.

Hence, an unequivocal electrophysiological indicator of a genuine inhibition process still needs to be identified. Also, this quest has been hindered by the strong focus on rather few electrophysiological phenomena, such as the N200 and P300 or narrowly defined effects in selected frequency bands. Potentially reflecting this limited focus, EEG activity patterns that could with high reliability be traced back to the right IFG or preSMA are not known yet [Huster et al., 2013]. Altogether, this leaves us with a severely limited understanding of the entirety of processes driving successful inhibition, with many electrophysiological features

potentially undetected that may be of relevance for successful inhibition or response adaptation.

This study set out to characterize electrophysiological features associated with response inhibition tasks in their entirety. To this means we developed a pipeline of consecutive analysis steps where we (i) computed time-frequency representations of the single-trial EEG data, (ii) estimated coherent functional networks and their spatio-spectral activity patterns, (iii) identified features reflecting differences in network activity between go and stop trials in specific frequency bands and time frames, and (iv) applied a single-trial classification to identify and rank these features in accordance with their contribution to the classifier's performance. We furthermore computed correlations of neural feature values with behavioral measures to highlight features of relevance for neural and behavioral stopping performance. This procedure allowed us to assess not only evoked activity, as reflected in ERPs, but also induced activity, that is, neural activity not strictly time-locked to certain events. Further, we would disentangle activity patterns otherwise showing strong spatio-temporal overlap, thereby dissociating multifaceted phenomena as seen in the beta band, presumably comprising oscillatory activity of at least inferior frontal and motor regions. Although our procedure was developed to specifically discover EEG characteristics that would remain undetected with conventional analyses, we of course also expected to detect networks already implicated in response inhibition tasks. As such, it was expected that fronto-central delta and theta activity would excel as strong predictors of the trial type, but also that the relevance of parieto-occipital activity could be replicated. We further expected an dissociation of beta band activity according to frontal and rather motoric patterns, and also that at least one of the corresponding networks would be associated with activity of the inferior frontal and pre-supplementary cortices.

METHODS

Participants

Seventeen healthy volunteers participated in the study. Datasets of three participants were excluded from further analysis, due to an insufficient number of artifact-free trials in at least one condition. The final sample consisted of 14 healthy participants (7 female; mean age = 24.83 ± 2.83 years). All participants were right-handed and had normal or corrected-to-normal vision. Volunteers participated in return for payment and gave written informed consent prior to study participation. The study was approved by the ethics committee of the University of Oldenburg.

Stimuli, Task and Behavioral Data Analysis

Participants performed a modified stop signal task while seated in an electrically shielded and sound attenuated cabin. Stimuli were presented on a 24″ TFT screen placed behind an electrically shielded window. Viewing distance to the screen was 100 cm such that the visual angle of stimuli was about 1.7°. Task presentation and performance was controlled using the Presentation software (version 14.08, Neurobehavioral Systems). Responses were given via a dedicated button box.

Stimuli consisted of centrally presented arrows pointing to the left or the right. On a go trial, the target arrow was displayed in purple and subjects had to indicate the direction the arrow was pointing via button press. On stop trials, the color of a given target would change from purple to orange, instructing the participants to withdraw from responding. Stimulus-color assignments were counterbalanced across subjects. The delay between the onset of target presentation and color change (stop signal delay, SSD) was tracked and adapted by an algorithm according to a staircase procedure, used to achieve a response rate of approximately 50%. SSD values were adapted as increments or decrements in steps of 50 ms depending on whether a subject's response on the last stop trial was correct (successful inhibition) or not (failed inhibition), respectively [Band et al., 2003]. SSD adaptations were tracked separately for left- and right-hand responses. The presentation time for stimuli was the same for go and stop stimuli (with respect to the relevant hand). In stop trials, go stimuli were presented for the length of the SSD. Stop stimuli immediately followed the go stimuli and were presented for 216 ms. Total presentation time of go stimuli in go trials was determined by the SSD plus 216 ms. Whenever none of the previously described stimuli was presented, a fixation cross was shown in the middle of the screen instead.

Trials were grouped in blocks of 30 with an average trial duration of 1,360 ms. Each block contained 20 go and 10 stop trials and feedback on the participant's performance was given after each block, instructing subjects to respond faster or more accurately when detecting average response times larger than 550 ms or a rate of failed stops of more than 50% with stop signal trials, respectively. Participants completed a total of 30 blocks with self-paced breaks after every fifth block.

To assess the behavioral performance of the participants, the mean reaction time (RT) to go stimuli, the percentage of failed stop trials (unsuccessful inhibition), as well as the stop signal reaction time (SSRT) were calculated. The SSRT was computed by subtracting the average stop signal delay (SSD), that is, the time between the presentation of the go and the stop stimuli in stop trials, from the mean RT on valid go-trials [$SSRT_{mean}$ according to Band et al., 2003]; thus, shorter SSRTs indicate faster or more efficient inhibitory processing. Since no statistically significant differences between response hands were observed, behavioral data were averaged across response hands.

EEG: Recording and Preprocessing

EEG data from 31 channels were recorded using a BrainAmp (Brainproducts, München, Germany). Electrodes were mounted in a flexible lycra-electrocap according to

