ORIGINAL ARTICLE

Direct Recordings from Human Anterior Insula Reveal its Leading Role within the Error-Monitoring Network

Julien Bastin[1,2], Pierre Deman[1,2], Olivier David[1,2], Maëlle Gueguen[1,2], Damien Benis[1,2], Lorella Minotti[2,3], Dominique Hoffman[3], Etienne Combrisson[4,5], Jan Kujala[6], Marcela Perrone-Bertolotti[1,7], Philippe Kahane[2,3], Jean-Philippe Lachaux[5] and Karim Jerbi[5,8]

[1]University of Grenoble Alpes, F-38000 Grenoble, France, [2]Inserm, U1216, F-38000 Grenoble, France, [3]Neurology Department, CHU de Grenoble, Hôpital Michallon, F-38000 Grenoble, France, [4]Center of Research and Innovation in Sport, Mental Processes and Motor Performance, University of Lyon I, Lyon, France, [5]DYCOG Lab, Lyon Neuroscience Research Center, INSERM U1028, UMR 5292, University Lyon I, Lyon, France, [6]Department of Neuroscience and Biomedical Engineering, Aalto University, 02150 Espoo, Finland, [7]CNRS, LPNC, UMR 5105, F-38040 Grenoble, France and [8]Psychology Department, University of Montreal, Montreal, QC, Canada

Address correspondence to Julien Bastin, Institut des Neurosciences de Grenoble, Bâtiment Edmond J. Safra des Neurosciences, Chemin Fortuné Ferrini, Université Joseph Fourier, Site Santé La Tronche, BP 170 38042 Grenoble Cedex 9, France. Email: julien.bastin@ujf-grenoble.fr

Abstract

The ability to monitor our own errors is mediated by a network that includes dorsomedial prefrontal cortex (dmPFC) and anterior insula (AI). However, the dynamics of the underlying neurophysiological processes remain unclear. In particular, whether AI is on the receiving or driving end of the error-monitoring network is unresolved. Here, we recorded intracerebral electroencephalography signals simultaneously from AI and dmPFC in epileptic patients while they performed a stop-signal task. We found that errors selectively modulated broadband neural activity in human AI. Granger causality estimates revealed that errors were immediately followed by a feedforward influence from AI onto anterior cingulate cortex and, subsequently, onto presupplementary motor area. The reverse pattern of information flow was observed on correct responses. Our findings provide the first direct electrophysiological evidence indicating that the anterior insula rapidly detects and conveys error signals to dmPFC, while the latter might use this input to adapt behavior following inappropriate actions.

Key words: anterior cingulate cortex, EcoG, gamma, performance monitoring, saliency

Introduction

The ability to monitor our own errors is a critical aspect of learning that is mediated by coordinated activity in the dorsal medial prefrontal cortex (dMPFC, including anterior cingulate cortex [ACC], and presupplementary motor area [preSMA]) and anterior insula (AI) (Ullsperger and von Cramon 2001; Holroyd and Coles 2002; Ito et al. 2003; Wessel et al. 2012). However, the functional hierarchy and precise spatio-temporal dynamics of error-monitoring signals in the human brain remains debated. A classical view postulates a leading role of anterior cingulate cortex (ACC). This largely stems from scalp-electroencephalography (EEG) studies that have identified the ACC as the generator of an error-related negativity potential occurring around 100 ms after behavioral errors (Gehring et al. 1993; Dehaene et al. 1994; Debener et al. 2005). An alternative hypothesis posits that signals coming from the AI may be the actual input to the error-monitoring network (Sridharan et al. 2008; Ham et al. 2013) and may reflect an error-awareness signal (Ullsperger et al. 2010; Klein et al.

2013). In parallel, a recent study suggests that responses in supplementary motor area during error-monitoring occur systematically before ACC responses, which might be an indication that SMA plays a leading role in the error-monitoring network (Bonini et al. 2014).

These partly diverging accounts of information flow within the error-monitoring circuitry may be due to differences in experimental paradigms and spatial sampling but also to differences in the techniques used to assess brain's response to errors. For instance, functional neuroimaging and scalp-EEG differ in terms of spatio-temporal resolution, but they also measure different underlying physiological processes. Such discrepancies may explain why it has been difficult so far to reach a consensus on the dynamic properties of the error-detection network and the role of the insula therein.

In the present study, we sought to overcome some of these limitations by directly recording neural signals with depth electrodes implanted in key structures of the error-monitoring network. To this end, we recorded intracerebral electroencephalography (stereotactic electroencephalography, sEEG) in 6 epileptic patients, while they performed a stop-signal task (SST). This task has been used in numerous previous studies to probe the neural correlates of the error-monitoring network (Ito et al. 2003; Ramautar et al. 2006; Sharp et al. 2010). Using sEEG, we were able to assess both local activity and inter-areal connectivity modulations during error processing with both high spatial and temporal precision. Key to this study was the opportunity to simultaneously record AI and dmPFC structures (ACC or preSMA) in most patients. Precisely, all patients had electrodes in AI and, in 3 out of the 6 patients; the AI recordings were obtained simultaneously with signals from ACC or preSMA (or both). This allowed us to probe the role of AI by measuring error-related changes in (1) its local activity and (2) its long-range effective connectivity within the error-monitoring network. We found that error commission was associated with an increase in broadband gamma activity (BGA, 50–150 Hz) in AI as well as a remarkable reversal of information flow within the error-monitoring circuitry with AI driving dmPFC activity on error trials but dmPFC driving AI on correct responses. Taken together, our findings reveal the broadband electrophysiological underpinnings of human AI activity following behavioral errors and, most importantly, they posit AI as the key error-signaling component within the error-monitoring network.

Materials and Methods

Patients

Intracerebral recordings were obtained from 6 neurosurgical patients with intractable epilepsy (4 female, mean age: 29.5 ± 3.5 years.) at the Epilepsy Department of the Grenoble University Hospital. In order to localize epileptic foci that could not be identified through noninvasive methods, neural activity was monitored in lateral, intermediate, and medial wall structures in these patients using stereotactically implanted multilead electrodes (sEEG). Electrode implantation was performed according to routine clinical procedures and all target structures for the presurgical evaluation were selected strictly according to clinical considerations with no reference to the current study. Patients selected for the study were those whose electrodes sampled anterior insular cortices. All participants had normal or corrected to normal vision and provided written informed consent. Experimental procedures were approved by the Institutional Review Board and by the National French Science Ethical Committee (CPP Sud-Est V no. 09-CHUG-12).

Electrode Implantation

Eleven to fifteen semi-rigid, multilead electrodes were stereotactically implanted in each patient. Electrodes had a diameter of 0.8 mm and, depending on the target structure, contained 10–15 contact leads 2 mm wide and 1.5 mm apart (DIXI Medical Instruments). All electrode contacts were identified on each patient's individual postimplantation MRI. Each subject's individual preimplantation MRI was co-registered with the postimplantation MRI to determine the anatomical location of each contact and to compute all contacts coordinates in the Montreal Neurological Institute (MNI) space using standard Statistical Parametric Mapping (SPM) algorithms. Visual inspection of the contacts locations was also used to check whether each sEEG contact was located in gray or white matter.

Intracerebral EEG Recordings

SEEG recordings of 128 contacts in each patient were conducted using a commercial video-sEEG monitoring system (System Plus, Micromed, Italy). The data were bandpass-filtered online from 0.1 to 200 Hz and sampled at 512 Hz, using a reference electrode located in white matter. Each electrode trace was subsequently re-referenced with respect to its direct neighbor (bipolar derivations with a spatial resolution of 3.5 mm) in order to achieve high local specificity by canceling out effects of distant sources that spread equally to both adjacent sites through volume conduction. All electrodes exhibiting epileptiform signals were discarded from the present study. This was achieved in collaboration with the medical staff and was based on visual inspection of the recordings and by systematically excluding data from any electrode site that was a posteriori found to be located within the seizure onset zone. We report data from the following sites: AI in the dorsal part of anterior short and accessory insular gyri (Craig 2009), preSMA corresponding to the anterior part of Brodmann's area 6 (Picard and Strick 1996) and ACC corresponding to Brodmann's area 24, including the dorsal ACC and the rostral ACC (McCormick et al. 2006).

SST Experimental Procedures

Patients completed on average 414 ± 41 trials of a SST, which was composed of 2 types of trials (GO and STOP trials, Fig. 1A) presented in a randomized order. The SST generates a large number of errors during stop trials. Comparison of neural responses across correct and incorrect button presses allowed us to identify electrode contacts that responded in an error-monitoring selective manner equivalent to that typically observed in the functional magnetic resonance imaging (fMRI)-defined error-monitoring network. The majority (70%) of trials were GO trials, which consisted of an imperative GO cue (a white circle subtending 1° of visual angle) that prompted patients to press a button with the right index as fast as they could. During STOP trials (30% of trials), patients were asked to inhibit their prepared movement in response to a STOP cue (a white cross subtending 1° of visual angle) that unpredictably followed the GO cue after a variable delay (the stop-signal delay, SSD). GO cues were preceded by a variable fixation period (range: 500–1500 ms) and vanished after button press or after 1000 ms during GO trials and after the SSD during STOP trials. STOP cues disappeared after 500 ms or after incorrect button press, in case of errors.

