Interactions between frontal and posterior oscillatory dynamics support adjustment of stimulus processing during reinforcement learning

Irene van de Vijver [a,b,*], Joram van Driel [a,c], Arjan Hillebrand [d], Michael X Cohen [a]

[a] University of Amsterdam, Department of Psychology, Amsterdam, The Netherlands
[b] Radboud University, Behavioural Science Institute, Nijmegen, The Netherlands
[c] Vrije Universiteit, Department of Cognitive Psychology, Amsterdam, The Netherlands
[d] Department of Clinical Neurophysiology and Magnetoencephalography Center, VU University Medical Center, Amsterdam, The Netherlands

ABSTRACT

Reinforcement learning (RL) in humans is subserved by a network of striatal and frontal brain areas. The electrophysiological signatures of feedback evaluation are increasingly well understood, but how those signatures relate to the use of feedback to guide subsequent behavioral adjustment remains unclear. One mechanism for post-feedback behavioral optimization is the modulation of sensory processing. We used source-reconstructed MEG to test whether feedback affects the interactions between sources of oscillatory activity in the learning network and task-relevant stimulus-processing areas. Participants performed a probabilistic RL task in which they learned associations between colored faces and response buttons using trial-and-error feedback. Delta-band (2–4 Hz) and theta-band (4–8 Hz) power in multiple frontal regions were sensitive to feedback valence. Low and high beta-band power (12–20 and 20–30 Hz) in occipital, parietal, and temporal regions differentiated between color and face information. Consistent with our hypothesis, single-trial power-power correlations between frontal and posterior-sensory areas were modulated by the interaction between feedback valence and the relevant stimulus characteristic (color versus identity). These results suggest that long-range oscillatory coupling supports post-feedback updating of stimulus processing.

1. Introduction

During reinforcement learning (RL), experienced associations between stimuli (or situations), actions, and outcomes guide behavioral adaptation. Optimal performance requires stimuli to become associated with the action that yields the best outcome. Whereas outcome evaluation in the brain has been extensively documented (see, e.g., Garrison et al., 2013; Lee et al., 2012; Maia, 2009), how this affects the processing of, and association between, stimuli and actions has received less empirical investigation. In the current study we investigated how feedback evaluation interacts with stimulus representations to optimize stimulus processing for future behavior.

Previous studies have outlined a reward learning circuit consisting of midbrain dopaminergic structures, the striatum, orbitofrontal cortex (OFC), and medial frontal cortex (MFC), which is regulated by structures such as (dorso)lateral prefrontal cortex (dlPFC; for reviews see, e.g., Haber and Knutson, 2010; Maia, 2009; Rushworth et al., 2011). Stimulus representations, and stimulus-action associations, are most likely stored

and (re)activated in task-relevant cortical areas (Nyberg et al., 2001, 2000; Rösler et al., 2007; Slotnick, 2009). Thus, behavioral improvement requires cooperation between the learning circuit and other task-relevant brain areas. Indeed, fMRI research has demonstrated that perception-related activity in visual areas is adjusted with learning (Luft et al., 2015). Additionally, functional connectivity between the striatum and task-specific visual areas increases during rewards (Schiffer et al., 2014), and functional connectivity between the striatum, frontal cortex, motor cortex and visual processing areas changes over the course of learning (den Ouden et al., 2010; Horga et al., 2015).

An important question remains how this modulation of stimulus representations is realized. Changes in oscillatory activity within brain regions, and synchronous fluctuations in oscillatory activity in spatially distinct brain regions, may provide underlying mechanisms (Akam and Kullmann, 2010; Fries, 2005; Gregoriou et al., 2009; Wang et al., 2010; Womelsdorf et al., 2007). Frontal theta-band oscillations (4–8 Hz) support the processing of negative feedback and implementation of behavioral adjustments (Cohen, Wilmes and van de Vijver, 2011), whereas

frontal beta-band oscillations (15–30 Hz) may signal behaviourally relevant rewards (Marco-Pallarés et al., 2015). Additionally, theta-band connectivity between MFC, lPFC, and motor cortex is higher after negative than positive feedback (Cavanagh et al., 2010; Luft et al., 2013; van de Vijver et al., 2011). Oscillatory coupling between frontal and posterior cortical areas shows a stronger increase after errors than after correct responses during both cognitive control tasks (Cohen et al., 2009; Cohen and van Gaal, 2013; van Driel et al., 2012) and working memory tasks (Cashdollar et al., 2009; Palva et al., 2010; Salazar et al., 2012; Sauseng et al., 2005). Thus, oscillations play a role in feedback evaluation, and provide an important mechanism for long-range communication between frontal and posterior cortical areas when adjustments in behavior or maintained information are required.

We therefore hypothesized that feedback affects the interactions between sources of oscillatory activity in the frontal learning circuit and task-relevant stimulus processing areas. To investigate whether feedback specifically influenced task-relevant visual brain areas, we designed a RL task with colored face stimuli. The colors were informative to learn the correct responses in one half of the experiment, and the face identities in the other half. Because the visual stimuli were identical in both situations, differences in brain activity and connectivity could be explained only by the associations that were being learned, and, thus, the visual feature that was attended. We expected a double dissociation, such that post-feedback connectivity between the learning network and color-processing areas would increase when color stimuli were informative for learning, whereas connectivity between this network and face-processing areas would increase when face stimuli were informative. We measured MEG while participants performed the learning task, and estimated the sources of oscillatory brain dynamics using beamforming techniques. Interactions between sources were reconstructed using power-power correlations. This allowed us to separately investigate face- and color-processing areas, and their interactions with the learning network.

2. Materials and methods

2.1. Participants

13 young adults from the University of Amsterdam campus (4 male, 1 left-handed), ranging in age from 18 to 30 years (M 22.2, SD 3.51), participated in three sessions on separate days: (1) a behavioral session to practice the RL task, (2) an MEG session at the VU University Medical Center during which they performed the learning task and a localizer task, and (3) an MRI session at the Amsterdam Spinoza Center for Neuroimaging to acquire structural T_1 and DTI scans. The participants underwent telephone screening according to standard MEG and MRI exclusion criteria from the centers where the scans were acquired. Participants received a financial compensation of €42.50. Task performance was not rewarded. All procedures were completed in compliance with relevant laws and institutional guidelines and were approved by the local ethics committees (VUmc Medical Ethical committee for MEG session, University of Amsterdam psychology Ethics Committee for behavioral and MRI sessions).