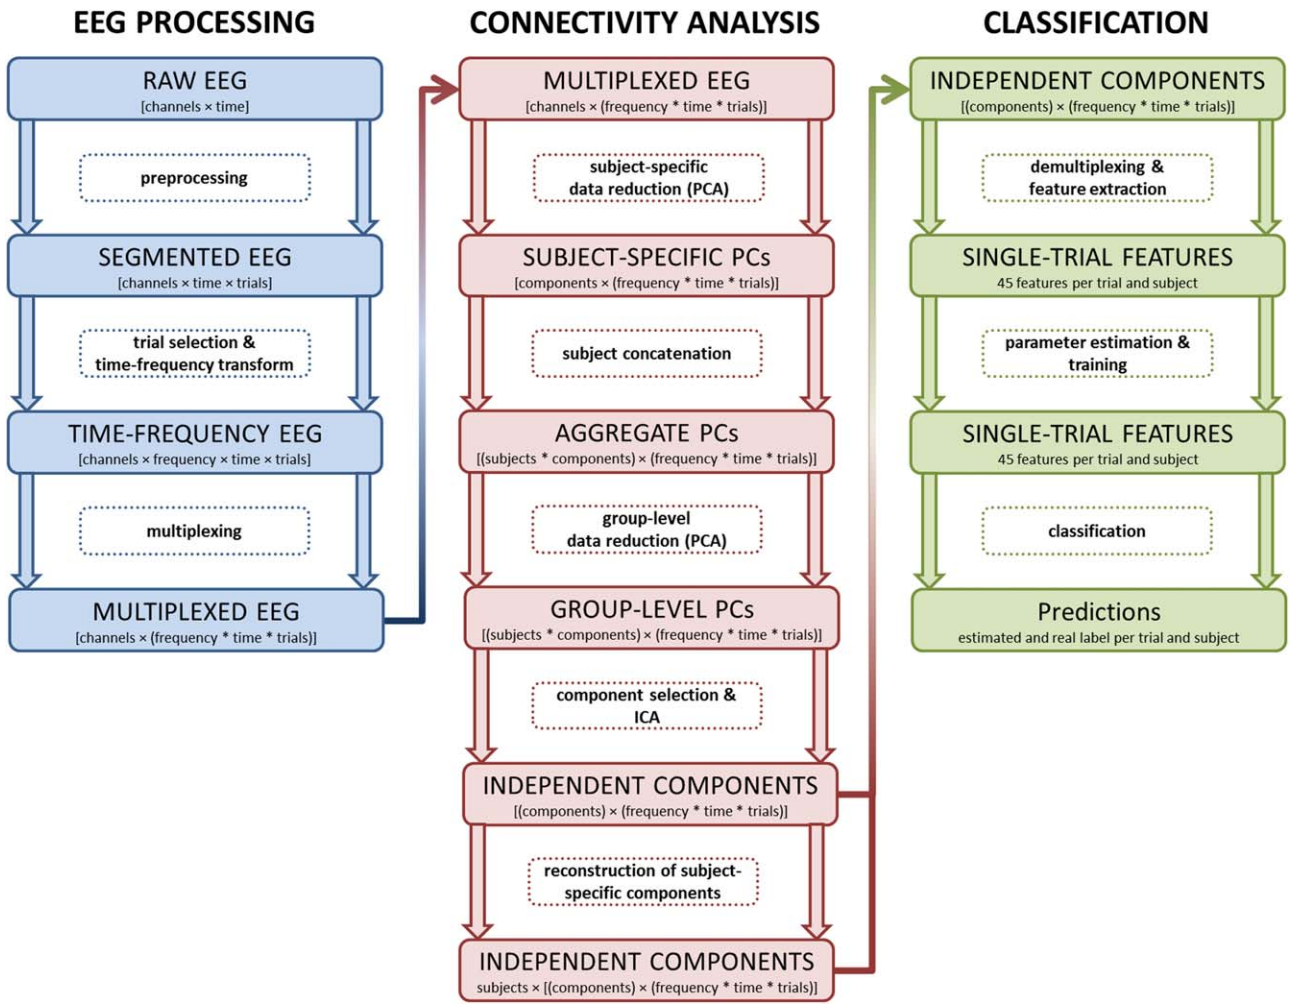

Figure 1.

Schematic of the processing workflow. This figure depicts the main stages of the EEG processing pipeline divided into three major stages: EEG preprocessing, functional connectivity analysis using group-level data decomposition, and the single-trial classification procedure. Data formats and their organization are highlighted in large rectangles, with the transformatory processing steps in-between. [Color figure can be viewed at wileyonlinelibrary.com]

the 10–10 system for electrode placement [easycap, Falk Minow Services, Munich, Germany; Chatrian et al., 1985]. In addition, the electrooculogram (EOG) was recorded from one channel placed below the left eye. EEG and EOG recordings were taken continuously from 0.01 up to 1,000 Hz at a sampling rate of 2,500 Hz. Impedances were kept below 5 kΩ and matched for homologous sites with a maximum deviation of 500 Ω. All data were recorded against a reference-electrode positioned on the tip of the nose. A ground electrode was placed on the forehead. Off-line, data were filtered from 0.1 to 75 Hz and re-sampled to 500 Hz. Ocular, muscular, as well as technical artifacts were rejected by computing an ICA for every subject's data, rejecting those components representing artifactual sources and back-projecting the thereby cleaned data. Correct go and successful stop trials were extracted, aligned to go stimuli with go trials and to stop stimuli with stop trials, from 1.5 seconds pre- to 1.5 seconds post-stimulus presentation. To further control for artifacts, trials with amplitudes higher than 80 μV were automatically rejected. The EEGLAB software package was used for pre-processing of the data [version 9.0.8; Delorme and Makeig, 2004]. Please refer to Figure 1 for an overview over the whole processing pipeline.

EEG: Trial Selection, Time-Frequency Transform, and Multiplexing

After preprocessing and segmentation of the EEG, an adjusted number of go and stop trials was randomly

selected from each subject's available data, since the number of trials per condition has to be same for each subject in order to compute a group ICA. This is, because the data sets of single subjects later get vertically concatenated prior to a group-wise data reduction step, meaning that these data sets get stacked in the dimension containing electrodes/components [for details see Eichele et al., 2011; Huster et al., 2015]. To this means, a randomized selection of 176 trials from each subject (44 go and 44 stop trials for each the left and right hand conditions) was drawn and it was ensured that the resulting averages were representative for the neural responses prior to trial selection. That is, the go- and stop-ERPs derived after trial selection did not statistically differ from those ERPs computed from all trials; this was assessed by means of point-wise two-tailed t-tests while utilizing a false discovery rate (FDR) correction for multiple comparisons.

For each of these trials a time-frequency decomposition using Morlet wavelet convolution was computed. A frequency-range from to 30 Hz was covered in 100 frequency steps. Increasing numbers of cycles were used, starting with 1 cycle for the lowest frequency up to a maximum of 15 cycles for 30 Hz. Then, frequency-specific log-transformed power values for each time bin were divided by the mean power of a baseline from −1,000 to −100 ms, resulting in the so-called event-related spectral perturbation [ERSP, in dB; Delorme et al., 2004]. Time-frequency transforms were computed using EEGLAB's newtimef-function.

For each subject the data were restructured to a two-dimensional matrix with rows corresponding to channels and columns corresponding to the multiplexed time-frequency bins. Multiplexing (or muxing) refers to the mixing of several signals into a single or lower number of signals. Since the data matrices after time-frequency transform are 4-dimensional ([channels × frequency × time × trials]), yet standard ICA and PCA work on 2-dimensional matrices. This is perfectly valid since PCA and ICA do not rely on any specific ordering of samples along the main dimension for data decomposition (in EEG usually time). The resulting multiplexed EEG matrix per subject was of size [channels × (frequency * time * trials)].

Functional Connectivity Analysis: Concept

To identify functionally coherent networks from the EEG data, a group-level independent component analysis (ICA) was set up. ICA is a powerful machine learning algorithm that blindly decomposes EEG data into statistically independent sources [Hyvärinen and Oja, 2000]; as such, it is a commonly used method for the estimation functional connectivity in functional magnetic resonance imaging and EEG [e.g., Bastos and Schoffelen, 2016; Joel et al., 2011]. In context of ICA, connectivity is based on the estimation of statistical dependencies across electrodes caused by the common activity of a source, that is, a single cortical area, a network of brain regions, or other generators of physiological activity (e.g., muscles or eye movements). Mostly, ICA is applied on a single subject level but the resulting components and their order are not necessarily similar across subjects, impeding further inferences on a group level. Recently, group ICA approaches have been developed that are able to overcome these limitations by identifying a single set of independent components for all subjects. We applied a self-implemented group ICA approach, adapted after Eichele et al. [2011] and Huster et al. [2015]. In short, group ICA as used here relies on a combination of two steps for data reduction, one at the level of individual subjects and a second applied to an aggregate data matrix built from the output of the previous subject-specific reduction. Lastly, ICA is applied to the output of the second or group-level data reduction. Resulting group independent components correspond to statistically independent sources representative for the sample as a whole.

Functional Connectivity Analysis: Whitening and Data Reduction

Whitening (i.e., zero-meaning, variance homogenization, decorrelation) and data reduction are common preprocessing steps prior to the actual application of ICA to reduce the complexity of the problem and render its solution for large amounts of data feasible. The combination of PCA and z-standardization of the input data are procedures commonly applied to achieve this preprocessing. A PCA computes decorrelated components from the input data and inherently ranks them in descending order according to their variance explanation. Data reduction usually then corresponds to the extraction of the first x components, with x being lower than the number of input channels. The number of selected principal components can, for example, be based on the decision to account for a total of 80% or 90% of the variance. A channel-wise z-transformation, that is, a unification of the scaling across channels, prior to PCA not only completes the whitening procedure, but may further help to retain smaller effects that otherwise may not contribute much to the overall variance (such as small amplitude fluctuations in higher frequencies) and may otherwise get lost during data reduction. As mentioned earlier, data reduction in context of group ICA is conducted consecutively both at the level of single subjects as well as the group-level.