The ability to stop a response is related to the SSD value (the longer the SSD, the more difficult it is to stop). The SSD was varied from trial to trial to adjust task difficulty using a staircase procedure: if subjects succeeded in withholding the response in a STOP trial, the SSD increased by 50 ms; if they failed, SSD decreased by the same amount of time. The staircase procedure ensured that
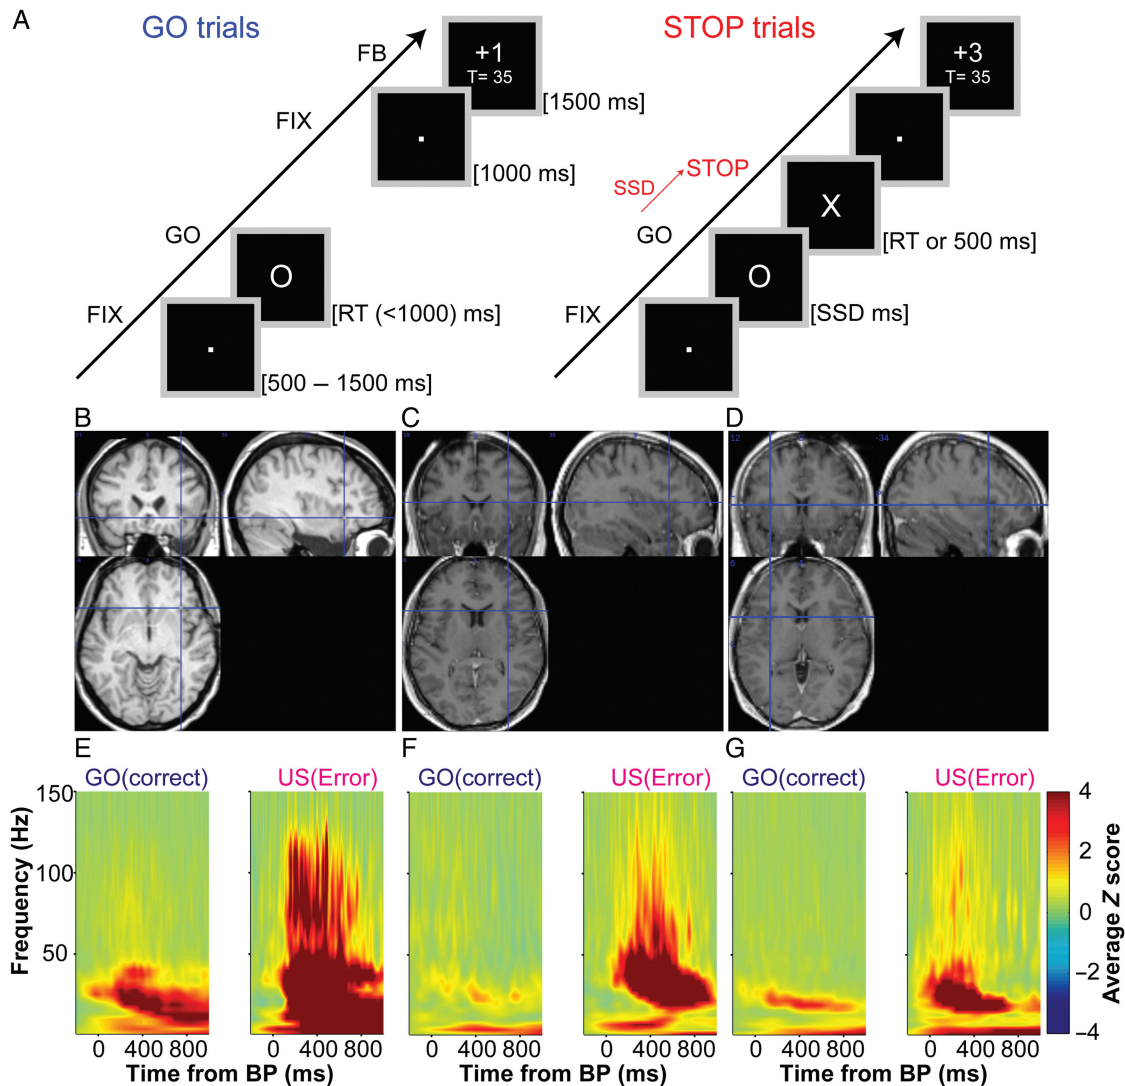
Figure 1. Depth electrode electrophysiological recordings in the AI following errors and correct responses. (A) During the Stop-signal paradigm, participants had to rapidly press a button after GO cues and try to withhold their prepared response after stop cues. (B–D) Illustrative anatomical locations of AI recording in 3 patients (B: P1; C: P2; D: P3). (E–G) Striking power increase in broadband gamma (50–150 Hz) and in the beta band (15–35 Hz) in the AI after errors in this 3 patients (E: P1; F: P2; G: P3). These illustrative TF maps represent increases and decreases in spectral power compared with a prestimulus baseline level (baseline, from [−700 to −200] ms).

the subjects would fail in withholding their response in around 50% of the stop trials. To maintain high attentional and motivational level, a feedback lasting 1500 ms appeared 1000 ms after each trial to indicate success or failure to patients in the current trial. Successful trials were accompanied by a score increase of 1 point for GO trials and of 3 points for STOP trials. Unsuccessful trial feedbacks consisted of losing 1 point for GO trials and of 3 points for STOP trials, respectively. During feedback, both score increase/decrease in the current trial and the total of points won previously were displayed. Visual stimuli were delivered on a 19-in. TFT monitor with a refresh rate of 60 Hz, controlled by a PC with Presentation 14.1 (Neurobehavioral Systems, Albany, CA). The monitor was placed 70 cm away from the subject's eyes. All patients responded to the tasks through right-hand button presses.

Signal Preprocessing and Time–Frequency Power Analysis

The first step of our analysis was to determine the typical response to errors in the AI. To do this, we used the freely available ELAN analysis toolbox (Aguera et al. 2011), developed at the Brain Dynamics and Cognition Laboratory (INSERM U1028, Lyon, France) to compute standard time–frequency (TF) wavelet decomposition of power (Lachaux et al. 2003), which allowed us to identify major TF components of interest across correct vs. incorrect button presses. The frequency range was 1–150 Hz and the time interval included a 1000-ms prestimulus baseline and lasted until 1000 ms after button presses. Normalized power values (Z scores were computed using a −700 to −200 ms prestimulus baseline) were computed and averaged across trial types for each sEEG contact-pairs to visualize task-induced power modulations.

Computation of Single-Trial Instantaneous Broadband Envelopes

Broadband gamma activity (BGA) was extracted with the Hilbert transform of sEEG signals using custom Matlab scripts (Mathworks Inc., MA, USA) as follows: the sEEG signals were first bandpass filtered in 10 successive 10-Hz-wide frequency bands

(e.g., 10 bands, beginning with 50–60 Hz up to 140–150 Hz). For each bandpass-filtered signal, we computed the envelope-using standard Hilbert during the entire experiment with a time resolution of 15.625 ms. Again, for each band, this envelope signal (i.e., time-varying signal amplitude) was divided by its mean across the entire recording session and multiplied by 100 for normalization purposes. Finally, the envelope signals computed for each consecutive frequency bands (e.g., 10 bands of 10 Hz intervals between 50 and 150 Hz) were averaged together, to provide one single time-series (the BGA) across the entire session, expressed as percentage of the mean. This time series was then divided into trial-wise responses, which were analyzed as detailed below. We performed similar computations to estimate functional modulation in the other frequency bands (4–8 Hz; 8–12 Hz; 15:35 Hz; For 15:35 Hz activity, we used 5 successive 5-Hz-wide frequency bands).

Identification of Sites in the AI with Error-Specific Responses

We first used the BGA responses to identify error-selective contacts in the insula for theoretical reasons. For each sEEG contact, we calculated the average BGA during the post-button press interval (0–800 ms) for each of the correct and incorrect button presses. Error-related sites within the insula corresponded to contact-pairs for which the unpaired t-test on the activity between error and correct trials was positive and reached significance after controlling the P values for multiple comparisons over the number of available sEEG electrodes for each patient ($n = 95 \pm 3$ contact-pairs for each patient) using the False Discovery Rate -FDR- algorithm (Genovese et al. 2002). To exclude recording sites in AI with BGA responses that could be caused by a saliency effect (driven by the unexpectedness of STOP cues) rather than error processing, we performed a further analysis where BGA responses were time-locked to stop cues. We calculated the average BGA during the postSTOP-cue interval (0–800 ms) for each of the successful STOP and unsuccessful STOP trials. We then compared these values using an unpaired t-test and corrected the P-values for multiple comparisons over the number of available sEEG electrodes using the FDR. Contacts that showed significantly greater response to errors time locked to button presses as well as errors time-locked to stop cues were deemed to be error-related sEEG contacts. A similar type of conjunction analysis was used previously in a neuroimaging study to clarify the respective contribution of saliency versus error component in AI, preSMA and ACC (Wessel et al. 2012).

Response Onset and Peak Latencies

We examined the timing of error-monitoring signals. These signals were defined by the difference between correct and incorrect responses, time-locked to button presses or to stop cues. We calculated onset latencies for each contact by comparing BGA for the appropriate conditions for each 15.625 ms time bin in the [−250 to 1000 ms] time interval, using single-trial responses at each time bin as observations (unpaired t-tests across trials). The resulting P-values were corrected for multiple comparisons in the time domain with the FDR algorithm to take into account the 81 tests performed across time. The onset latency for a contact corresponded to the first time bin at which a significant corrected P value was observed. In addition, the duration of experimental effects had to be superior or equal to 90 ms (6 bins) to derive onset latency for a sEEG contact. We also calculated the peak latency for each contact, which corresponded to the time point at which the activity observed after errors reached its maximum.