2.2. Tasks and behavior

2.2.1. Probabilistic learning task

Participants performed a probabilistic RL task in which they learned the correct associations between six stimuli and two response buttons by trial-and-error (Fig. 1a). On each trial, a stimulus was presented until the participant pressed a button (maximum 1200 ms). After a delay of 500 ms, feedback was displayed for 400 ms. An inter-trial interval of 1500 ms separated successive trials. Stimuli consisted of six neutral faces of young males (Ebner, 2008), which were colored red, blue, green, yellow, purple, or orange (Fig. 1b). Per block, participants were instructed to learn either the associations between the faces and the
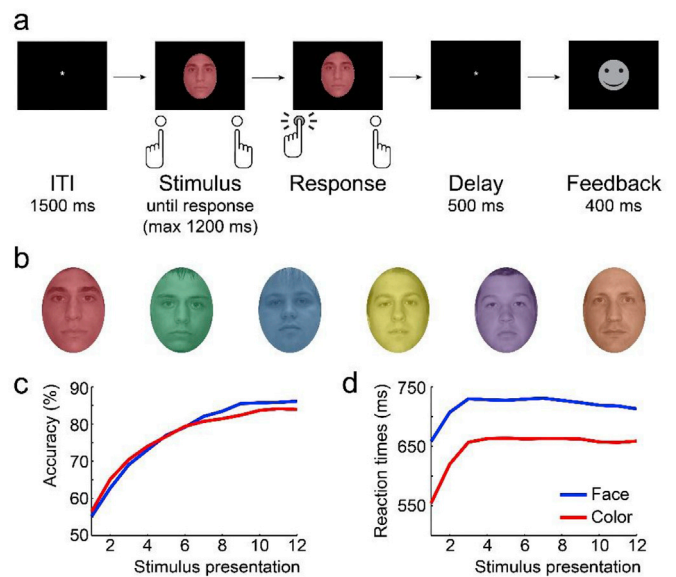

Fig. 1. Reinforcement learning task and behavioral performance. (a) Sequence of events in an example trial. (b) The six different facial stimuli and six colors they could be presented in. During the learning task, all combinations of faces and colors were used. (c) Learning accuracy as percentage correct responses averaged over stimuli per condition and presentation blocks. (d) Reaction times over trials averaged over stimuli per condition and presentation blocks.

response buttons, or between the colors and the response buttons. Thus, although the stimuli were identical in the color and face blocks, the type of associations that were learned differed. We will refer to this as the color and face dimensions, respectively.

In each block, all six stimuli (in the relevant dimension) were presented 12 times. Per six trials, all stimuli were presented once in random order. All stimuli in the relevant dimension were combined with all options on the other dimension twice per block, once in each half. Thus, in a face block, each face was presented in each color twice. Per block, three stimuli were mapped to the left, and three to the right response button. Participants were informed that each block featured new mappings between stimuli and responses. Participants performed eight face and eight color blocks in alternating order. The dimension of the first block was counterbalanced over participants.

Feedback consisted of a smiley or a frowny face as positive or negative feedback, respectively, or the words 'too late' if no response was recorded during stimulus presentation. Participants received invalid feedback on 20% of the trials (negative feedback after a correct response or positive feedback after an incorrect response). In the MEG session, responses were given with the index fingers on two MEG-compatible button boxes. Before the learning task, participants adjusted the luminance of all colors on all faces, making them subjectively equiluminant. In the practice session, participants received instructions and practiced the task until their average accuracy was at least 60% for both stimulus dimensions. Every participant practiced at least two blocks per stimulus dimension.

2.2.2. Localizer task

During the MEG session, the learning task was preceded by a 1-back task, which was included to localize time-space-frequency clusters that differentiated between face and color information. The stimuli consisted of six gray-scaled faces, six gray-scaled houses, six colored patches, and six isoluminant gray patches. The faces and colors were the same as in the learning task, but whereas these features were combined in the learning task stimuli, here they were presented separately. Participants had to indicate whether each stimulus did or did not match the previous one by pressing a button with their right index or middle finger, respectively

(mappings counterbalanced over participants). Stimuli were presented for 1500 ms regardless of the response time. If no response was detected during stimulus presentation, the words 'too late' were presented for 400 ms. An inter-trial interval of 1500 ms separated the trials. All stimuli were presented five times in quasi-randomized order including 20% stimulus repetitions.

2.2.3. Behavioral analyses

To verify that participants learned the correct associations, we compared accuracy (defined as percentage correct) averaged over blocks, stimuli, and trials for the face and color dimensions separately against chance level performance (50%) using t-tests. To investigate whether participants continued learning over stimulus presentations and whether learning differed between the face and color dimensions, accuracy and reaction times (RTs) were averaged over stimuli and blocks and entered into separate ANOVAs with the factors Bin (trials 1–3, 4–6, 7–9, 10–12) and Stimulus dimension (Face, Color). Greenhouse-Geisser corrections were applied when appropriate, but uncorrected degrees of freedom are reported.

2.3. MRI data acquisition and preprocessing

Magnetic resonance imaging data were acquired on a Philips 3T scanner with a 32-channel array head coil. An anatomical T_1-weighted image (220 slices of 240×240 voxels, voxel size $1 \times 1 \times 1$ mm, TR = 8.219 s, TE = 3.79 s) was obtained to create the head model for the beamformer analysis (see below). Diffusion-weighted data and fMRI data recorded during two localizer tasks were also obtained but are not discussed in the current report. The T_1 scans were converted to a box-shaped image in mrilab (Elekta Neuromag Oy, version 1.7.25), and coregistered with the MEG data using the digital surface points recorded before the MEG scan (see below). To this end, we applied surface matching software developed by one of the authors (AH), which results in an estimated co-registration accuracy of approximately 4 mm (Whalen et al., 2008). For one participant MRI scans could not be acquired. Instead we identified which of the T_1 scans of the other participants best fitted the outline of this participant's scalp (as recovered from the digital surface points) and used this T_1 scan. For group-level analyses, all T_1 scans were normalized to the MNI standard brain with SPM8 (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/).

2.4. MEG data acquisition and sensor-level analyses

2.4.1. Data acquisition and preprocessing for sensor-level analyses

MEG data were recorded with a 306-channel whole-head MEG system (Elekta Neuromag Oy, Helsinki, Finland) in a magnetically shielded room (Vacuumschmelze, Hanau, Germany) while participants were in a supine position. This system consists of 102 triple-sensor detector units, each containing two (orthogonal) planar gradiometers and a magnetometer. The head position relative to the sensors was recorded continuously using four head-localization coils. Before entering the scanner, the head-localization coil positions and the outline of the participant's scalp (~500 points) were digitized with a 3D digitizer (Fastrack, Polhemus, Colchester, VT, USA). MEG data were recorded at 1250 Hz. An antialiasing filter of 410 Hz and a high-pass filter of 0.1 Hz were applied online.

Other internal and external noise and artifacts were removed offline with the temporal extension of Signal Space Separation (tSSS) in MaxFilter (Elekta Neuromag Oy, version 2.2.10; Taulu and Simola, 2006; Taulu and Hari, 2009), using a subspace correlation limit of 0.9 and a sliding window of 10 s. For sensor-level analyses, tSSS included a transformation of all data to the head position of the participant with the most optimal position in the scanner (based on visual inspection), to equalize sensor locations across participants. Malfunctioning channels were identified by visual inspection, and excluded and recalculated with tSSS

(M 5.23, SD 2.39 channels).