First, the single-subject data sets containing the multiplexed single-trial and time-frequency transformed EEG separately undergo channel-wise z-transformation and data reduction via PCA, leaving us with a data matrix of size [components × (frequency * time * trials)] for every subject. The first 15 components were extracted from each of these data sets for further processing because on average the first 15 principal components (PCs) explained about 90% of the total variance.

Then, the selected principal components of all subjects got concatenated vertically to build an aggregate data matrix containing the extracted PCs; the resulting matrix thus generally is of size [(subjects * components) × (frequency * time * trials)]. This necessitates that the subject-specific data sets are of the same size, that is, they contain the same number of trials, and that the order of trials with respect to the conditions they belong to is consistent across data sets as well. A second PCA is computed on this concatenated data matrix after a component-wise z-transform, which provides us with group-level PCs. Of these group-level PCs, again the first 15 orthogonal components were extracted and z-standardized.

Functional Connectivity Analysis: ICA, Single-Subject Reconstruction

Ultimately, the selected group-level PCs were subjected to extended infomax ICA as implemented in EEGLAB [Lee et al., 1999]. This resulted in 15 group-level independent components (ICs) representing time-frequency patterns that capture activations coherently expressed across subjects.

The single-subject activations corresponding to these 15 ICs were reconstructed by matrix multiplication of the weight matrices that resulted from the two PCAs and the ICA, and the original multiplexed time-frequency-transformed single-trial data of each subject. Note that this is easily possible since the first PCA coefficient matrix, coding for the transformation from EEG to PCs, is subject-specific, and the second group-level PCA coefficient matrix can be partitioned into subject-specific divisions coding the transformation of the subject-specific PCs to the group-level PCs. The ICA coefficient matrix, coding the transform of group-level PCs to ICs, naturally is invariant across subjects. These matrices and their inverses further define the scalp mapping of ICs, that is, their contributions to the electrodes in every subject. Further details can be found in Eichele et al. (2011) or Huster et al. (2015).

Source Analysis of IC Topographies

Brain sources were computed for those group ICs that stood out as containing highly predictive features, that is, features with high absolute weights as determined in the final training used for model estimation (see description of the classification procedure below). The scalp maps of these ICs were computed by averaging the individual IC scalp topographies across subjects and subjecting these averaged IC topographies to inverse modeling. Source reconstructions were estimated according to the Bayesian approach for distributed source modeling as implemented in SPM12 (www.fil.ion.ucl.ac.uk/spm). Head compartment meshes for the scalp as well as the outer and inner scull were generated using the template T1-image provided with SPM12. Electrode locations were defined and integrated in the head model based on a standard label matching procedure. A forward model was set up based on the coregistration of the head compartment meshes and three spherical shells. The source space was confined to a segmented representation of the cortex with 8,196 vertices computed from the T1-weighted template image. Sources were reconstructed for a time window of 20 ms centered at a component's peak activity, with source orientations being normal to the cortical surface. The statistical maps containing the normalized Maximum Intensity Projection (MIP), that is, the spatial source activity profiles, were used to determine the local maxima for selected ICs. Further details on the procedures for source reconstruction as implemented in SPM can, for example, be found in Litvak et al. [2011].

Feature Generation and Single-Trial Classification

To generate and select meaningful features for single-trial classification, clusters of time-frequency bins within each group IC were identified that differed significantly between stop and go trials. First, IC data were averaged across trials for each subject and paired *t*-tests comparing stop- and go-related activity were calculated for each time-frequency bin, that is, every combination of time and frequency. An FDR correction was applied on the resulting *P*-values to correct for multiple comparisons. Then, clusters concurrently spanning larger ranges in time and frequency were defined as groups of adjacent time-frequency bins with significant differences between conditions; that is, all significant bins that could be connected through a path uninterrupted by a non-significant bin would belong to the same cluster (e.g., IC 2 in Figure 2 shows three clusters; the clusters are separated by non-significant time-frequency bins that are masked green). In addition, a cluster had to encompass a minimum of 100 adjacent and significant bins. A total of 45 clusters were extracted according to this procedure, and each cluster would correspond to a feature for the single-trial classification.

To derive the feature vectors for every single trial, a mask was computed from each cluster and the mean activity (i.e., the mean event-related spectral perturbation) within each of these clusters was extracted from all trials and subjects. Thus, every trial was characterized by a vector (the feature vector) containing 45 values (thus 45 features), each corresponding to the mean activity within a cluster. In addition, the values of each feature were linearly scaled to the range $[-1, +1]$ to correct for amplitude differences across frequencies.

A Support Vector Machine (SVM) was used for the classification of go and stop trials based on the trials' feature vectors and their labels (either go or stop). SVM classification is a method that constructs a hyperplane in a multidimensional space that maximizes the margin between classes of data. Here, C (or type 1) Support Vector