Computation of Event-Related Potentials

Event-related potentials (ERPs) were measured by averaging data segments across epochs in the time domain of all bipolar derivations after applying a zero-phase low-pass Butterworth filter with a cutoff frequency of 15 Hz.

Whole-Brain Analyses

To generate statistical images summarizing our findings at the group level, we took the following approach. First, we computed for each patient and for each sEEG site, a Welsh t-test between error and correct trials time-locked to button presses during the 800 ms following the motor responses (either in the gamma or in the beta bands). FDR correction was then performed to correct for multiple comparisons across the number of sEEG contacts of each patient. A 3-mm isotropic 3D image of the resulting P-values was then obtained by assigning the P-values to all voxels in a sphere of 10 mm radius centered on the MNI position of the corresponding sEEG site. Such spatial smoothing is needed because of the low spatial coverage within-subject and high-sampling variability across subjects. In addition to the spatial heterogeneity, there is also a fair amount of variability across patients in the temporal domain (this is why we used an 800 ms window for the analysis). To display the voxels activated at the group-level, we used individual statistical images and count within each voxel the number of patients showing a statistically significant change in activity. A voxel was considered to be significant if the FDR corrected P-value was < 0.05 in at least 2 patients. This contrast between unsuccessful suppression (US) and a correct button press (GO) conditions is thus denoted by "US > GO" (see Fig. 5).

A similar approach was used to display cortical distributions of error-related activities that rule out a pure saliency effect, that is, activity increases merely generated by the occurrence of an unexpected stop cue. To this end, we used a conjunction analysis that depicted the number of patients with a statistically significant change in power not only time-locked to the button press (failed suppression vs. correct button presses) but also time-locked to the stop cue (failed suppressions vs. successful suppression [SS]). Similarly, to the response-locked analysis, the stop-cue-locked responses were assessed using a Welsh t-test and a 0–800 ms time interval ($P < 0.05$ FDR corrected). This conjunction approach identifies voxels where 2 contrasts are statistically significant: (a) US vs. correct button press (GO), and (b) US vs. SS, and is thus denoted by "US > GO & US > SS" (see Fig. 5).

Finally, previous activations from functional neuroimaging studies were added to the final images, representing the rAI, pre-SMA, and ACC regions (Ramautar et al. 2006; Klein et al. 2007; Sharp et al. 2010; Wessel et al. 2012; Ham et al. 2013). We used the MNI coordinates reported in these 5 landmark papers to determine on these brain maps the areas typically involved in error-monitoring. All voxel <10 mm from the MNI coordinates reported in these publications were colored in light gray (see Fig. 5). For display purposes, 3D images were projected on an inflated cortical mesh of a canonical brain using routines from the Statistical Parametric Mapping (SPM 12) software. Note that we projected all left sEEG sites onto the right hemisphere. Although this assumes the absence of lateralization effects, it was useful in order to visualize and summarize the activity across the 6 patients of this study using these whole-brain representations. The within-subject statistical tests for all AI sites are summarized in Supplementary Table S1.

Measuring Task-Related Network Interactions

The dynamics of information flow between anterior insular and medial prefrontal components of the error-monitoring network

were probed using Granger causality (GC) (Granger 1969). In practice, GC was estimated between the unfiltered raw SEEG signals simultaneously recorded in AI and key structures of the dorsomedial prefrontal cortex (ACC and/or preSMA). This was possible because in 3 out of 6 patients, AI recordings were obtained simultaneously either with ACC or preSMA (or both). In particular, participants P2 and P4 were implanted in AI and in dACC (Fig. 6), and participants P3 and P4 were implanted in AI and in preSMA (Fig. 7). A careful identification of the anatomical locations of the recording sites is of the highest importance, in particular when it comes to delineating the borders between preSMA and dACC (Amiez et al. 2013, 2015). Supplementary Figure S1 shows the coordinates and locations of the preSMA and dACC sites overlaid on the individual subject MRI. The detection of statistically significant BGA responses described in the first part of the study provided a practical and functionally relevant electrode-selection procedure for the subsequent GC analysis. Only data from insular sites that showed significant error-related BGA responses were used to assess the error-related modulations of connectivity with ACC and/or preSMA. The GC analysis was conducted using a MATLAB Toolbox (Seth 2010). In the analysis, Granger Causality was computed across trials separately in the GO (correct button press) and US (unsuccessful STOP, i.e., erroneous button press) conditions; the computations were performed using partially overlapping 200 sample-long windows (ca. 390 ms given the 512 Hz sampling frequency), with window centers spanning the range from 293 ms preceding button press to 780 ms after button press. Each time window was detrended, the ensemble mean was subtracted, and first-order differencing was applied. The optimal model order was determined by evaluating the Akaike (Akaike 1974) and Bayesian (Seth 2005) information criteria (AIC and BIC respectively). While the former failed to converge, the BIC yielded an optimal model order of 25. Subsequently, the Granger Causality terms were computed, and the difference of influence (DOI) measure was computed. This represents the overall net directionality between 2 sites A and B, obtained by assessing the difference between GC measured from A to B and the GC measured from B to A. The statistical significance of DOI terms was determined using permutation tests (200 permutations), a standard nonparametric technique widely used in neuroscience. We used the implementation of permutation test provided with the Granger Causality toolbox (Seth 2010). Furthermore, we also ran the same analysis, with different window size and model order parameters. This led to comparable results and confirmed robustness of the GC analysis. Note that while we often refer to the coupling captured by Granger Causality throughout the manuscript as neuronal interaction (with a putative dominant direction of influence), it is important to keep in mind that what is actually estimated is a directed statistical measure of functional connectivity.

Computing Task-Related Temporal Modulation of Heartbeat Rate

The time course of the heart-beat rate was estimated from the electrocardiogram (ECG), which was also sampled at 512 Hz. The R peaks of the ECG's QRS complex were identified using a dedicated detection algorithm provided as part of the FMRIB plug-in to EEGlab (Niazy et al. 2005; Wessel et al. 2011). Next, the duration of the inter-beat interval was extracted by computing the difference between the latencies of every 2 successive R peaks. The time series of heart rate in beats per minutes (bpm) was then estimated by linear interpolation of the inter-beat intervals. The data were individually checked for artifacts via visual

inspection and subsequently segmented into 1-s nonoverlapping intervals from one second preceding the response to 3 s following the response. For each subject, we thus obtained estimates of heartbeat rate at 4 time intervals around button presses. By separately averaging these heart rate time series for correct and erroneous button presses, we were able to compare heart rate values in correct (GO trials) and incorrect (US trials) conditions. Statistical significance for these comparisons was assessed via unpaired t-tests applied to each of the 4 time intervals.

Results

Intracerebral EEG data were collected from a total of 559 intracerebral sites across 6 patients with intractable epilepsy (4 women; mean age, 30 ± 3.5 years) while they performed a SST (Fig. 1A). The average error rate was $3.8 \pm 1.5\%$ during GO trials and $53.9 \pm 2.6\%$ during STOP trials. This group performance is close to the expected 50% level given the staircase procedure used to adjust dynamically stop-trial difficulty as a function of subject performance (see Materials and Methods for details). Patients committed only very few omission errors during GO trials ($n = 5$, 0, 40, 6, 23, and 4 omission errors in participants P1–P6 respectively). The mean response time computed over incorrect stop trials was 495 ± 61 ms, which was significantly shorter ($P = 0.0036$) than the mean response time over correct GO trials (598 ± 66 ms). Overall, all subjects performed well and the behavioral results were comparable with those reported in previous sEEG studies based on a stop-signal task (Swann et al. 2009, 2012).

Compared with correct button presses (GO trials), erroneous button presses (i.e., unsuccessful stop, US trials) were associated with a strong broadband increase in neuronal population response in anterior insular cortex of all 6 patients (Fig. 1B–G. In the following, we primarily focus on BGA (50–150 Hz) because previous studies on humans have shown that this neural marker is correlated with neuronal firing (Manning et al. 2009) as well as with fMRI BOLD responses (Lachaux et al. 2007; Jerbi et al. 2009). Nevertheless, for the sake of completeness, task-related changes in the lower frequency bands are also examined and discussed.

In order to precisely assess the dynamics of AI activity during error-monitoring, we quantified BGA on a trial-by-trial basis and across conditions (Fig. 2). We found a robust increase in BGA in AI when participants had to refrain from pressing a button, but failed to do so (US). This increase was absent when the subjects actually performed a required button press (GO trials). This effect was highly consistent across trials and patients (Fig. 2). This BGA increase was found to be statistically significant in 22 AI recording sites across all patients (Fig. 2; t-test on single-trial BGA between 0 and 800 ms posterror vs. correct responses; all t-values >3.4; $P < 0.05$ corrected, see Materials and Methods section). This error-monitoring response in AI started 44 ± 14 ms after button press and lasted slightly longer than half a second (mean effect duration was 632 ± 70 ms and the effect peaked on average at 274 ± 48 ms following the erroneous button press). These latencies were determined as the time at which responses to incorrect button presses (US) became statistically different from responses to correct button presses. Note that if we only selected a single representative recording site per patient, the above latencies remained largely unchanged (57 ± 15 ms, 553 ± 72 ms and 254 ± 66 ms, for the mean onset, duration and peak latencies, respectively).