All subsequent processing steps for sensor-level analyses of the learning task data were performed in Matlab (The Mathworks, Natick, MA, USA) using EEGLAB functions (Delorme and Makeig, 2004). After the data were converted to EEGLAB format, they were filtered with a 0.5 Hz high-pass and a 350 Hz low-pass two-way least-squares FIR filter, downsampled to 625 Hz, epoched from −1.5–4.0 s around stimulus onset, and baseline corrected using a −400 to −200 ms pre-stimulus baseline. Epochs containing artifacts were manually selected and removed (M 7.27%, SD 4.20% of trials, range 1.66–17.36%). The last epoch of each block was excluded because some recordings ended prematurely. After independent component analysis (standard EEGlab ICA routine, type jader), components containing blinks were manually selected and removed. This procedure was performed separately for the magnetometers and the two sets of gradiometers, to acquire the most specific blink components (M 1.15, SD 0.37, range 0–2 components).

2.4.2. Reinforcement learning task trial selection

Only trials in which participants received valid feedback were included. Trials were sorted into four conditions: 'face positive' (face stimulus, positive feedback; FP), 'face negative' (FN), 'color positive' (CP), and 'color negative' (CN). Large differences in trial counts between conditions can influence the estimation of oscillatory power and connectivity (Cohen, 2014). Therefore, we determined the number of trials in the condition with the lowest trial count. If the trial count in any other condition exceeded this number by more than 20 trials, trials were randomly selected from that condition until this threshold was reached. The final trial count was on average 106.2 (SD 25.4), 95.1 (29.3), 106.2 (25.4), and 93.4 (23.7) for the FP, FN, CP, and CN conditions, respectively. The smallest number of trials in any participant in any condition was 42.

2.4.3. Sensor-level analyses and statistics

Custom-written Matlab routines were used to extract time-frequency dynamics from the signals of the two separate sets of gradiometers (magnetometer data were not used for sensor-level analyses). First, stimulus-locked epochs were convolved with a family of complex Morlet wavelets, defined as Gaussian-windowed complex sine waves: $e^{i2\pi ft} e^{-t^2/2\sigma^2}$, where i is the complex operator, t is time, f is frequency, which increased from 1 to 50 Hz in 40 logarithmically spaced steps, and σ defines the wavelet width in each frequency band. This parameter was set to $s/(2\pi f)$, with s increasing from 3 to 8 cycles over the 40 frequencies. This is in line with procedures used in comparable studies (see, e.g., Canavagh et al., 2012a; Luft et al., 2013; van de Vijver et al., 2011), but allows a higher temporal resolution at lower frequencies and a more stable estimation at higher frequencies (Cohen, 2014). Next, per epoch we separately extracted stimulus-locked and feedback-locked estimates of power from the resulting complex signal Z (power time series: p(t) = real $[Z(t)]^2$ + imag $[Z(t)]^2$). Finally, per sensor, the feedback-locked power values from the two gradiometers were summed and converted to a decibel (dB) scale (10*log10 [power/baseline]), using a stimulus-locked −500 to −200 ms pre-stimulus baseline.

In EEG studies of similar RL processes, theta-band oscillations over MFC peak between 200 and 500 ms after feedback presentation (see, e.g., Cavanagh et al., 2010; van de Vijver et al., 2011). We therefore focused our sensor-level analyses on this time window. Because oscillatory brain dynamics can manifest at different scalp locations in MEG compared to EEG research (Ahlfors et al., 2010), we selected sensors based on the topography of condition-averaged post-feedback theta power. At these sensors (over two medial frontal and two lateral frontal areas, see Fig. 2a), we compared average power in a window of 4–8 Hz and 200–500 ms post-feedback with ANOVAs including the factors Feedback valence (positive, negative feedback) and Stimulus dimension (Face, Color).

Fig. 2. Effect of feedback valence on theta power at the sensor level. (a) Post-feedback condition-averaged theta-power increase (300–400 ms) indicated two medial frontal sensors and two lateral frontal sensor clusters showing large increases in theta power. (b) Feedback-locked condition-averaged power at the four sensors/clusters (Llfc sensors: 222/3, 232/3, 322/3, 342/3; Rlfc sensors: 1232/3, 1312/3, 1322/3, 1342/3). The dashed box indicates the time-frequency window (4–8 Hz, 200–500 ms) used for the selection of the data displayed in (c) and for statistics. (c). Theta power was significantly larger after negative compared to positive feedback at the lateral frontal sensor clusters, but not at the medial frontal sensors. Error bars represent the SEM (Llfc = left lateral frontal cortex, Rlfc = right lateral frontal cortex, 422/3 = sensors 422 and 423 combined (same rule applied to the names of the other sensors); neg = negative, pos = positive, FB = feedback).

2.5. MEG data analysis: source-level

2.5.1. Preprocessing for source-level analyses

We excluded the previously identified malfunctioning channels from the data and removed artifacts and noise with tSSS using the same settings as for the sensor-level analyses, but without transforming to a standard head position.

All subsequent preprocessing steps for source-level analyses were performed in Matlab, using Fieldtrip functions (version 20140723; Oostenveld et al., 2011). After the data were converted to Fieldtrip format, preprocessing included the application of a 0.5 Hz - 350 Hz band-pass two-way Butterworth filter, downsampling to 625 Hz, the creation of −1.5–4.0 s stimulus-locked epochs, and baseline correction using a −400 to −200 ms pre-stimulus baseline. The same trials were removed as for the sensor-level analyses. Note that for the source-level analyses we did not remove eye blinks, as beamforming (see below), particularly in combination with tSSS (Hillebrand et al., 2013), already reduces artifacts (e.g. Adjamian et al., 2009, but see Hipp and Siegel, 2015). This was confirmed by control analyses further examining the possible influence of eye blinks on the beamformer results (reported in the Supplementary Material). Moreover, removing all trials containing post-feedback blinks would have resulted in severely reduced trial counts in multiple participants.

The same preprocessing was applied to the localizer data for the source-level analyses, but stimulus-locked epochs now ranged from −1.5–3.0 s. On average, 5.45% of trials (SD 4.89%, range 0–15.83%) were removed. The final trial count was on average 28.5 (SD 1.3), 28.2 (2.0), 28.7 (1.3), and 28.0 (2.2) for the face, house, color, and gray conditions, respectively. The smallest number of trials in any participant in the face and color conditions was 26 (the house and gray conditions were not included in subsequent analyses).

2.5.2. LCMV beamforming to reconstruct source-level power

All beamforming-related processing steps were performed using the Fieldtrip toolbox (Oostenveld et al., 2011). For each participant, we first estimated the forward model, the projection of all voxels onto all sensors (now including both magnetometers and gradiometers). Specifically, we constructed a semi-realistic single-shell head model (according to Nolte, 2003) based on the individual's T_1 scan, which was divided into 5 × 5x5 mm voxels (M 11371.1, SD 982.3, range 10302-13089 voxels). At each voxel the lead fields were estimated for an equivalent current dipole, for 3 orthogonal orientations.