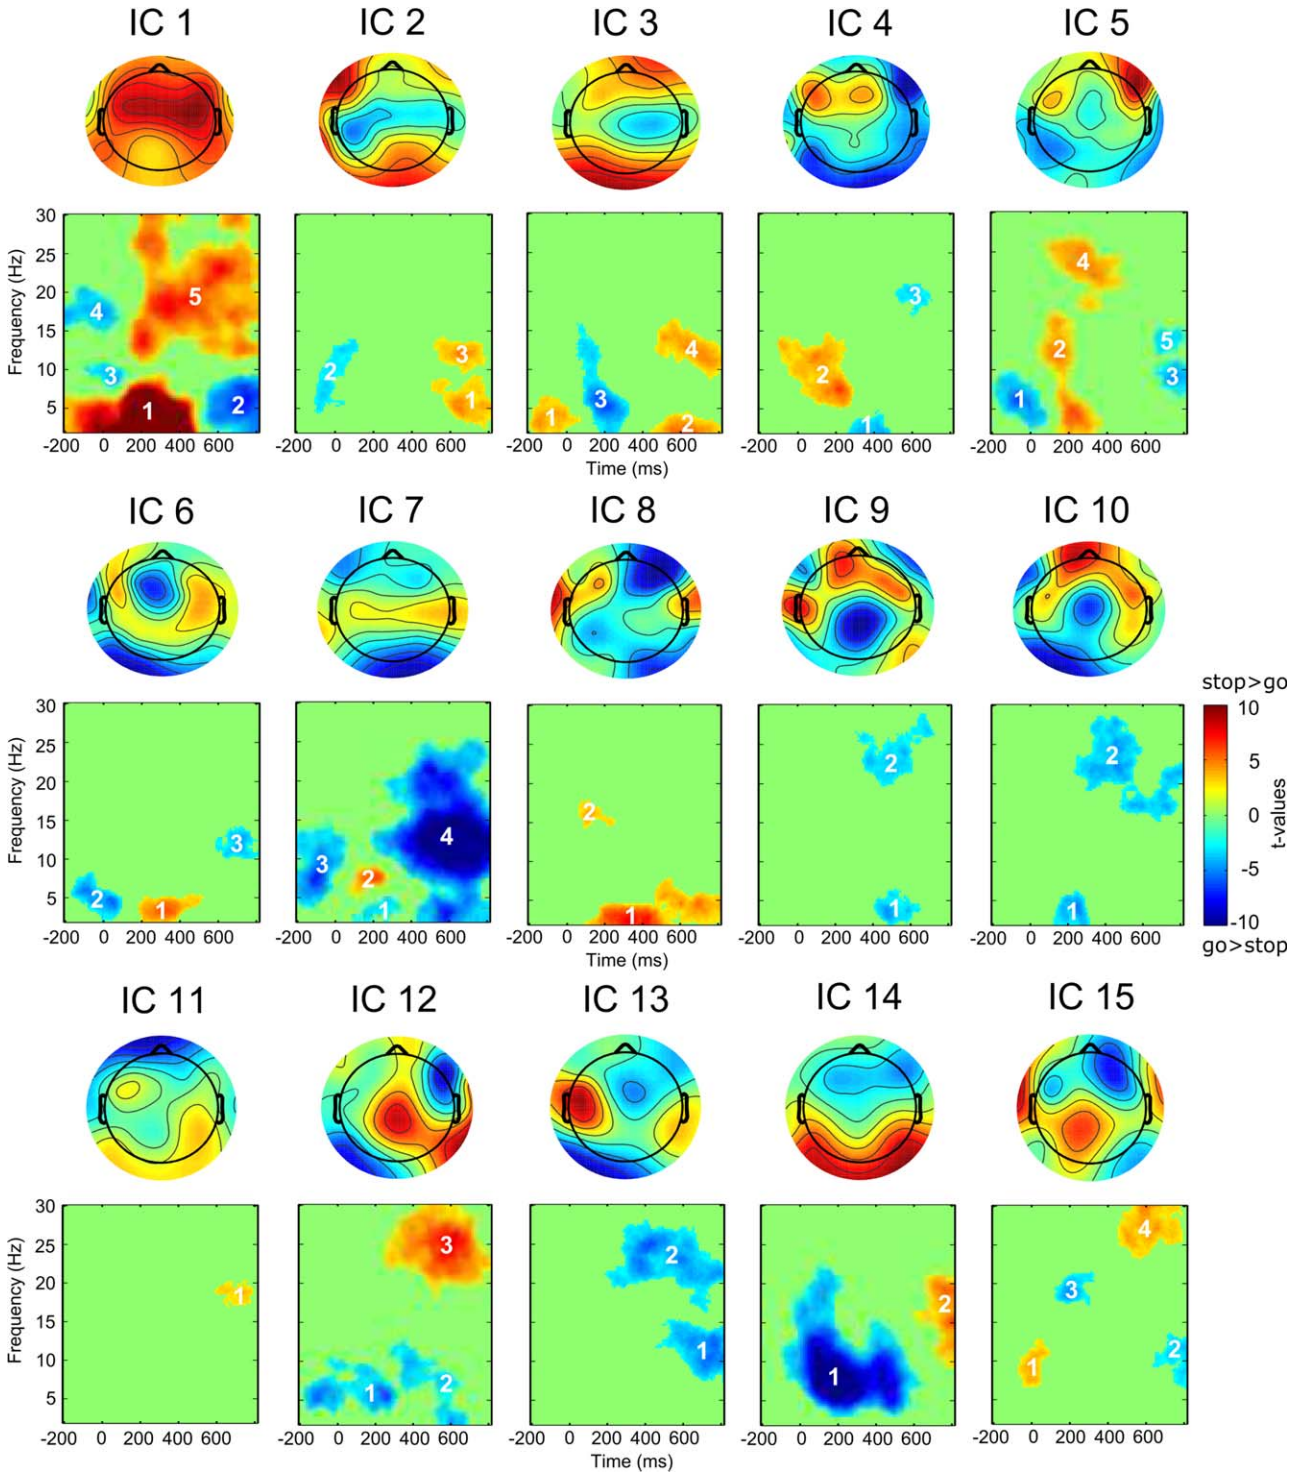

Figure 2.

Group independent components and significant clusters for the stop-go difference. Features can be identified by the combination of the component's and the cluster's number (e.g., 1–5 = first component, cluster number five). [Color figure can be viewed at wileyonlinelibrary.com]

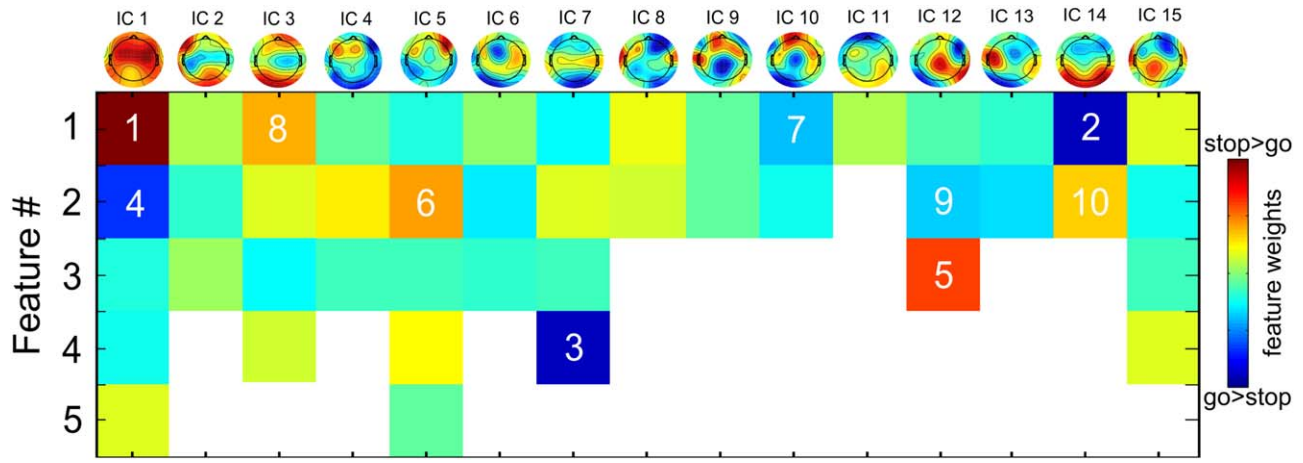

Figure 3.
Color-coded feature weights with colors also indicating the direction of the stop-go difference. The ranks of the 10 best predictors are added on top of the color-coding (indicating the best predictor). [Color figure can be viewed at wileyonlinelibrary.com]

Classification (C SVC) with a linear kernel was used as implemented in the LIBSVM software package for MATLAB interfacing [Version 3.14, Chang and Lin, 2011]. The training and test data for SVC comprised 88 go- and 88 stop-trials from fourteen subjects, that is, 2,464 trials in total, and 45 features for each of these trials along with labels, indicating whether the trial was a go or stop trial.

Initially, a grid search was conducted to estimate the best-fitting value for parameter C, which determines the penalty for classification errors during training. The parameter was fixed to the value providing the best classification accuracy as estimated using a 10-fold cross-validation; that is, the available trials were divided into 10 non-overlapping samples and the elements of each sample were predicted from a classifier trained on the trials of the remaining 9 samples. The parameter showing the best classification accuracy (i.e., the highest percentage of correct classifications) was then used in a leave-one-out cross-validation, in which every trial once gets predicted from the training sample containing all other trials. The final classification accuracies for go and stop trials were derived from this cross-validation procedure. This procedure is well-established for the evaluation of a classifier's performance providing an almost unbiased estimate of the generalization error while being computationally feasible [e.g., Hsu et al., 2010; Theodoridis and Koutroumbas, 2008]. At last, the SVM classifier was trained on all available trials, and from this final linear SVM model a weight vector with one weight for each feature was derived. Comparable to regression weights, this weight vector provides information about the relevance of each feature for classification [Chang and Lin, 2008]. Hence, the resulting feature weights were ranked in descending order of their absolute size, thereby providing us with a feature list ordered according to the relative relevance of the features.

Post-Hoc Assessment of Brain-Behavior Associations

To further test for the potential behavioral relevance of the extracted features, averaged activity of the ten best predicting features across trials were computed for every subject. Go- and stop-related mean feature values were then correlated with go-RT and SSRT, respectively, and resulting p-values were corrected for multiple comparisons using the false discovery rate procedure.