Figure 2G shows the anatomical location in each patient of the AI electrode contacts that showed the strongest error-related BGA response (mean MNI coordinates: $x = \pm 35$, $y = 17$, $z = 3$). These AI sites were located within the accessory and anterior short gyri of the insular agranular cortex that functionally correspond to the
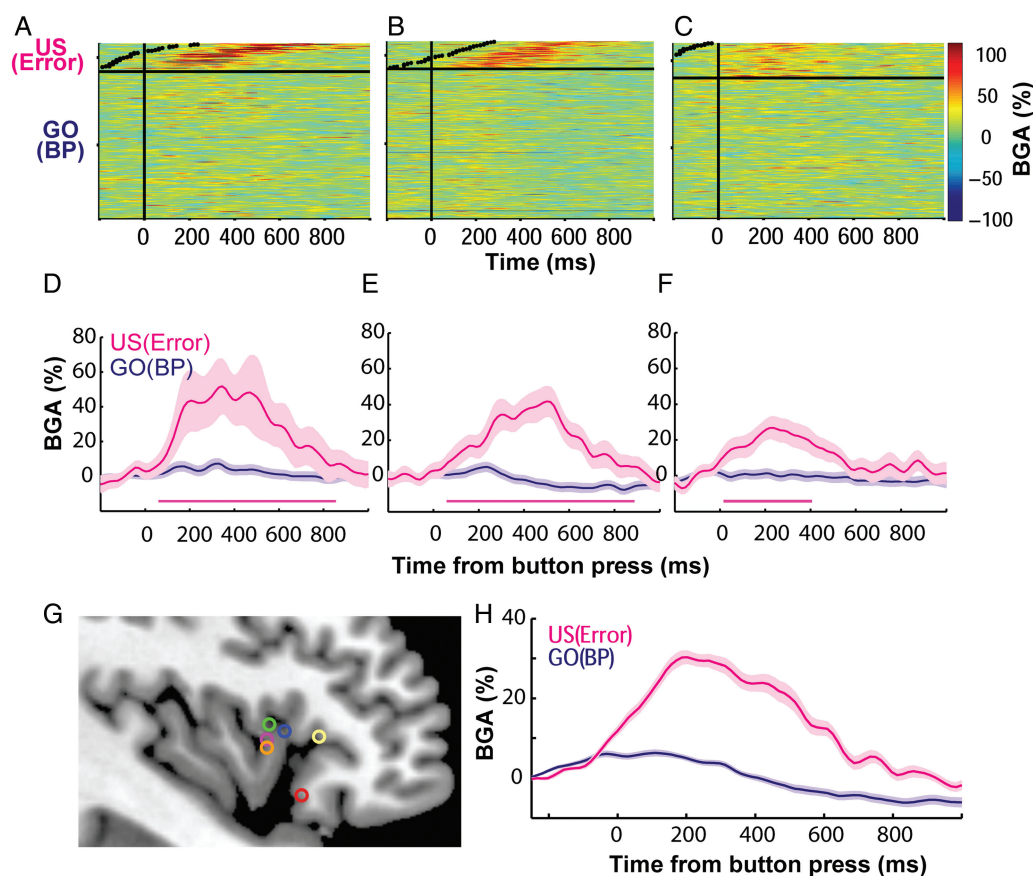
Figure 2. Single-trial gamma-band activities in the AI following errors and correct responses. (A–C) Single-trial gamma-band activity (BGA, expressed in percentage of signal change) in 3 illustrative AI contacts during errors and correct button presses (A: Patient 1, B: Patient 2, C: Patient 3; MNI coordinates of each AI contact are indicated in Fig. 1). The single-trials in condition US are arranged in ascending order of appearance of the stop signal (indicated by black dots). (D–F) Profile of gamma-band activities aligned to button presses after errors (pink) and correct responses (blue) for the same 3 AI contacts (D: P1, E: P2, F: P3). Shaded areas represent ± 1.96 SEM across trials. Pink horizontal lines indicate segments with statistically significant differences between correct and error trials. (G) Anatomical location of most selective AI contacts that respond preferentially to errors in the AI (*n* = 6 patients; one AI contact/patient: each circle corresponds to one patient). (H) Grand-average profile of gamma-band activity (*n* = 6 AI contacts) after errors (pink) and correct responses (blue). Shaded areas represent ±1.96 SEM across patients.

cognitive subregion of the insula (Dupont et al. 2003; Kurth et al. 2010). Figure 2*H* shows the grand-average error-related BGA across these sites. Note that this group-level overview remained similar when all AI contacts with significant responses were used (*n* = 22) rather than one contact per patient. Furthermore, removing the functional selection bias did not affect the group-level response, suggesting that the error-monitoring effect in AI was very robust (i.e., this was demonstrated by keeping in the group-level analysis all sEEG sites within AI instead of selecting only the sites that statistically showed the error-monitoring effect individually). In theory, the AI response following errors might simply be a response to the saliency of the stop cue (an unexpected and behaviorally significant event). But if this were the case, then the same response should be observed after STOP signals even when patients successfully withhold their response. Therefore, to rule out that the observed significant BGA increases in the failed STOP trials across the 22 insular sites could be entirely explained by stimulus saliency, we directly compared insular responses in correct STOP trials (SS) to those observed in incorrect STOP trials (US). Since the salient stop signal was present in both, an account exclusively based on stimulus saliency would predict comparable insular responses in both conditions or even a strongest response during SS trials (because patients might have paid more attention to the unexpected stop cue

when they were successful). This is not what we observed: we found that although BGA activity was enhanced in both cases (in line with the involvement of AI in saliency processing, Fig. 3), the amplitude of the BGA response was significantly higher after US compared with SS in 15 out of the 22 AI contacts (unpaired *t*-tests on single-trial BGA between 0 and 800 ms post-Stop-cue onset performed for each AI site, *P* < 0.05). In other words, out of the 22 AI sites that display significant BGA increases for US compared with GO trials (i.e., analysis time-locked to button press), a large portion (15 sites) also showed a significant BGA increase for US compared with SS (i.e., analysis time-locked to stop-cue). Thus, the temporal profile of BGA response (Fig 3) suggests an additive effect of saliency and error monitoring in the AI: the initial BGA component could reflect saliency processing while the higher peak and more sustained response may reflect more specific error-monitoring processes.

Note that in the US condition the Stop cue disappeared on button press, whereas in SS the STOP cue remained visible for 500 ms. One could therefore ask whether the higher gamma response in AI in the US condition (compared with SS) could be due to the subtle difference in the visual stimulus. This interpretation is unlikely given that the extinction of the stimulus upon button press in the GO condition did not elicit a comparable gamma response. In addition, if the gamma response was related
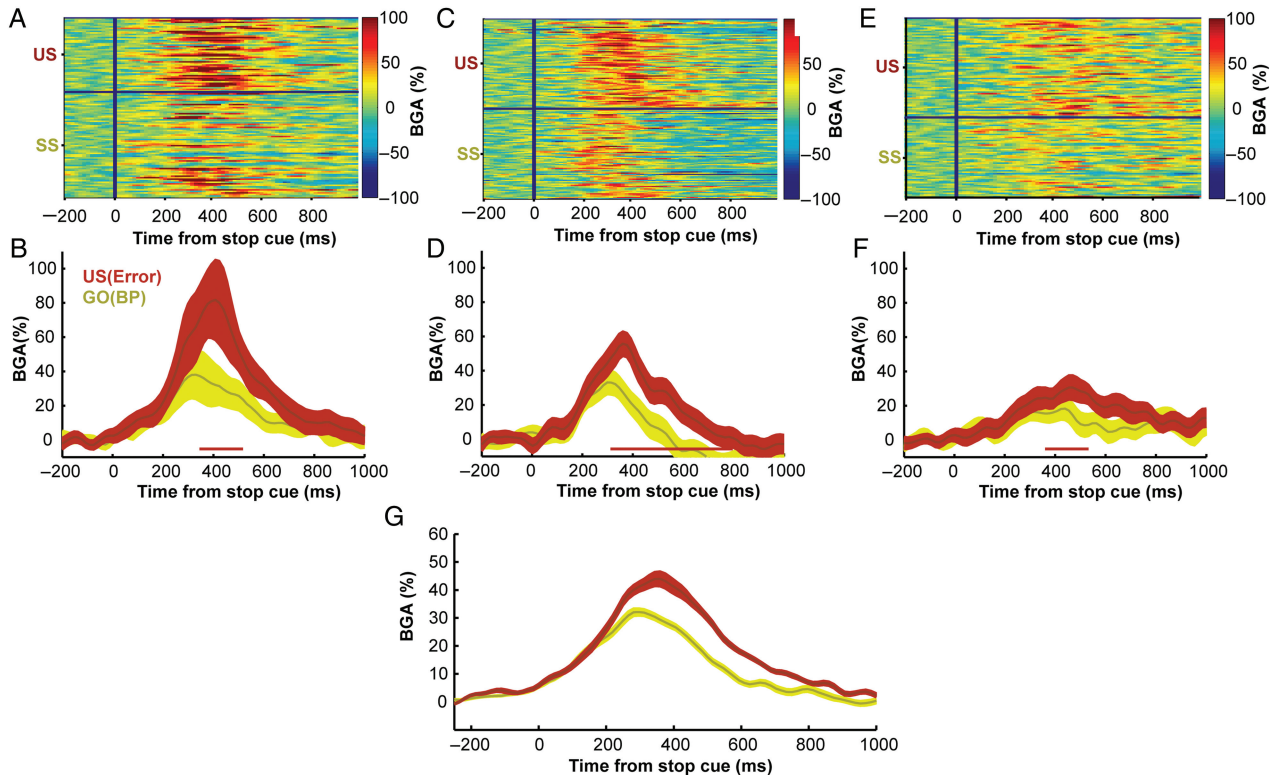
Figure 3. Stop-cue-related gamma-band responses in successful versus US. Single-trial gamma-band responses in the AI following errors and correct inhibition. (A,C,E) Single-trial gamma-band responses during errors and correct stop trials in 3 illustrative patients (P1–P3: MNI coordinates of each AI contact is indicated in Fig. 1). (B,D,F) Temporal profile of mean gamma-band amplitude aligned to stop cues for the same patients (P1–P3) during errors (orange) and correct stop trials (yellow). Shaded areas represent ± 1.96 SEM across trials. Orange horizontal lines indicate segments with statistically significant differences between correct and error trials. (G) Grand-average profile of gamma-band activity ($n = 6$ AI contacts) after errors (orange) and correct inhibition (yellow). Shaded areas represent ± 1.96 SEM across patients.