Next, we applied a linear constrained minimum variance (LCMV)

beamformer to estimate the spatial distribution of the electrical brain sources underlying the measured magnetic fields. Based on the measured data and the forward model, the beamformer estimates the relative contribution (or weight) of the variance in the signal at each sensor to the activity at a particular location, and sequentially reconstructs the signal at each voxel using these voxel-specific weights. We applied the beamformer separately for each frequency band. Because the covariance matrices that serve as the sensor-level input to the LCMV beamformer are computed in the time domain on band-pass filtered data, a smaller time window can be used to compute the covariance compared to frequency-domain beamformers such as Dynamical Imaging of Coherent Sources, without constraining the frequency resolution (but see Brookes et al., 2008). This allows a sliding-window approach, where the beamformer is applied on consecutive time windows to assess fluctuations over time (Dalal et al., 2008).

Thus, to acquire a full time-frequency representation, we repeatedly applied the beamformer from −800 to 1000 ms around the stimulus (localizer) or feedback (learning task) in 50 ms steps and for seven frequency bands: delta (2–4 Hz), theta (4–8 Hz), alpha (8–12 Hz), low beta (12–20 Hz), high beta (20–30 Hz), low gamma (30–50 Hz), and high gamma (50–80 Hz; see Fig. 3 for a schematic representation of the beamforming procedure). The data were band-pass filtered with the *filtfilt* Matlab function to obtain frequency band-specific signals (Fig. 3a). Note that this analysis approach yields a limited spectral resolution, which does not allow us to make claims about precise frequencies beyond the boundaries we selected a priori. The temporal window of data included in each estimation step to compute the covariance matrix decreased in width with increasing frequency, from 400 ms for the delta band, 300 ms for the theta, alpha, and beta bands, to 200 ms for the gamma bands. Per voxel, only the most prominent of the three estimated weights (orientations) was selected for signal reconstruction. The regularization parameter lambda (of the *ft_sourceanalysis* Fieldtrip function) was set to 1% in order to deal with the reduced rank of the covariance matrix caused by tSSS.

For both tasks, we also estimated baseline power in a pre-stimulus window. This window also decreased in width with increasing frequency, but always ended −100 ms pre-stimulus. The beamformer weights were computed on the combined baseline and task data (Fig. 3b), and then used to extract baseline and task-relevant power separately (Fig. 3c; Gross et al., 2013). This implied that the baseline was recomputed in combination with each time-frequency window of the task data. Finally, reconstructed power in each time-frequency window was converted to dB scale (10*log10 [power/baseline]; Fig. 3d). This correction

Fig. 3. LCMV beamforming procedure and sequence of events in source-level analyses. (a–c) LCMV beamforming procedure: (a) Sensor-level MEG data were first band-pass filtered into seven frequency bands. (b) Per frequency band, the beamformer was applied separately for consecutive, partially overlapping time windows (see 3c). Window length decreased with increasing frequency. Per frequency, the task-data time window was matched by a (non-sliding) baseline window of the same length. (c) I – The beamformer weights were computed per task-data time window. For each computation, the data in the window were combined with the baseline data. II - The calculated weights were then used to reconstruct the task data at the source level in this time window, and also the baseline data at the source level. This implies that a separate source-level baseline dataset was reconstructed for each task-data time window using separate weights. III – Finally, per task-data time window, the baseline data as reconstructed for each separate task-data time window were used to baseline correct those task data using a decibel conversion. (d–f) Sequence of events in source-level analyses: (d) To extract clusters of voxels that were specifically sensitive to face or color information, we compared activation during face versus color stimuli in the localizer task. To extract clusters of voxels that were specifically sensitive to positive or negative feedback, we compared activation during positive versus negative feedback in the learning task. ROIs were defined as clusters of 50 contiguous voxels around the voxel with the peak difference between conditions. (e) Per ROI, post-feedback activation in four conditions was compared with an ANOVA with factors Feedback valence (positive or negative feedback) and Stimulus dimension (face or color informative for learning). (f) Per frontal-posterior ROI pair, activation was correlated over trials for each condition and again compared with an ANOVA with factors Feedback valence and Stimulus dimension (BF = beamformer, ROI = region of interest).

not only allows a more specific focus on post-feedback dynamics, but also helps to avoid a depth bias due to non-uniform projection of white noise.

To determine task-relevant clusters of voxels (see below) we ran four separate beamformers for the face and color trials in the localizer task and the trials featuring positive and negative feedback in the learning task (see Fig. 3 for a schematic representation of the analysis pipeline). We also examined the use of one common set of weights per comparison. To this end, we computed the weights using the combined data from the two conditions in a comparison (faces and colors, or positive and negative feedback) and then computed the power per voxel for the two conditions separately. This procedure did not qualitatively change the resulting clusters, although the spatial extent of the significant difference between conditions tended to be larger with common weights than with condition-specific weights.

2.5.3. Source-level statistics to determine task-relevant clusters

To define task-relevant clusters, we established which clusters of voxels differentiated significantly between face and color stimuli in the localizer task (note that contrasts of faces vs. houses and colored vs. gray patches did not result in significant clusters), and which clusters of voxels differentiated significantly between positive and negative feedback in the learning task (Fig. 3d). Note that although we used the same learning-task data to subsequently test our hypothesis, the contrast used to select task-relevant clusters focused on the main effect of feedback valence whereas our key hypothesis was focused on an interaction between feedback valence and stimulus dimension.

We first normalized the beamformer results to standard MNI space with the Matlab toolbox Nutmeg (Dalal et al., 2004) and applied a mask (based on the FSL Harvard-Oxford atlas, http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/) to exclude voxels in the ventricles. This resulted in a matrix of 14608 $5 \times 5 \times 5$ mm voxels. We only examined beamformer results from 0 to 1000 ms post-stimulus (localizer) or post-feedback (learning task) to decrease the number of statistical comparisons.

For both tasks and for each time-frequency window we applied a three-step procedure to find clusters of voxels showing significant power differences with regard to the comparison of interest (similar to Cohen and Ridderinkhof, 2013). First, we selected voxels demonstrating a significant power change during task performance, using permutation testing. In each of 1000 iterations, condition-averaged power values per voxel were multiplied by −1 for a randomly chosen number of participants and *t*-values were computed (average baseline-corrected dB-values against zero; multiplication of a condition difference by −1 is equivalent to permuting the order of the condition subtraction). The highest *t*-values per iteration were entered into a distribution against which we tested the *t*-values of the real data. Next, within the task-sensitive voxels, we examined task-relevant comparisons (color-face, positive-negative feedback), again with permutation testing. In each of 1000 iterations, condition-difference power values per voxel were multiplied by −1 for a randomly chosen number of participants, and *t*-values were computed (condition difference against zero). Per voxel, all *t*-values were combined into a distribution of *t*-values against which we tested the real condition-difference *t*-value. Per iteration, we also stored the voxel count

of the largest significant cluster. Finally, these cluster sizes were used to generate a null distribution for maximum-statistic cluster-based correction for multiple comparisons (Maris and Oostenveld, 2007).