RESULTS

Behavioral Data

Participants demonstrated a mean RT in go trials of 594.4 ms (std = 152.2) and committed errors in only 3.7% (std = 3.9) of all go trials. In 46% of the stop trials (std = 10) participants failed to successfully inhibit their responses. The average SSRT was estimated to be 220.2 ms (std = 51.1). Hence, behavioral performance measures were well within the range of values expected with this kind of stop signal paradigm.

Functional Connectivity and Feature Extraction

Each of the 15 components estimated from the multi-subject EEG data is characterized by its topography and time-frequency profile. Figure 2 depicts independent component topographies and the outcome of the contrast of the stop versus go comparison. Based on this contrast, 45 features were derived, each constituting a cluster of adjacent and connected time-frequency bins significantly differing between the two conditions.

TABLE I. Local maxima of independent component source reconstructions

IC	X	Y	Z	Region	BA
1	−16	−54	68	Left precuneus	5
	18	−56	68	Right superior parietal	5
	−58	−3	6	Left superior temporal	48
	51	−11	14	Right rolandic operculum	48
	−49	−50	−25	Left inferior temporal	37
	−48	3	−35	Left inferior temporal	20
	43	3	−37	Right inferior temporal	20
3	−30	28	46	Left middle frontal	9
	39	46	11	Right middle frontal	45
	51	36	4	Right inferior frontal (triangularis)	45
	63	−11	−22	Right middle temporal	21
	−58	−14	−23	Left middle temporal	21
5	51	38	2	Right inferior frontal (triangularis)	45
	52	−12	−36	Right inferior temporal	20
7	32	28	41	Right middle frontal	9
	43	−14	35	Right postcentral	3
	−42	−14	34	Left postcentral	3
	−20	−95	5	Left middle occipital	18
	23	−92	4	Right calcarine	18
	−49	34	−4	Left inferior frontal (orbitalis)	47
	51	36	4	Right inferior frontal (triangularis)	45
10	−16	58	20	Left superior frontal	10
	−73	−16	22	Left postcentral	48
	−50	32	0	Left inferior frontal (triangularis)	47
12	13	−25	68	Right precentral	4
	53	13	12	Right inferior frontal (opercularis)	48
14	−40	−26	37	Left inferior parietal	3
	51	36	4	Right inferior frontal (triangularis)	45
	43	−23	38	Right postcentral	3
	−61	−29	8	Left middle temporal	22
	60	−35	7	Right middle temporal	22
	−39	−88	−4	Left middle occipital	19
	41	−66	0	Right middle temporal	37

Single-Trial Classification

The single-trial prediction based on a support vector machine with a linear kernel received an average accuracy of 84.7%. Of the go trials, 85.55% were correctly identified, whereas 14.45% were wrongly classified as stop trials. Of the stop trials, 83.85% were correctly assigned a stop label, whereas 16.15% were assigned the incorrect go label. An overview of all features with their color-coded weights can be found in Figure 3, which also denotes the 10 best predicting features as ranked by their absolute weights. These ten features contributing highest to single-trial classification will shortly be highlighted at this point, by providing concise descriptions of spatial, spectral, and temporal characteristics; condition-specific local maxima and minima will be mentioned only if they deviate from the go versus stop contrast as depicted in Figure 2. Table I lists the local maxima of source reconstructions for those seven components that provided the 10 strongest predictors, and Figure 4 depicts the source reconstructions for five components with focus on the inferior frontal region. At last, Figure 5 provides an overview of the 10 best predictors, ordered by their temporal evolution and spectral content.

#1: IC 1 contributed two features that were ranked as first and fourth best predictors. Feature 1 was ranked first and captured much of the power in delta and theta frequency bands that was stronger in stop than go trials with a stop-related maximum at 273 ms at 3.4 Hz. Based on its topography and spectral composition, it is likely that this feature captures much of the fronto-central N200 and P300 ERPs. In accordance with an association with the P300, this component seems to be generated in a widely distributed network of parietal and temporal regions. Its activity in the go condition correlated negatively with response times ($r = -0.57$).

#2: IC 14 feature 1 was characterized by a marked parieto-occipital topography, and correspondingly showed the expected effects in the lower frequency ranges including alpha and theta frequencies. Activity of this feature was lower in stop trials, with a maximum difference at 212 ms and 7 Hz. The local minimum in the stop condition was found at 8.5 Hz and 432 ms though.

#3: IC 7 feature 4 was characterized by more negative activity in stop than in go trials, and a local maximum in stop trials at 12.7 Hz at 700 ms. Driven by a widely distributed network of regions including inferior and middle frontal, postcentral, as well as occipital regions, its prominent late activity difference between conditions may indicate an adaptive function for attentional and behavioral adaptation.

#4: IC 1 also contributed its feature 2, which showed lower activity in the delta and theta band in stop as compared with go trials after 600 ms, thereby possibly corresponding to some rebound in the delta/theta frequency range causing this inversion relative to feature 1 of the same component.

#5: Feature 3 of IC 12, characterized by increased beta band activity in stop relative to go trials with a maximum difference at 25.2 Hz and at 602 ms, was scored fifth best predictor. Sources were located in the right precentral region and the opercular area of the inferior frontal cortex.

#6: Feature 2 of IC 5 not only was a high ranking predictor, but its go-related activity also correlated significantly with reaction times ($r = 0.86$). It seems that differences between go and stop conditions of this feature were driven by a stronger depletion of go-related activity in the higher theta and lower alpha band at around 273 ms, with a local minimum at 224 ms and 5.6 Hz. Correspondingly, more negative going activity in the go condition was associated with faster responses. Dominant sources for this component were found in the right inferior frontal (triangularis) area and temporal regions.

#7: IC 10 feature 1 was driven by lower activity in the delta frequency band in go as compared with stop trials between 200 and 300 ms post stop-stimulus presentation. The topography, which showed a central to frontal transition with a shift from negative to positive potential

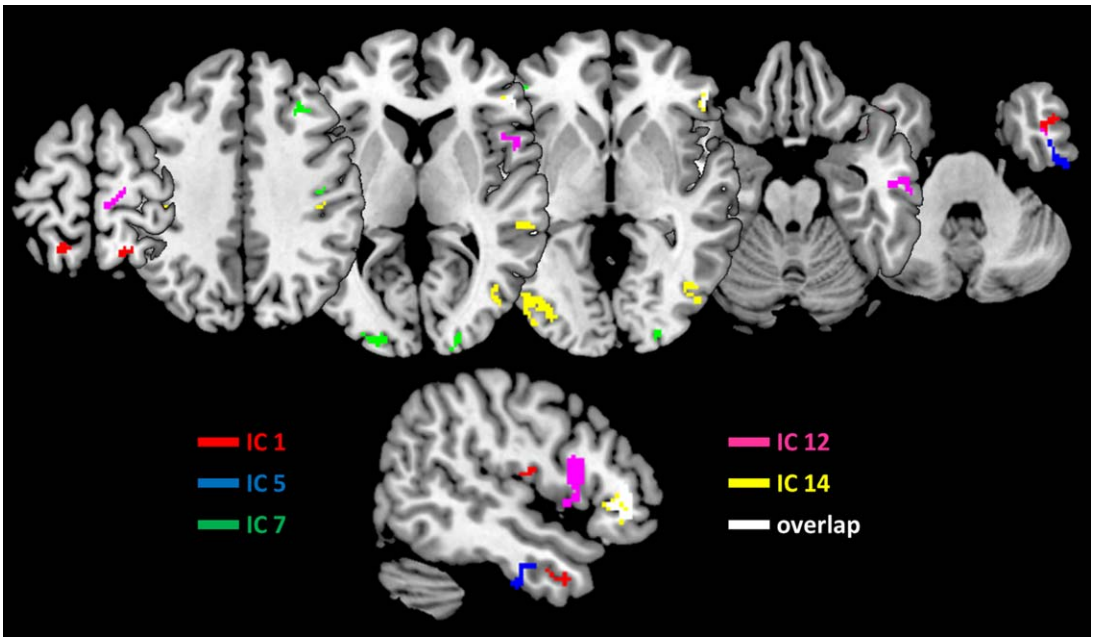

Figure 4.