to the duration of cue presentation, we would have expected to find a higher response for SS compared with US, which is the opposite of what we observed. Nevertheless, we performed a control analysis where we repeated the US versus SS analysis but this time restricting the US trials to the trials where the participants pressed the button even before the appearance of the Stop cue. This guaranteed that the sensory stimulus was identical across the 2 conditions. Because such trials are rare, this analysis was only possible using data from 2 participants (P1: $n = 13$ US trials and P3: $n = 45$ US trials). Critically, the Stop-cue-related BGA activity was significantly greater ($P < 0.05$, FDR corrected) during US trials compared with SS trials for these 2 patients. This confirms that the AI GBA Stop-cue effect (US > SS) is preserved when we control for visual stimulus properties. Alternative interpretations based on motor-components can also be ruled out given the absence of a statistically significant BGA response in the GO trials (see Fig. 2).

The focus here on using BGA (50–150 Hz) as a neural marker of error-monitoring is mainly motivated by the established evidence for tight correlations between BGA and neuronal spiking and also between BGA and BOLD activity (Logothetis et al. 2001; Niessing et al. 2005; Lachaux et al. 2007; Ray and Maunsell, 2011). However, as can be seen in the TF representations (Fig. 1E–G) error-related power increases are also prominent at lower frequencies. We therefore also ran post hoc exploratory analyses in the theta (4–8 Hz), alpha (8:12 Hz), and beta (13–35 Hz) bands (Fig. 4) for all $n = 22$ AI contacts identified above. Significant error-related power increases were observed in the beta (13–35 Hz) band ($n = 20$ AI contacts when the analysis was time-locked to

button presses; and $n = 14$ contacts when the analysis was time-locked to stop cues). The results were not consistent across patients in the theta (4–8 Hz) and alpha (8:12 Hz) bands (Fig. 4).

Moreover, we asked whether another type of error, failure to press the button on Go trials (i.e., omission errors) was also associated with gamma power increases. However, because these errors are very rare events, we were only able to test this hypothesis in one participant (P3) who committed 40 omission errors. Remarkably, we found that AI activity in the gamma frequency range was higher after omission errors than after correct GO trials during an 800 ms time interval after feedback onset (AI contact P3, unpaired t-test, $t = 4,85$, $P < 0.05$, FDR corrected). No such effect was present in the beta frequency range. Although highly interesting, this result needs to be taken with a lot of caution because we were only able to test it in one participant given that the study was not designed to explore omission errors.

Although our primary focus was on the AI, we also performed exploratory whole-brain analyses to identify regions that were differentially activated during error-monitoring. Figure 5A shows that errors were associated with significant increases in BGA mainly in AI and dmPFC (ACC and preSMA), in addition to smaller clusters in the dorsolateral prefrontal cortex (including the inferior frontal gyrus) and in the dorsal and ventral premotor cortex. The rest of the active clusters reasonably matched components of the saliency network (Menon and Uddin 2010). Figure 5B shows that a conjunction analysis of the gamma activity corresponding to errors time-locked to button presses and stop cues revealed an effect only in the AI, ACC and ventral premotor sites. This
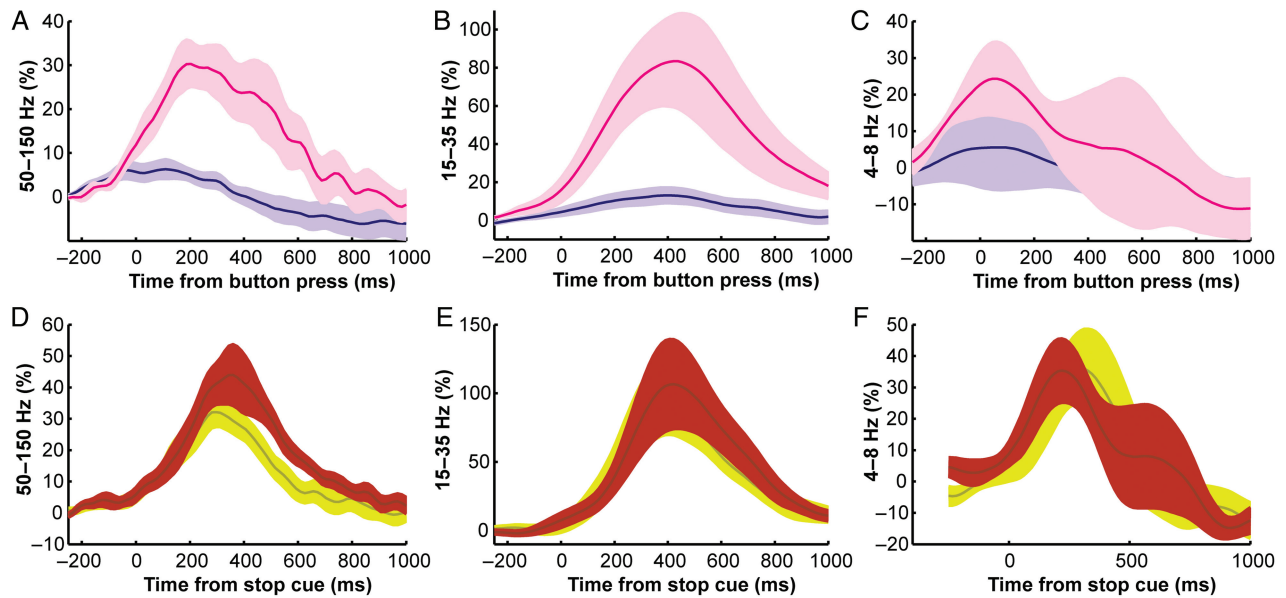
Figure 4. Grand-average electrophysiological responses in the AI following errors and correct responses. (A–C) Grand-average profile of theta (4–8 Hz), alpha (8–12 Hz), and beta (15–35 Hz) amplitudes ($n = 6$) following errors and correct button presses ($t = 0$ corresponds to button press). (D–F) Grand-average profile of theta, alpha, and beta amplitudes ($n = 6$) following correct and incorrect stop trials ($t = 0$ corresponds to stop-cue presentation). The format and conventions are the same as in Figure 2 (panels A–C) and 3 (panels D–F).
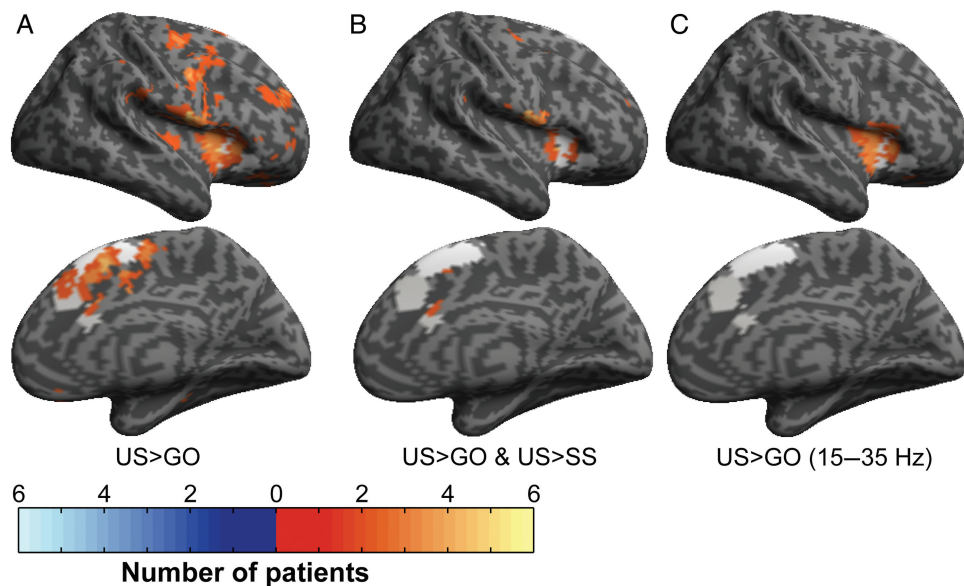
Figure 5. Whole-brain analyses. (A) Dynamics of the error-monitoring signal across all patients in the gamma band (errors > correct button presses, $P < 0.05$, FDR corrected). The upper views depict the right hemisphere (lateral surface) and the lower views the mesial surface of the right hemisphere. (B) Dynamics of the error-monitoring effect in the gamma band with a saliency mask (errors > correct button presses and US > SS trials in the [0–800 ms] time interval, $P < 0.05$, FDR corrected). (C) Dynamics of the error-monitoring signal across all patients in the beta band (errors > correct button presses, $P < 0.05$, FDR corrected). $t = 0$ corresponds to button press. Hot colors indicate activity increase after errors whereas cold colors indicate a decrease of activity after errors.

suggests a more specific encoding of error-related processes at these sites because they were activated in the gamma band when the error-signal was time-locked to button presses but also when contrasting unsuccessful and successful stop trials (corrected $P < 0.05$). Finally, a different picture emerged when estimating the dynamics of error-related activities within the beta band (13–35 Hz): error commission was associated with a clear increase of beta power that was anatomically restricted to the AI (Fig. 5C).