For all three steps, only voxels or clusters that exceeded the 95th percentile of the permutation-based distribution were considered significant ($p < 0.05$). To correct for the fact that we ran separate beamformers for multiple time-frequency windows, we applied two additional thresholds: we only included voxels that adhered to a threshold of $p < 0.005$ (to correct for the seven frequency bands) and showed significant differences over four consecutive time windows (to correct for the multiple time windows). All reported clusters adhere to these thresholds unless otherwise specified. Based on our hypothesis, we only selected frontal clusters that were sensitive to feedback valence, and posterior clusters that were sensitive to the stimulus dimension for further analyses. Although all clusters are reported in Table 1 and Table 2, only clusters that were included in further analyses are displayed in Figs. 4 and 5.

2.5.4. Test of modulation of post-feedback adjustment by stimulus dimension

Because the whole-brain comparisons resulted in substantially large clusters (ranging from 106 to 15014 voxels in the selected clusters, see Tables 1 and 2), we first selected more spatially restricted regions of interest (ROIs). Per cluster we selected the voxel with the highest t-value and the 49 spatially closest voxels (based on the Euclidean distance) that also showed significant differences, resulting in 50-voxel ROIs (6.25 cm^3). Large clusters with multiple, spatially separate peaks could result in multiple ROIs. Importantly, we did not investigate connectivity (see below) between ROIs originating from the same cluster, which could show spuriously strong correlations.

To investigate the interaction between feedback valence and stimulus-type specific processing, we ran beamformers on the learning task data for the four conditions of interest: FP, FN, CP, and CN (Fig. 3e). We then ran the beamformers only for two time windows, 100–500 and 500–900 ms after feedback presentation (combined with a 400 ms baseline window), to focus on post-feedback processes and avoid the temporal smoothing introduced by the initial, more exploratory sliding-window approach. These windows were based on previous demonstrations of learning-related differences in post-feedback theta-band power between 200 and 500 ms and in beta-band power between 500 and 800 ms (van de Vijver et al., 2011). Here, we used a window size of 400 ms to also allow accurate estimation of delta-band power. We used the estimated beamformer weights and filtered data per condition to reconstruct baseline and post-feedback power in each voxel of each ROI for all trials separately. Finally, we averaged these power values over voxels per ROI.

To study the effects of feedback valence and stimulus dimension on local power (Fig. 3e), we averaged power over trials and ROI voxels, per frontal and posterior ROI, and performed a dB conversion. Per ROI and time window, we compared dB-corrected power in the four learning conditions in the frequency band of interest (the frequency band in which the cluster was significant in the whole-brain analysis) in separate

ANOVAs with factors Feedback valence (positive, negative) and Stimulus dimension (color, face). Thus, per ROI we performed two ANOVAs for the two separate time windows.

To investigate whether frontal feedback evaluation was related to posterior stimulus processing, we investigated the similarity of trial-to-trial power fluctuations in pairs of distant ROIs (Fig. 3f; similar to, for example, Bruns et al., 2000; Mazaheri et al., 2009). Such cross-trial power-power correlations allow for more flexibility in the timing of the connectivity (as opposed to, for example, phase-based connectivity measures), and also allow for the assessment of the similarity in oscillatory signals with different frequencies (i.e. cross-frequency power-power coupling; see Siegel et al., 2012). Per ROI, we averaged activation over the 50 voxels and performed frequency-, time window-, and condition-specific baseline subtractions for every trial. This correction ensured that correlations could not be driven by general, large-scale trial-to-trial power fluctuations. We chose not to apply a dB conversion per trial because such a non-linear correction can unintentionally enlarge the influence of single trials with extreme power values. Next, for every frontal–posterior ROI pair, condition, and time window, trial-to-trial power fluctuations in the frequency band of interest were rank-correlated and Fisher's Z transformed. Finally, for each ROI pair we examined whether correlations differed significantly between the conditions of interest: per pair and time window we entered the correlations into separate ANOVAs with factors Feedback valence (positive, negative) and Stimulus dimension (color, face). Thus, per ROI pair we performed two ANOVAs for the two time windows. For all ANOVAs, we only report effects that are significant at a threshold of $\alpha = 0.005$.

3. Results

3.1. Behavior

Participants learned the correct associations: average accuracy was significantly higher than chance for both the face (t (12) = 13.276, $p < 0.001$) and the color stimuli (t (12) = 15.541, $p < 0.001$) (Fig. 1c). Additionally, accuracy increased over trial bins (F (3,10) = 101.761, $p < 0.001$). This increase was significant for all successive bin pairs (all p's < 0.002). Accuracy did not differ between stimulus dimensions (F (1,12) = 0.073, $p = 0.792$) nor did bin interact with stimulus dimension (F (3,10) = 2.012, $p = 0.130$), indicating that the extensive training on the task successfully equalized learning performance between the dimensions.

RTs also differed over trial bins (F (3,10) = 10.695, $p < 0.001$; Fig. 1d). Participants were significantly faster in bin 1 compared to bin 2 (F (1,12) = 27.673, $p < 0.001$). RTs did not differ between bins 2 and 3 (F (1,12) = 0.004, $p = 0.953$) or between bins 3 and 4 (F (1,12) = 3.615, $p = 0.082$). Participants responded faster when learning color than face stimuli (F (1,12) = 53.912, $p < 0.001$). There was no significant interaction effect of bin and stimulus dimension on RTs (F (3,10) = 2.711, $p = 0.082$).

Table 1
Clusters demonstrating a significant power difference between face and color stimuli in the localizer task. * = significant at p < 0.05 without correction for multiple frequencies and timepoints (Occip. = Occipital).

	Frequency (Hz)	Cluster location	Cluster size (voxels)	Power				Peak location (mm)			Peak time (ms post stimulus)	Duration (samples)
				Face		Color						
				M	SD	M	SD	x	y	z		
1	2–4	Premotor areas	742	0.926	0.385	0.149	0.317	−25	14	48	550	6
2	12–20	Occip., parietal cortex	1607	−1.659	0.817	−0.808	0.815	−35	−71	−7	750	6
3	20–30	Occip. cortex	399	−1.352	0.971	−0.618	1.041	20	−91	3	700	5
4	20–30	Cerebellum	16	−0.744	0.772	−0.130	0.370	−25	−61	−47	700	4
5*	20–30	Temporal cortex	213	−0.273	0.576	−0.765	0.582	55	−31	−17	550	5

Table 2

Clusters demonstrating a significant power difference between positive and negative feedback in the reinforcement learning task.* = significant at p < 0.05 without correction for multiple frequencies and timepoints (PFC = prefrontal cortex, fb = feedback).