Source reconstructions for major independent components. Shown are the upper 25% of the MIPs for each of the five components, as well as their overlap. The lower row depicts a sagittal cut through the right inferior frontal region. Interestingly, only IC 12 showed an association with the opercular region, whereas all other components but IC1 at least partially originated from triangular and orbital regions of the inferior frontal cortex. [Color figure can be viewed at wileyonlinelibrary.com]

amplitudes, reflects the activity of a left-hemispheric network with generators in superior and inferior frontal as well as postcentral cortices.

#8: IC 3 feature 1 also was driven by delta band activity with differences between conditions occurring already before presentation of the stop stimulus. Here, however, stop-related activity was more pronounced than that of go trials. Generators were estimated to reside in bilateral middle frontal and temporal, as well as the right inferior frontal cortex.

#9: Stop-related activity of feature 2 of IC 12 correlated with SSRTs ($r = 0.57$). This feature was characterized by lower activity in stop trials predominantly spanning a frequency range from lower alpha to the delta band, and a temporal window between 200 and 650 ms. Its local maxima were at 9.8 and 2.5 Hz at 371 and 602 ms, respectively.

#10: IC 14 feature 2 represented a late increase in lower beta band activity in stop trials, with a peak difference at 19.2 Hz and 761 ms. Main generators of this component were found in bilateral temporal and occipital regions, as well as the right triangularis area of the inferior frontal gyrus.

Brain-Behavior Associations

Testing for correlations between the condition-specific activity of the 10 best ranking features on the one, and the mean go RTs and SSRTs on the other hand, revealed some significant correlations that survived the correction for multiple testing (see Fig. 6). Go-related activity of feature 1 of IC 1 and feature 2 of IC 5 correlated with go RTs at $r = -0.57$ ($P < 0.05$) and $r = 0.86$ ($P < 0.001$), respectively. Feature 2 of IC 12 showed correlations between stop-related activity and the SSRT at $r = 0.57$ ($P < 0.05$). Correlations between SSRTs and fronto-central activity in the delta and theta frequency range, which would have corresponded to feature 1 of IC 1, expected to occur at around 200 ms post stop-signal presentation, were not found.

DISCUSSION

Using a combination of functional connectivity estimation and single trial prediction on EEG data of a stop signal task, this study not only validated already known features contributing to major differences between go- and stop-related activity, but furthermore identified a number of additional electrophysiological phenomena, some of which were also associated with behavioral performance measures. A major limitation of current approaches to the study of behavioral inhibition is the focus on a very limited number of EEG phenomena and regional activations, such as the N200 and P300 ERPs [e.g., Huster et al., 2013], or the right inferior frontal cortex [Aron, 2007]. Since a full

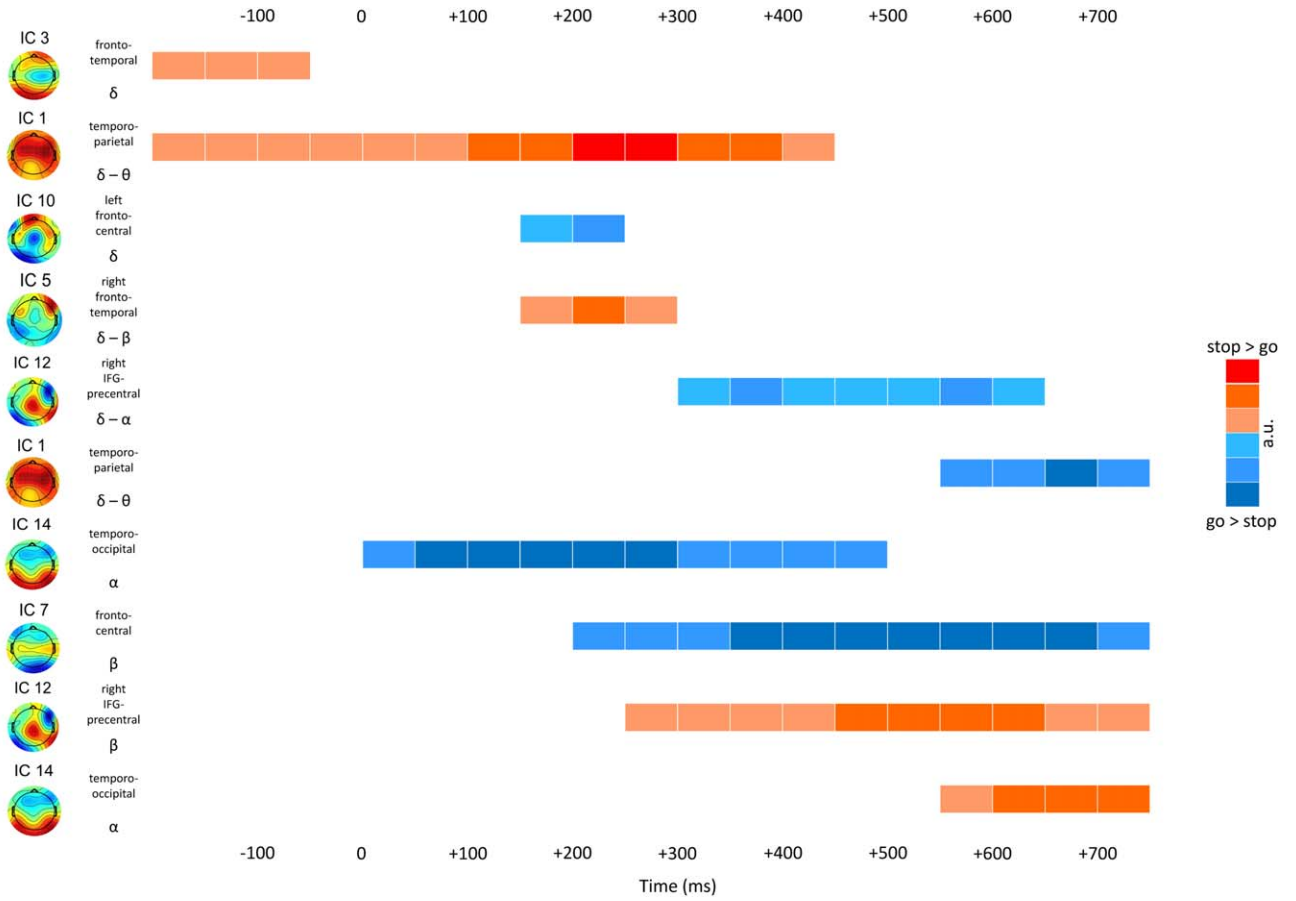

Figure 5.

Schematic of the temporal progression of network activities. Each row corresponds to one of the ten highest ranking features as indicated by the single-trial classification. Each feature is characterized by its component topography and underlying network, the feature's dominant frequency band, as well as a simplified schematic indicating the direction and strengths of the stop versus go difference. [Color figure can be viewed at wileyonlinelibrary.com]

understanding of the neurocognition of behavioral inhibition necessarily relies on the understanding of associations between neural and cognitive mechanisms driving performance in relevant task contexts, our goal has to be the identification of the entirety of processes seen in inhibition tasks. Thus, this work is the first to widen our perspective by uncovering neural processes supporting motor inhibition that so far have been overlooked.