Note that this beta band activation was restricted to a yet smaller portion of AI following the same conjunction analysis performed for gamma responses in order to rule out stimulus saliency effects (not shown, but see Supplementary Table S1 for details).

Next, to probe the dynamic contribution of AI to error-processing from a network perspective, we assessed the directionality of information flow by measuring GC between AI and simultaneously recorded key dmPFC structures. In particular,
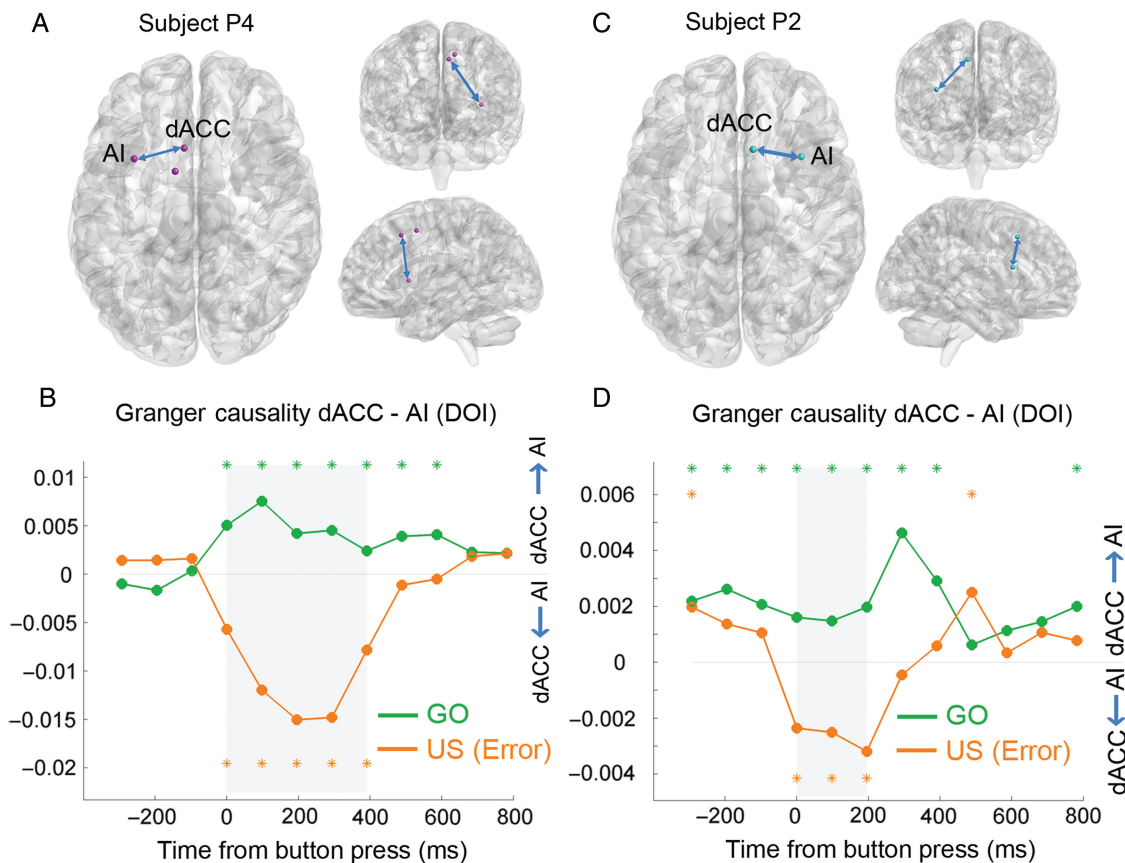
Figure 6. GC between AI and dorsal anterior cingulate cortex (dACC) following errors and correct responses. (A) Anatomical locations of AI and dACC recording sites used for GC analysis in patient P4. (B) Direction of influence (DOI) measured as the net difference between GC measured in both directions between dACC and AI. The blue arrows indicate that positive DOI values represent a predominant flow from dACC to AI, while negative DOI values represent a predominant flow from AI to dACC. The shaded areas depict the time window for which DOI was both significantly positive in one condition and significantly negative in the other condition (i.e., statistically significant directionality reversal, $P < 0.05$). (C) Same as in panel (A) but for patient P2. (D) Same as in panel (B) but for patient P2.

we probed GC between AI and either dorsal ACC (dACC), preSMA or both, depending on availability of simultaneous recordings in these structures. This was possible in 3 out of 6 patients (see Materials and Methods for details). The causality analyses revealed a remarkable change in the pattern of information flow between AI and dmPFC after errors: while after correct button presses the Granger DOI (direction of influence) indicated a prominent drive from dACC and preSMA toward AI, the reverse relationship was observed after behavioral errors (failed inhibition), with a prominent drive now from AI to dACC and preSMA (Figs 6 and 7). Interestingly, these DOI reversals (permutation tests, $P < 0.05$) occurred first in the dACC–AI interaction (within 400 ms following button press), and then for the preSMA-AI interaction (between 600 and 800 ms after button press). These DOI reversal windows were determined as the time points for which DOI switches from being statistically significant in one direction on correct button press trials to being statistically significant in the opposite direction on error trials.

In order to relate our findings to previous noninvasive EEG literature, we also computed event-related potentials (ERPs) on correct and incorrect button presses. Similar to GBA analysis, ERPs were calculated for AI sites across all 6 patients and among these also for available ACC sites (2 patients) and preSMA sites (2 patients). Consistent error-related negativity potentials (ERn) were found in AI (Supplementary Fig. 2) and in ACC (Supplementary Fig. 3), but not in preSMA. Although, the peaks on the

recorded ERP were not easy to identify in all patients; on average, the ERP in AI peaked at 239 ± 27 ms ($n = 6$ patients, see Supplementary Fig. 2). Interestingly, in the 2 individuals that had sEEG recordings simultaneously in AI and ACC, we observed that the AI peak preceded ACC peak in one patient while the other patient showed the opposite pattern (Supplementary Fig. 3). However, when comparing the latency at which error potentials began to differ from correct response potentials ($P < 0.05$), we found that the onset of the effect in AI preceded that of ACC in both patients (P2 and P4). As such, the temporal order of the onset of ERP difference (error vs. correct responses) in AI and ACC is in agreement with the GC results. ERP peak latencies, however, seem to be a less robust indicator of directionality.

Finally, we analyzed the electrocardiogram from all patients during a 4-s window around button presses (-1 s to $+3$ s) to verify whether error trials were associated with interoceptive responses. As expected from previous studies, we found a significant deceleration in heart rate after errors compared with correct trials in all 6 patients (Supplementary Fig. 4).