|	Frequency (Hz)	Cluster location	Cluster size (voxels)	Power				Peak location (mm)			Peak time (ms post fb)	Duration (samples)
				Negative fb		Positive fb						
				M	SD	M	SD	x	y	z		
1	2–4	Anterior PFC, rostral cingulate zone, temporal cortex	15014	0.728	0.334	0.277	0.378	40	−41	−7	450	17
2	4–8	Anterior PFC	106	0.986	0.313	0.513	0.539	0	59	−2	300	5
3	8–12	Occipital cortex	2082	−1.007	0.715	−0.475	0.494	−25	−66	48	150	4
4*	4–8	Pre-SMA	535	0.487	0.338	0.214	0.330	25	19	48	400	7
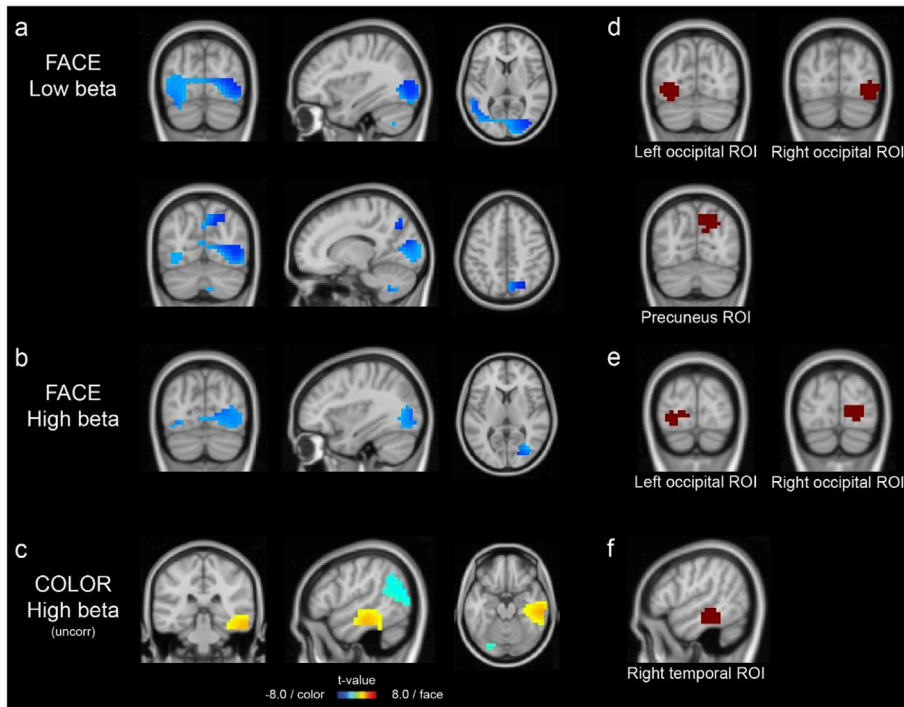
Fig. 4. Clusters of source-level voxels showing a significant difference between face and color stimuli in the localizer task. (a) Low beta power showed a larger decrease for face compared to color stimuli in a cluster comprising left and right occipital cortex and extending into right parietal cortex (upper row: x = 35.0, y = −78.1, z = 6.3 mm, time = 600 ms; lower row: x = 14.8, y = −68.8, z = 45.2 mm, time = 650 ms). (b) A similar larger decrease for faces compared to colors was seen in high beta power in a largely overlapping occipital cluster (x = 35.0, y = −75.7, z = 6.1 mm, time = 700 ms). (c) High beta power decreased more after colors compared to faces in right inferior and medial temporal cortex at p < 0.05 without correction for multiple frequencies and timepoints (x = 52.0, y = −33.1, z = −18.5 mm, time = 450 ms). (d) Three ROIs in the low beta band were selected, in left and right occipital cortex and right parietal cortex centering on the precuneus. (e) Two ROIs were selected in the high beta band in bilateral occipital cortex. (f) One high beta-band ROIs was selected within the temporal cluster. Brain images are displayed according to neurological convention (all selected times are windows of peak difference between conditions; uncorr = significant at p < 0.05 without correction for multiple frequencies and timepoints).
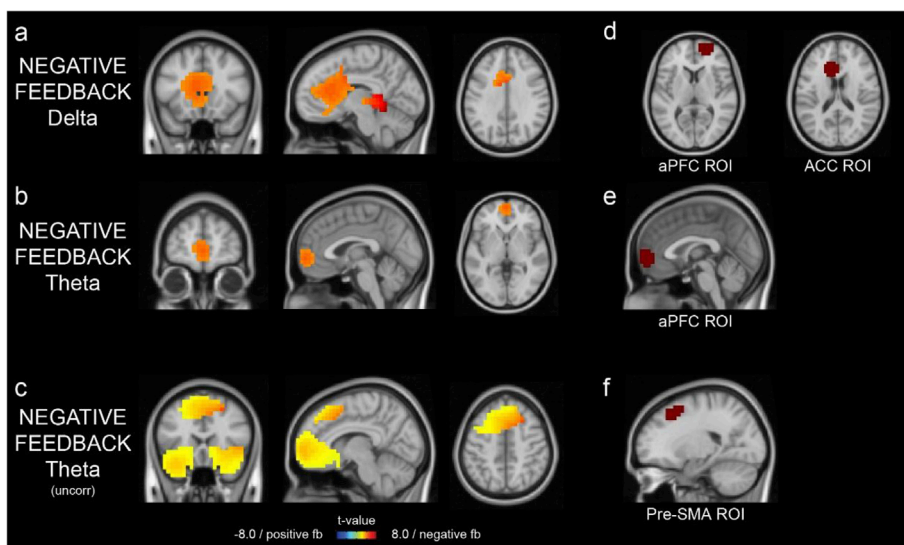

Fig. 5. Clusters of source-level voxels showing a significant difference between positive and negative feedback in the learning task. (a) Delta power was larger after negative compared to positive feedback in a large network comprising frontal, temporal, and parietal regions (x = −7.8, y = 16.5, z = 32.0 mm, time = 450 ms). (b) Theta power was larger after negative compared to positive feedback in anterior prefrontal cortex (aPFC; x = −1.6, y = 57.6, z = −0.6 mm, time = 300 ms). (c) Theta power was larger after negative compared to positive feedback in pre-supplementary motor cortex (pre-SMA) at p < 0.05 without correction for multiple frequencies and timepoints (x = 8.5, y = 19.6, z = 46.7 mm, time = 400 ms). (d) Two delta-band ROIs were selected, in right aPFC and left anterior cingulate cortex (ACC)/corpus callosum. (e) One theta-band ROI was selected within the aPFC cluster. (f) One theta-band ROI was selected within the pre-SMA cluster. Brain images are displayed according to neurological convention (time = window of peak difference between conditions; uncorr = significant at p < 0.05 without correction for multiple frequencies and timepoints, fb = feedback).

3.2. Sensor-level frontal theta-band power

To investigate whether the current MEG recordings were related to known feedback-related oscillatory dynamics measured with EEG, we examined sensor-level post-feedback theta power. At two midfrontal sensor locations, the condition-averaged post-feedback theta power (Fig. 2a) was increased. Although power was numerically larger after negative than positive feedback at both sensors 722/3 (sum of gradiometers 0722 and 0733) and 422/3 (Fig. 2b and c), this difference was not significant (sensor 722/3: F (1,12) = 2.569, p = 0.135; sensor 422/3: F (1,12) = 4.069, p = 0.067). At both locations, there were also no significant effects of stimulus dimension or interaction effects between feedback valence and stimulus dimension (all p-values > 0.1).