By using group ICA on time-frequency transformed data we studied not only evoked but also induced electrophysiological responses. We extracted 15 functional networks of which seven contributed the ten best predicting features differentiating between go- and stop related brain responses. Figure 5 provides an overview of activity profiles of these features, ordered by their temporal and spectral progression. Fronto-temporal and temporo-parietal networks, associated with low-frequency bands, showed early activity differences already before stop-stimulus onset, and later temporally overlapped with further low-frequency activity features between 200 and 300 ms. A sustained depression of temporo-occipital alpha activity coincided with stop-stimulus presentation; beta-band activity did not start to differ between conditions before 150–200 ms post stimulus presentation. This overview indicates that early differences may simply be driven by variations in preparation, stimulus processing and evaluation, whereas time frames from 150 ms onward show an abundance of diverging processes, including the activity of a network putatively associated with the right operculari of the IFG as well as the precentral region. Altogether, network profiles indicate strong overlap in terms of spatial (topographies), spectral, as well as temporal characteristics, making it unlikely that analyses conducted in electrode space alone, even when confined to specific frequency bands, will ever provide a pure measure of a neurocognitive process reflecting inhibition or its control.

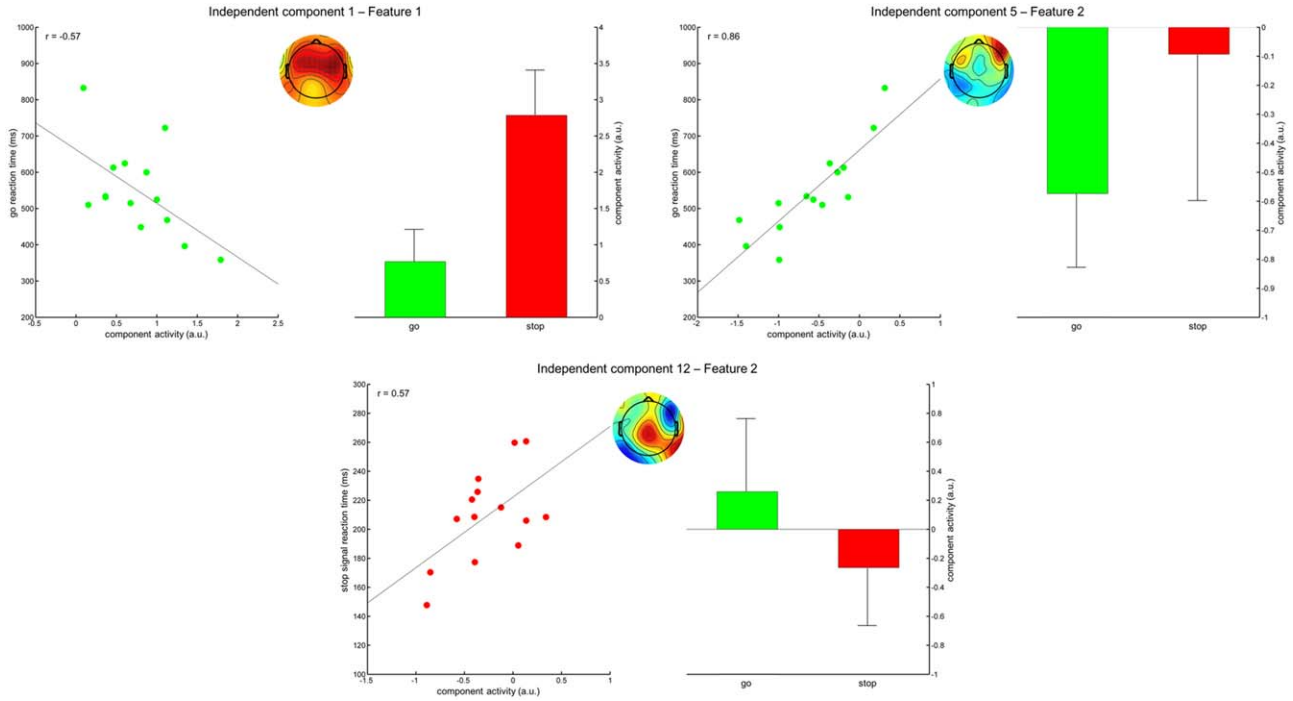

Figure 6.

Correlations of go- and stop-related component activity with go-RTs and SSRTs. Shown are the correlation plots including regression line and the resulting correlation coefficient. In the middle of each section, the topographic map of the corresponding component is depicted. To the right of each section, the mean activities for the go and stop conditions are shown, including their standard deviations. [Color figure can be viewed at wileyonlinelibrary.com]

Inspecting some networks closer, we found that one IC reflected the expected fronto-central activity in the delta and theta frequency range (Feature 1 of IC 1), which was higher in stop trials and originated from a constellation of widely-distributed generators. This component furthermore exhibited negative correlations of go-related activity and mean RTs; higher go-related activity thus was associated with shorter reaction times. In contrast to our expectations though, we could not find the expected correlations with SSRTs reported earlier, where longer SSRTs were associated with higher stop-related component activity around 200–210 ms [Huster et al., 2014], or the P300 onset latency [Wessel and Aron, 2015]. Feature 2 of IC 12, on the other hand, showed a positive correlation with SSRTs, a topography with a centro-parietal maximum that overlaps with the P300, pronounced effects in the lower frequencies, as well as onset of differences between conditions starting at around 200 ms. Thus, because of its spatio-temporal overlap with more centro-parietal activity, it would cause amplitude variations as observed in previous work [Huster et al., 2014]. More interestingly even, this component was the only one revealing a generator in the opercular region of the right inferior frontal gyrus, which is implicated to be the key region for behavioral inhibition and also exhibits correlations with SSRTs [Aron et al., 2014]. Thus, these studies altogether imply a low frequency process, evident at around 210 ms with a centro-parietal local maximum to be correlated with SSRTs, and that the previously reported P3 amplitude or latency variations may be caused by spatio-temporally overlapping activity originating from the opercular region of the right inferior frontal cortex.

Another interesting finding refers to feature 2 of IC 5, which also showed high positive correlations of activity in go-trials with mean response times. Stronger depletion of activity at about 15 Hz and below seen at around 200 ms post go-stimulus presentation, thus leading to a relative increase in stop as compared with go-trials, was associated with faster mean RTs. The same component showed increased activity in the same frequency range in stop trials at stop signal presentation, which led to generation of feature 1 of this component. This suggests that both features might correspond to the same process, that is, a decrease in activity relative to baseline at around 200 ms in go trials, which corresponds to the time frame the stop signal would be expected. The generators of this component were traced back to right triangularis region of the inferior frontal gyrus, as well as the right inferior temporal region. Although the exact function of the right triangular

and orbital regions of the inferior frontal cortex are not yet fully clarified, research from other domains suggests that homologue regions in the left hemisphere serve as hubs for the access of task-relevant information and a post-retrieval selection process to resolve conflict among concurrently active representations [e.g., Badre and Wagner 2007]. Applied to response inhibition, the right triangular and orbital regions may integrate response alternatives based on stored task representations, further accumulating information to resolve the conflict between response execution and suppression based on expectations and the actual occurrence of the stop signal.

In this context it is important to note that from the seven components we highlighted (ICs 1, 3, 5, 7, 10, 12, and 14) all but ICs 1 and 12 were associated with (predominantly right) activity in the inferior frontal triangularis and orbitalis regions. Among those was IC 14, of which feature 1 captured the expected depletion in alpha activity, here being more marked in stop trials. Source reconstructions indicated an additional association of IC 14 with parietal and associative occipital cortices, which would be in accordance with the assumption that attention-driven processing of the stop-stimuli may feed an accumulative decision module in the inferior frontal cortex.