Discussion

We found that committing an error (failure to inhibit a button press in a SST) is associated with a rapid increase of electrophysiological activity in the AI as well as a reversal of the direction of information flow between AI and key nodes of the
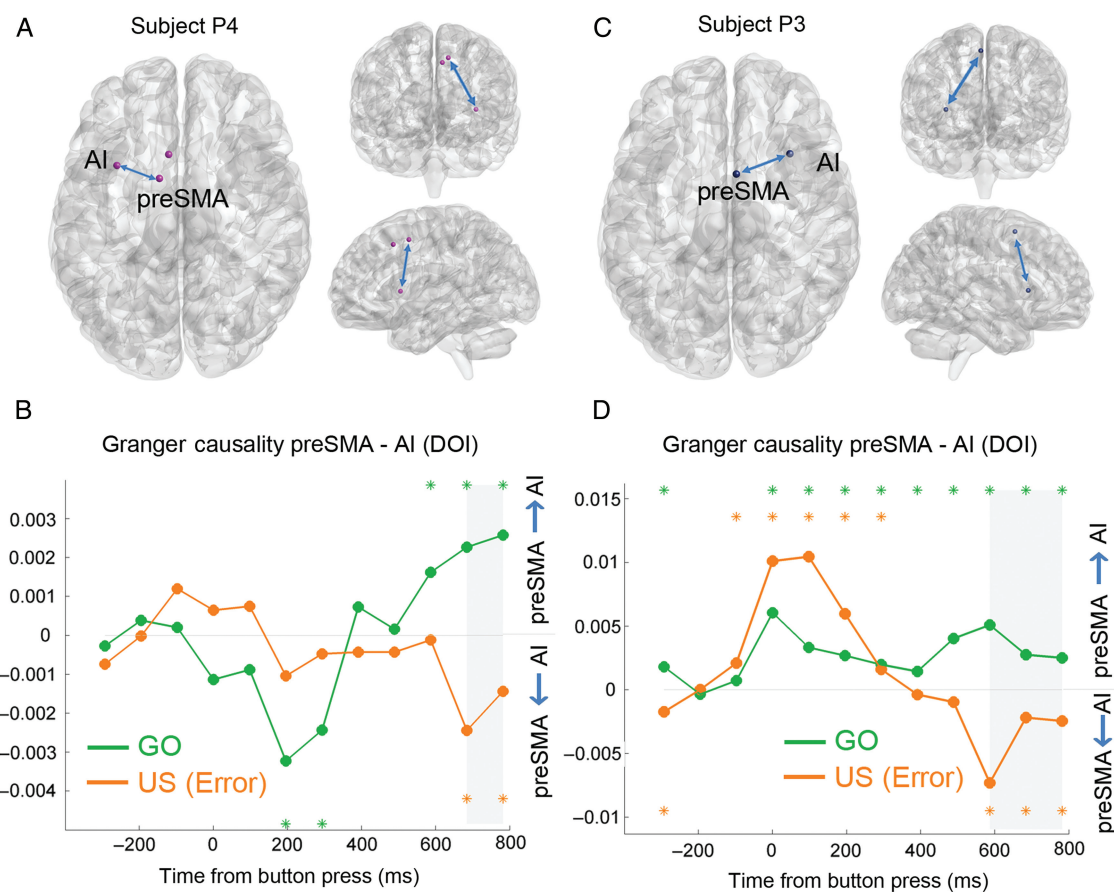
Figure 7. GC between AI and preSMA following errors and correct responses. (A) Anatomical locations of AI and preSMA recording sites used for GC analysis in patient P4. (B) Direction of influence (DOI) measured as the net difference between GC measured in both directions between preSMA and AI. The blue arrows indicate that positive DOI values represent a predominant flow from preSMA to AI, while negative DOI values represent a predominant flow from AI to preSMA. The shaded areas depict the time window for which DOI was both significantly positive in one condition and significantly negative in the other condition (i.e., statistically significant directionality reversal, $P < 0.05$). (C) Same as in panel (A) but for patient P3. (D) Same as in panel (B) but for patient P3.

dorsomedial prefrontal cortex. Our findings provide direct electrophysiological evidence for a leading role of AI within the saliency-error-monitoring networks.

Disentangling pure saliency-driven from pure error-monitoring processes is an intricate endeavor and is only possible up to a given extent. Realizing that one has committed an error is, in most contexts, a highly relevant event that may induce emotional reactions and may also require behavioral adaptation. The error-monitoring network and the salience network are thereby inherently related to one another. Converging evidence suggests that the activity of the Insula and ACC nodes of the salience network is modulated by the extent of "subjective salience", which may be cognitive, homeostatic, or emotional (Damasio et al. 2000; Critchley et al. 2004; Naqvi et al. 2007; Seeley et al. 2007; Craig 2009; Shenhav et al. 2013). So, to which extent do our findings allow us to distinguish between neuronal processes specifically mediating error-monitoring or, more generally representing salient events? To address this question, we first discuss how our results relate to an objective form of saliency (such as the presence of an unexpected stimulus, e.g., the mere appearance of the STOP signal) and then we separately discuss "subjective saliency", which may also be related to interoceptive processes. First of all, note that the hypothesis that AI activity modulations are purely driven by salient task-relevant stimuli cannot fully explain our results because, if this were the case, we would have

found similar responses for successful and unsuccessful inhibition in AI when performing an analysis time-locked to the stop cue (defacto an external salient signals). Our data are not consistent with this idea since AI exhibited stronger responses on error trials, in addition to an early increase of activity induced by the salient cues. This electrophysiological pattern may support an account of AI activity that would incorporate over time the effects of stimulus saliency and error-related signals (such as error-awareness).This interpretation is consistent with several non-invasive neuroimaging studies that either show similar functional responses in AI using an identical paradigm (Ramautar et al. 2006; Sharp et al. 2010) or that report a similar gradient of response within AI with higher BOLD responses when contrasting error to novelty processes (Wessel et al. 2012) or conscious to unconscious errors (Ullsperger et al. 2010; Klein et al. 2013; Charles et al. 2014).

Distinguishing pure error-related components from subjective saliency effects that arise in the context of error awareness is by definition a much more complex endeavor. For example, we cannot exclude the possibility that error-related activity in AI may reflect a rapid subjective saliency signal specifying to the other brain regions that more attention is required after errors to optimize task performance. In line with this hypothesis, patients' heart rate decreased after errors, as expected (Wessel et al. 2011). However, we failed to find any clear relationship

between AI activity (measured in the BGA, beta band or evoked potentials) and this heart rate decrease. This was slightly surprising since AI direct electrical stimulation has been shown to induce changes in heart rate (Oppenheimer et al. 1992). Such findings are closely related to the hypothesis that one functional role of AI is to signal an enhanced subjective feeling of such interoceptive signals that may contribute to our sense of agency (Craig 2009).

We focused on broadband gamma as a marker of neural activity in AI because of its robust relationship with BOLD and extracellular neuronal activity (Lachaux et al. 2007; Jerbi et al. 2009; Manning et al. 2009). However, further analyses revealed an unexpected and very specific increase of beta band power in addition to BGA within AI. This extends to AI previous findings from electrophysiological studies that also found an increase of beta band power during executive control in the inferior frontal gyrus and in the subthalamic nucleus (Swann et al. 2009; Bastin et al. 2014). The pattern of event-related responses was less consistent across participants (and with previous reports) in particular in terms of the ERP waveform. Discrepancies with the literature may be explained by a number of in part subtle differences. The established ERN waveforms are typical of scalp-EEG recordings, direct recordings in ACC may differ. In depth SEEG signals, ERP waveforms are also sensitive to the choice of the reference site (often taken in white matter) and to the re-referencing procedure (bipolar vs. monopolar) if any. Subtle differences in experiment design can also underlie differences in the properties of error-related negativity. Moreover, noninvasive and invasive investigations, both in human and nonhuman primates, have shown a great degree of variability in the observed ERN waveforms in error-related analyses (Godlove et al. ; Reinhart et al. ; Emeric et al. 2008; Jung et al. 2010; Buzsaki and Wang 2012; Bonini et al. 2014). That said, our ERP results clearly dissociated errors from correct trials in AI and this encoding of error-related signals occurred earlier in AI than in ACC. Our data also suggest that the temporal order of error-related ERP onset across structures provides a better match with the GC findings than an account based on the sequence of error-related ERP peak latencies.

Previous sEEG explorations of the performance and error-monitoring system in humans have sought to infer the dynamics and hierarchy within the network by investigating the temporal order of activation of its nodes (Jung et al. 2010; Bonini et al. 2014). Furthermore, only few neuroimaging studies have explored the dynamics of the error-monitoring circuitry by explicitly assessing the directionality of interactions among the nodes of the network. GC has been applied to fMRI in a Stop-Signal-Task to investigate the origin of the activation in the ventrolateral prefrontal cortex (VLPFC) that occurs during posterror slowing (Ide and Li 2011). The results provide evidence for the functional role of a cerebello-thalamo-cortical pathway in error-related cognitive control. A parallel stream of research based on Bayesian approaches using dynamic causal modeling (DCM) has also provided interesting insights into the effective connectivity within the error processing networks (Ham et al. 2013). The results of the latter study are consistent with our findings that AI plays a driving role in the error-processing network. Using DCM, Ham et al. (2013) investigated the causal interactions of the 3 key cortical nodes within the SN, which are dACC, left and right AI. They found that input onto the SN came through the right AI, and that this was the only node showing intrinsic connectivity to the other 2 parts of the SN. Both these fMRI findings and the direct electrophysiological data presented here are in line with the hypothesis that the AI acts as a "cortical-outflow hub" regulating activity in other brain regions (Sridharan et al. 2008; Menon and Uddin 2010). An important limitation of fMRI-based explorations

of network dynamics is that the hemodynamic response is insensitive to rapid interactions between network nodes. Applying DCM to fMRI data assesses the longer-lasting time-varying properties of the network, rather than its instantaneous interaction dynamics. Our ability to detect fast and transient reversals of causality (between AI on one hand, and ACC and PreSMA on the other) illustrates the added value of the high temporal and spectral resolution obtained with sEEG data and highlights the complementarities between different modalities.

The reversal of information flow between AI and dmPFC on erroneous button presses was revealed here using time-domain GC. The link to the local BGA findings is thus not straightforward.

The very early onset of error-monitoring signals in AI as captured by broadband power modulations, and the reversal of its directionality relationship with ACC and preSMA, as revealed by GC, are consistent with a leading role of AI over dmPFC after errors (Sridharan et al. 2008; Ham et al. 2013). Our findings are also in agreement with predictions from a recent model of ACC function (Shenhav et al. 2013) that postulates that AI sends motivationally relevant (unexpected and erroneous) signals to ACC to be used to adjust cognitive control signals. In parallel, AI input to preSMA might enable preSMA to adjust the ongoing speed-accuracy trade-off in the motor system to switch from a fast/habitual response to more controlled/slower behavior (Isoda and Hikosaka 2007). Note however that, although statistically significant, the reversals observed here between AI and preSMA were slightly less consistent across participants than the clear-cut reversals observed between AI and dACC.