Next, we examined two more lateral frontal sensor clusters that also demonstrated increased condition-averaged theta power (see Fig. 2a). Per cluster, we averaged the data over sensors. Post-feedback theta power was significantly larger after negative compared to positive feedback at both clusters (left: F (1,12) = 6.857, p = 0.022, right: F (1,12) = 12.268, p = 0.004; Fig. 2b and c). Additionally, theta power was larger for the color compared to the face dimension at the right (F (1,12) = 6.126, p = 0.028), but not at the left cluster (F (1,12) = 0.221, p = 0.647) (not shown). Feedback valence and stimulus dimension did not interact at either cluster (p-values > 0.5). Thus, we replicated the known effect of feedback valence on theta power at lateral rather than midfrontal sensors.

3.3. Task-relevant source-level power differences

3.3.1. Localizer task

Four clusters of voxels demonstrated significant differences between face and color stimuli during the localizer task (Table 1, clusters 1–4). Delta-band power was larger for face than for color stimuli in a frontal cluster located anterior to the motor cortex, demonstrating the largest difference approximately 550 ms post-stimulus (not shown). Three posterior clusters showed larger power decreases for face compared to color stimuli. An extensive cluster in the low beta band comprised large parts of occipital and parietal cortex (Fig. 4a), and a smaller cluster in the high beta band overlapped with this cluster in occipital cortex (Fig. 4b). A second high beta-band cluster was located in the cerebellum (not shown). All three posterior clusters demonstrated the largest effect around 700–750 ms post-stimulus.

We selected the two posterior beta-band clusters centering on occipital cortex for further analyses and created ROIs per cluster around each separate peak. This resulted in three low-beta band ROIs in left lateral occipital cortex (peak voxel x = −35, y = −71, z = −7 mm), right lateral occipital cortex (40,-81,-7), and right precuneus (25,-71,43; Fig. 4d), and two high-beta band ROIs in left medial occipital cortex (−35,-76,-7) and right medial occipital cortex (20,-91,3; Fig. 4e).

At p < 0.05 but without correction for multiple frequencies and timepoints, we also found a cluster at the posterior side of the right inferior temporal cortex that demonstrated a larger high-beta band power decrease compared to baseline for colors compared to faces (Table 1 cluster 5, Fig. 4c). Because no posterior color-specific clusters survived thresholding, we included this cluster in subsequent analyses and created a ROI around its peak voxel (x = 55, y = −31, z = −17 mm; Fig. 4f).

3.3.2. Reinforcement learning task

Three clusters demonstrated significant effects of feedback valence in the learning task (Table 2, clusters 1–3). Power was larger after negative than positive feedback in a delta-band cluster covering a large part of frontal cortex, thalamus, and temporo-occipital areas (Fig. 5a), peaking 450 ms after feedback. Similarly, power was larger after negative feedback in a theta-band cluster located medially in anterior prefrontal cortex (aPFC), with the largest effect 300 ms after feedback (Fig. 5b). Finally, power decreased more after negative than positive feedback in a

posterior alpha-band cluster located mainly in occipital cortex (not shown). Although the largest effect in this cluster appeared 150 ms post-feedback, it already differentiated between positive and negative feedback 250 ms before feedback presentation.

We selected the frontal delta- and theta-band clusters for further analyses. Within the large delta-band cluster we selected two ROIs: one in anterior cingulate cortex (ACC; peak voxel x = 20, y = 59, z = 8 mm) and one in right aPFC (−10,24,18; Fig. 5d). Although the delta cluster contained additional frontal peaks, they all showed a maximal power difference before 200 ms post-feedback. The anterior frontal theta-band cluster contained only one peak resulting in one ROI (0,59,-2; Fig. 5e).

At p < 0.05 but without correction for multiple frequencies and timepoints, we found another large theta-band power cluster in the pre-supplementary motor area (pre-SMA; Table 2 cluster 4, Fig. 5c). Because post-feedback theta-band differences at this location are considered a core feature of feedback processing (Cohen et al., 2011), we included this cluster for further analyses and selected a ROI around its peak voxel (x = 25, y = 19, z = 48 mm; Fig. 5f).

3.4. Effects of feedback valence and stimulus dimension on local oscillatory power

We first investigated the effects of feedback valence and stimulus dimension on local post-feedback power per ROI. In the posterior ROIs, the only significant effect was a stronger suppression of low beta-band power in the precuneus ROI after negative compared to positive feedback (100–500 ms; negative −0.89 dB, positive −0.46 dB; F (1,12) = 12.989, p = 0.004). There were no other effects of feedback valence in the posterior ROIs (all p-values > 0.1), nor were there effects of stimulus dimension (all p-values > 0.17) or interaction effects of feedback valence and stimulus dimension (all p-values > 0.1). Because the stimuli were no longer visible in the investigated time windows, the absence of an effect of stimulus dimension cannot be interpreted as a direct disconfirmation of the effect of stimulus type that we found in the localizer task. These results do suggest, however, that local post-feedback stimulus-type specific adjustments are not reflected in oscillatory power in the same frequency bands and neuronal circuits as the initial processing of these stimuli.

In the frontal ROIs, the only significant effect was a larger increase in delta power after negative compared to positive feedback in the ACC ROI in the early (100–500 ms; negative 0.56 dB, positive 0.10 dB; F (1,12) = 13.049, p = 0.004) and late post-feedback time windows (500–900 ms; negative 0.71 dB, positive 0.21 dB; F (1,12) = 24.113, p < 0.001). Although the delta and theta clusters in aPFC were based on a valence contrast, the effect of valence in the current ANOVAs was not significant at the threshold of $α$ = 0.005 (delta power early window F (1,12) = 8.938, p = 0.011, late window F (1,12) = 7.800, p = 0.016; theta power early window F (1,12) = 9.357, p = 0.010, late window F (1,12) = 4.796, p = 0.049). This difference between the initial valence contrast and the current tests may be due to the different statistical approaches, but the current selection of predefined ROIs and time windows may also have decreased the effect compared to the initial, fully data-driven contrast. There was no effect of feedback valence on theta power in the pre-SMA ROI either in the early (F (1,12) = 1.743, p = 0.211) or the late time window (F (1,12) = 0.503, p = 0.492). There were no main effects of stimulus dimension in any of the frontal ROIs (all p-values > 0.18), and no interaction effects between feedback valence and stimulus dimension that were significant at the threshold of $α$ = 0.005 (all p-values > 0.01).

3.5. Feedback-related correlated fluctuations in power at distant locations

Finally, we investigated the effects of feedback valence and stimulus dimension on power-power correlations between pairs of frontal and posterior ROIs. We found significant interactions between feedback valence and stimulus dimension on the correlations in three ROI pairs, all

between 500 and 900 ms post-feedback: between ACC delta power and precuneus low-beta power (F (1,12) = 16.498, p = 0.002), between aPFC theta power and left occipital high-beta power (F (1,12) = 13.418, p = 0.003), and between aPFC theta power and right temporal high-beta power (F (1,12) = 16.244, p = 0.002).