Whereas the processing procedures applied for the purpose of this study can well be considered highly innovative, this does not mean that they come without limitations. The application of group decomposition to spectral EEG data does allow for the concurrent analysis of evoked and induced activity, but at the same time hinders a direct comparison with phenomena in the temporal domain. As a consequence, we cannot for sure determine that IC 1 indeed corresponds to N2 and/or P3 activity, because a back-transform of the IC's activity to the temporal domain is not possible here. Also, the spatial precision of inverse modeling to identify cortical generators could further be improved by denser electrode placement, which does not mean that the source constellations found here are necessarily invalid. Yet, differentiations within regions, as suggested here for the IFG, would benefit much from this increase in reliability. The higher spatial sampling may also help to further identify and locate activity originating from the pre-supplementary motor area, which often is found active in response inhibition tasks. Not least, we focused here on the go versus stop contrast, because it is the one most commonly applied in this context. However, it has been argued that the successful versus failed stop contrast may be more adequate, assuming that it better separates actual inhibition-related from other neurocognitive processes [e.g., Kenemans, 2015]. The methods applied here may need further refinement though to be sufficiently sensitive to such potentially subtle effects, which may not necessarily be reflected in amplitude changes but may rather be manifested in small temporal differences.

To conclude, our data strengthen the association of the right opercular region with the actual inhibition process, and this is the first study to identify a spatio-spectral scalp EEG profile for this network, which was characterized by concurrent activity below 12 Hz and above 20 Hz. Triangular and orbital inferior frontal regions may rather play a role in the selection of the most appropriate response alternative, based on incoming information and its integration in context of task representations. However, these processes exhibit strong overlap in the spatial, temporal, and spectral domain, with most relevant features being associated with activity patterns below 15 Hz, and occurring at around 200 ms. Thus, simple analyses in the electrode space may not find such effects, or may wrongly attribute them to major EEG phenomena such as the P3 with which they only may indirectly be associated (e.g., with the P3 indicating the evaluation of earlier inhibitory processing). In addition, evaluative and response selection components, such as IC 1 and IC 5 also influence task behavior, likely in cooperation with adaptive mechanisms across trials. This study provides testable hypotheses that can be addressed in future work, optimally binding electrophysiological phenomena to cortical activations as observed in functional magnetic resonance imaging, thereby providing us with the opportunity to integrate models of fast cortical processing with the well-described subcortical systems for motor-control.

REFERENCES

Alegre M, Alvarez-Gerriko I, Valencia M, Iriarte J, Artieda J (2008): Oscillatory changes related to the forced termination of a movement. Clin Neurophysiol 119:290–300.

Aron AR (2007): The neural basis of inhibition in cognitive control. Neuroscientist 13:214–228.

Aron AR (2011): From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. Biol Psychiatry 69:e55–e68.

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

Badre D, Wagner AD (2007): Left ventrolateral prefrontal cortex and the cognitive control of memory. Neuropsychologia 45:2883–2901.

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

Band GPH, van der Molen MW, Logan GD (2003): Horse-race model simulations of the stop-signal procedure. Acta Psychologica 112:105–142.

Bastos AM, Schoffelen JM (2016): A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. Front Syst Neurosci 9:175. (2016):

Bengson JJ, Mangun GR, Mazaheri A (2012): The neural markers of an imminent failure of response inhibition. Neuroimage 59: 1534–1539.

Chang Y-W, Lin C-J (2008): Feature ranking using linear svm. JMLR: Workshop and conference Proceedings 3:53–64.

Chang C-C, Lin C-J (2011): LIBSVM: A library for support vector machines. ACM Trans Intell Syst Technol 2:1–27.

Chatrian GE, Lettich E, Nelson PL (1985): Ten percent electrode system for topographic studies of spontaneous and evoked EEG activity. Am J Electroencephalogr Technol 25:83–92.

Coxon JP (2006): Intracortical inhibition during volitional inhibition of prepared action. J Neurophysiol 95:3371–3383.

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

Eichele T, Rachakonda S, Brakedal B, Eikeland R, Calhoun VD (2011): EEGIFT: Group independent component analysis for event-related EEG data. Comput Intell Neurosci 129365.

Huster RJ, Enriquez-Geppert S, Lavallee CF, Falkenstein M, Herrmann CS (2013): Electroencephalography of response inhibition tasks: Functional networks and cognitive contributions. Int J Psychophysiol 87:217–233.

Huster RJ, Plis SM, Lavallee CF, Calhoun VD, Herrmann CS (2014): Functional and effective connectivity of stopping. Neuroimage 94:120–128.

Huster RJ, Plis SM, Calhoun VD (2015): Group-level component analyses of EEG: Validation and evaluation. Front Neurosci 9:254.

Hsu C, Chang C, Lin C (2010): A practical guide to support vector classification. http://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf

Hyvärinen A, Oja E (2000): Independent component analysis: Algorithms and applications. Neural Netw 13:411–430.

Joel SE, Caffo BS, van Zijl PC, Pekar JJ (2011): On the relationship between seed-based and ICA-based measures of functional connectivity. Magn Reson Med 66:644–657.

Kenemans JL (2015): Specific proactive and generic reactive inhibition. Neurosci Biobehav Rev 56:115–126.

Krämer UM, Knight RT, Münte TF (2011): Electrophysiological evidence for different inhibitory mechanisms when stopping or changing a planned response. J Cogn Neurosci 23:2481–2493.

Lavallee CF, Meemken MT, Herrmann CS, Huster RJ (2014): When holding your horses meets the deer in the headlights: Time-frequency characteristics of global and selective stopping under conditions of proactive and reactive control. Front Hum Neurosci 8:994.

Lee TW, Girolami M, Sejnowski TJ (1999): Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. Neural Comput 11: 417–441.

Litvak V, Mattout J, Kiebel S, Phillips C, Henson R, Kilner J, Barnes G, Oostenveld R, Daunizeau J, Flandin G, Penny W, Friston K (2011): EEG and MEG data analysis in SPM8. Comput Intell Neurosci 2011:852961.

Schmiedt-Fehr C, Basar-Eroglu C (2011): Event-related delta and theta brain oscillations reflect age-related changes in both a general and a specific neuronal inhibitory mechanism. Clin Neurophysiol 122:1156–1167.

Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009): Intracranial EEG reveals a time- and frequency-specific role for the right inferior frontal gyrus and primary motor cortex in stopping initiated responses. J Neurosci 29:12675–12685.

Swann NC, Cai W, Conner CR, Pieters TA, Claffey MP, George JS, Aron AR, Tandon N (2011): Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: ELECTROPHYSIOLOGICAL responses and functional and structural connectivity. Neuroimage 59: 2860–2870.

Tallet J, Barral J, Hauert C-A (2009): Electro-cortical correlates of motor inhibition: A comparison between selective and non-selective stop tasks. Brain Res 1284:68–76.

Theodoridis S, Koutroumbas K (2008): Pattern Recognition, 4th ed. Oxford: Elsevier Ltd.

Wessel JR, Aron AR (2015): It's not too late: The onset of the frontocentral P3 indexes successful response inhibition in the stop-signal paradigm. Psychophysiology 52:472–480.

van den Wildenberg WPM, Burle B, Vidal F, van der Molen MW, Ridderinkhof KR, Hasbroucq T (2010): Mechanisms and dynamics of cortical motor inhibition in the stop-signal paradigm: A TMS study. J Cogn Neurosci 22:225–239.