In summary, using rare intracerebral data simultaneously acquired from key nodes of the error-monitoring network, our study provides direct electrophysiological evidence in humans supporting the view that AI is the input to the brain's error-monitoring system. The error-related signals computed by AI may be used as priors by ACC and preSMA to compute more elaborate signals that mediate learning and adaptive behavior.

Supplementary Material

Supplementary material can be found at: http://www.cercor.oxfordjournals.org/.

Funding

Notes

References

Aguera PE, Jerbi K, Caclin A, Bertrand O. 2011. ELAN: a software package for analysis and visualization of MEG, EEG, and LFP signals. Comput Intell Neurosci. 2011:158970.

Akaike H. 1974. A new look at the statistical model identification IEEE Trans Automatic Control. 19:716–723.

Amiez C, Champod AS, Wilson CR, Procyk E, Petrides M. 2015. A unilateral medial frontal cortical lesion impairs trial and error learning without visual control. Neuropsychologia. 75:314–321.

Amiez C, Neveu R, Warrot D, Petrides M, Knoblauch K, Procyk E. 2013. The location of feedback-related activity in the midcingulate cortex is predicted by local morphology. J Neurosci. 33:2217–2228.

Bastin J, Polosan M, Benis D, Goetz L, Bhattacharjee M, Piallat B, Krainik A, Bougerol T, Chabardes S, David O. 2014. Inhibitory control and error monitoring by human subthalamic neurons. Transl Psychiatry. 4:e439.

Bonini F, Burle B, Liegeois-Chauvel C, Regis J, Chauvel P, Vidal F. 2014. Action monitoring and medial frontal cortex: leading role of supplementary motor area. Science. 343:888–891.

Buzsaki G, Wang XJ. 2012. Mechanisms of gamma oscillations. Annu Rev Neurosci. 35:203–225.

Charles L, King JR, Dehaene S. 2014. Decoding the dynamics of action, intention, and error detection for conscious and subliminal stimuli. J Neurosci. 34:1158–1170.

Craig AD. 2009. How do you feel–now? The anterior insula and human awareness. Nat Rev Neurosci. 10:59–70.

Critchley HD, Wiens S, Rotshtein P, Ohman A, Dolan RJ. 2004. Neural systems supporting interoceptive awareness. Nat Neurosci. 7:189–195.

Damasio AR, Grabowski TJ, Bechara A, Damasio H, Ponto LL, Parvizi J, Hichwa RD. 2000. Subcortical and cortical brain activity during the feeling of self-generated emotions. Nat Neurosci. 3:1049–1056.

Debener S, Ullsperger M, Siegel M, Fiehler K, von Cramon DY, Engel AK. 2005. Trial-by-trial coupling of concurrent electroencephalogram and functional magnetic resonance imaging identifies the dynamics of performance monitoring. J Neurosci. 25:11730–11737.

Dehaene S, Posner MI, Tucker DM. 1994. Localization of a neural system for error detection and compensation. Psychol Sci. 5:303–305.

Dupont S, Bouilleret V, Hasboun D, Semah F, Baulac M. 2003. Functional anatomy of the insula: new insights from imaging. Surg Radiol Anat. 25:113–119.

Gehring WJ, Goss B, Coles MGH, Meyer DE, Donchin E. 1993. A neural system for error detection and compensation. Psychol Sci. 4:385–390.

Genovese CR, Lazar NA, Nichols T. 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage. 15:870–878.

Granger CWJ. 1969. Investigating causal relations by econometric models and cross-Spectral methods. Econometrica. 37:424–438.

Ham T, Leff A, de Boissezon X, Joffe A, Sharp DJ. 2013. Cognitive control and the salience network: an investigation of error processing and effective connectivity. J Neurosci. 33:7091–7098.

Holroyd CB, Coles MG. 2002. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. Psychol Rev. 109:679–709.

Ide JS, Li CS. 2011. A cerebellar thalamic cortical circuit for error-related cognitive control. Neuroimage. 54:455–464.

Isoda M, Hikosaka O. 2007. Switching from automatic to controlled action by monkey medial frontal cortex. Nat Neurosci. 10:240–248.

Ito S, Stuphorn V, Brown JW, Schall JD. 2003. Performance monitoring by the anterior cingulate cortex during saccade countermanding. Science. 302:120–122.

Jerbi K, Ossandon T, Hamame CM, Senova S, Dalal SS, Jung J, Minotti L, Bertrand O, Berthoz A, Kahane P, et al. 2009. Task-related gamma-band dynamics from an intracerebral perspective: review and implications for surface EEG and MEG. Hum Brain Mapp. 30:1758–1771.

Jung J, Jerbi K, Ossandon T, Ryvlin P, Isnard J, Bertrand O, Guenot M, Mauguiere F, Lachaux JP. 2010. Brain responses to success and failure: Direct recordings from human cerebral cortex. Hum Brain Mapp. 31:1217–1232.

Klein TA, Endrass T, Kathmann N, Neumann J, von Cramon DY, Ullsperger M. 2007. Neural correlates of error awareness. Neuroimage. 34:1774–1781.

Klein TA, Ullsperger M, Danielmeier C. 2013. Error awareness and the insula: links to neurological and psychiatric diseases. Front Hum Neurosci. 7:14.

Kurth F, Zilles K, Fox PT, Laird AR, Eickhoff SB. 2010. A link between the systems: functional differentiation and integration within the human insula revealed by meta-analysis. Brain Struct Funct. 214:519–534.

Lachaux JP, Fonlupt P, Kahane P, Minotti L, Hoffmann D, Bertrand O, Baciu M. 2007. Relationship between task-related gamma oscillations and BOLD signal: new insights from combined fMRI and intracranial EEG. Hum Brain Mapp. 28:1368–1375.

Lachaux JP, Rudrauf D, Kahane P. 2003. Intracranial EEG and human brain mapping. J Physiol Paris. 97:613–628.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature. 412:150–157.

Manning JR, Jacobs J, Fried I, Kahana MJ. 2009. Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. J Neurosci. 29:13613–13620.

McCormick LM, Ziebell S, Nopoulos P, Cassell M, Andreasen NC, Brumm M. 2006. Anterior cingulate cortex: an MRI-based parcellation method. Neuroimage. 32:1167–1175.

Menon V, Uddin LQ. 2010. Saliency, switching, attention and control: a network model of insula function. Brain Struct Funct. 214:655–667.

Naqvi NH, Rudrauf D, Damasio H, Bechara A. 2007. Damage to the insula disrupts addiction to cigarette smoking. Science. 315:531–534.

Niazy RK, Beckmann CF, Iannetti GD, Brady JM, Smith SM. 2005. Removal of FMRI environment artifacts from EEG data using optimal basis sets. Neuroimage. 28:720–737.

Niessing J, Ebisch B, Schmidt KE, Niessing M, Singer W, Galuske RA. 2005. Hemodynamic signals correlate tightly with synchronized gamma oscillations. Science. 309:948–951.

Oppenheimer SM, Gelb A, Girvin JP, Hachinski VC. 1992. Cardiovascular effects of human insular cortex stimulation. Neurology. 42:1727–1732.

Picard N, Strick PL. 1996. Motor areas of the medial wall: a review of their location and functional activation. Cereb Cortex. 6:342–353.

Ramautar JR, Slagter HA, Kok A, Ridderinkhof KR. 2006. Probability effects in the stop-signal paradigm: the insula and the significance of failed inhibition. Brain Res. 1105:143–154.

Ray S, Maunsell JH. 2011. Network rhythms influence the relationship between spike-triggered local field potential and functional connectivity. J Neurosci. 31:12674–12682.

Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD. 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. J Neurosci. 27:2349–2356.

Seth AK. 2005. Causal connectivity of evolved neural networks during behavior. Network. 16:35–54.

Seth AK. 2010. A MATLAB toolbox for Granger causal connectivity analysis. J Neurosci Methods. 186:262–273.

Sharp DJ, Bonnelle V, De Boissezon X, Beckmann CF, James SG, Patel MC, Mehta MA. 2010. Distinct frontal systems for response inhibition, attentional capture, and error processing. Proc Natl Acad Sci USA. 107:6106–6111.

Shenhav A, Botvinick MM, Cohen JD. 2013. The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron. 79:217–240.

Sridharan D, Levitin DJ, Menon V. 2008. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. Proc Natl Acad Sci USA. 105:12569–12574.

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. 2012. 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.

Ullsperger M, Harsay HA, Wessel JR, Ridderinkhof KR. 2010. Conscious perception of errors and its relation to the anterior insula. Brain Struct Funct. 214:629–643.

Ullsperger M, von Cramon DY. 2001. Subprocesses of performance monitoring: a dissociation of error processing and response competition revealed by event-related fMRI and ERPs. Neuroimage. 14:1387–1401.

von Stein A, Sarnthein J. 2000. Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. Int J Psychophysiol. 38:301–313.

Wessel JR, Danielmeier C, Morton JB, Ullsperger M. 2012. Surprise and error: common neuronal architecture for the processing of errors and novelty. J Neurosci. 32:7528–7537.

Wessel JR, Danielmeier C, Ullsperger M. 2011. Error awareness revisited: accumulation of multimodal evidence from central and autonomic nervous systems. J Cogn Neurosci. 23:3021–3036.