Follow-up t-tests indicated that the directions of the interaction effects differed between the ROI pairs. ACC-precuneus correlations were significantly higher after negative feedback on color trials than face trials (t (12) = 2.499, = 0.028; Fig. 6a), whereas there was no difference after positive feedback (t (12) = −1.813, p = 0.095). In contrast, aPFC-left occipital power correlations were significantly more negative after negative feedback on color trials compared to face trials (t (12) = −2.355, p = 0.036; Fig. 6b), whereas there was no difference after positive feedback (t (12) = 1.713, p = 0.112). Finally, aPFC-temporal power correlations were significantly higher after positive feedback on color trials compared to face trials (t (12) = 3.108, p = 0.009; Fig. 6c), whereas there was no difference after negative feedback (t (12) = −0.731, p = 0.479).

There were no effects of stimulus dimension (all p-values > 0.039) or feedback valence (all p-values > 0.065) on power correlations between any frontal-posterior ROI pair and no other significant interaction effects (all p-values > 0.03), that were significant at the threshold of α = 0.005.

Taken together, the current results suggest that local power dynamics mainly dissociated between feedback processing (frontal delta/theta) and stimulus processing (posterior beta), while interregional power-correlation patterns showed specific interactions between the type of stimulus information and the feedback received.

3.6. Relation between long-distance power coupling and learning

To investigate whether frontal-posterior interactions were related to learning, we compared the correlations between the first (bins 1 and 2 from the behavioral analyses) and second part (bins 3 and 4) of the learning blocks. Per frontal-posterior ROI pair, we only selected the trials with the feedback valence for which that connection demonstrated a significant effect (negative feedback for the ACC-precuneus and aPFC-occipital connections, positive feedback for the aPFC-temporal connection). We split these trials into the first and second 50% for the two stimulus dimensions separately, and recomputed the correlations per stimulus dimension and block part. Finally, we entered these correlations into ANOVAs with factors Stimulus dimension (Color, Face) and Block part (First, Second). However, none of the connections demonstrated a significant interaction effect of stimulus dimension and block part (all p-values > 0.1), nor were there significant main effects of block part (all p-

values > 0.09).

4. Discussion

We investigated whether oscillation-based connectivity between the frontal RL network and posterior stimulus-processing areas reflected post-feedback adjustments in stimulus processing. After negative compared to positive feedback, theta-band power was larger over midfrontal and lateral frontal areas, and delta and theta power increased more in multiple frontal source-level clusters. Crucially, the effect of feedback valence on power-power correlations between frontal and visual cortical areas depended on the relevant stimulus dimension. These effects were present after feedback presentation, when the stimuli were not displayed. This suggests that the frontal learning network and task-relevant visual areas take part in post-feedback oscillation-based interactions to optimize future stimulus processing.

4.1. Feedback valence affects local frontal delta and theta power

The current effect of feedback valence on sensor-level theta power replicates previous EEG results (see, e.g., Cavanagh et al., 2010; Cohen et al., 2007; Mas-Herrero and Marco-Pallarés, 2014; van de Vijver et al., 2011). Whereas the location of the theta peak tends to be midfrontal rather than lateral prefrontal in EEG studies, the current peak location matches previous MEG findings on post-feedback oscillations (Doñamayor et al., 2012). Indeed, EEG and MEG have different sensitivity profiles for sources oriented radially versus tangentially with respect to the skull (Ahlfors et al., 2010; Hillebrand and Barnes, 2002). An anterior source-level cluster also demonstrated larger theta power after negative feedback, in line with EEG research showing that aPFC theta power reflects feedback valence and predicts behavioral adaptation and exploration (Cavanagh et al., 2012a,b; van de Vijver et al., 2014). Note, however, that eye blinks were not removed for the source-level analyses, although the tSSS and beamforming procedures likely removed most of this artifact (Adjamian et al., 2009; Hillebrand et al., 2013; but see Hipp and Siegel, 2015). Indeed, control analyses confirmed that eye blink artifacts did not affect our results.

Finally, an extensive delta-band cluster demonstrated a larger power increase for negative compared to positive feedback. Not much is known about frontal post-feedback delta power, although prior EEG studies reported increased delta-band power after response errors (Kolev et al., 2009, 2005), NoGo trials (Nigbur et al., 2011), and novel, unexpected events (Wessel and Aron, 2013). More generally, increased delta-band oscillations may reflect the processing of motivationally salient and


Fig. 6. Power correlations between frontal and posterior ROIs differed with feedback valence and stimulus dimension, between 500–900 ms post-feedback. (a) Correlations between delta power in the anterior cingulate cortex and low beta power in the precuneus were larger after negative feedback on color compared to face trials. (b) Correlations between theta power in anterior prefrontal cortex (aPFC) and high beta power in left occipital cortex were more negative after negative feedback on color trials compared to face trials. (c) Correlations between theta power in aPFC and high beta power in right temporal cortex were higher after positive feedback on color trials compared to face trials. Error bars represent the SEM (ACC = anterior cingulate cortex, aPFC = anterior prefrontal cortex, left occipital = left occipital cortex, right temporal = right temporal cortex; * significant at α < 0.05 in post-hoc t-test).

emotional stimuli, and may be negatively related to dopamine releases in the reward network (Knyazev, 2012).

4.2. Stimulus characteristics affect local posterior beta-band power

Beta power decreased in multiple posterior clusters depending on the stimulus characteristics in the localizer task. Alpha- and beta-power decreases are often found during stimulus presentation and may reflect changes in attentional processing (Bauer et al., 2012; Donner and Siegel, 2011; Worden et al., 2000). Larger power decreases reflect a stronger need to direct attention towards a stimulus, in line with the current larger beta-power decrease in occipital areas for the more complex face stimuli. As these condition differences in beta power peaked only 550–750 ms after stimulus presentation, the initial processing of color and face stimuli likely overlapped in brain location and frequency, and was not picked up by our contrast.

4.3. Frontal-posterior long-range connectivity reflects feedback valence and stimulus dimension

Post-feedback correlations in three frontal-posterior ROI pairs differed depending on a combination of stimulus dimension and feedback valence. Importantly, local power was only sensitive to stimulus information *or* feedback valence but did not reflect their combination. This suggests that the combination of stimulus and feedback information to update behavior critically relies on the information exchange *between* the learning network and task-relevant brain areas.

Connectivity fluctuations can arise in the absence of local power changes as they rely on the relative timing of the activity in the involved brain areas. This idea has been put forward most specifically in theories on phase coherence, describing how the effectiveness of connections depends on synchrony between the phases of their respective oscillations, either within or between frequency bands (Engel et al., 2001; Fries, 2015, 2005; Siebenhühner et al., 2016; Singer and Gray, 1995). Power-based representations of connectivity such as power-envelope correlations and the current power-power correlations are based on the notion that (anti)correlated fluctuations in activation magnitude shape long-distance communication in a similar way (Bruns et al., 2000; Mazaheri et al., 2009; O'Neill et al., 2015; Siegel et al., 2012).

Specifically, the increased correlations between frontal and posterior oscillations for particular stimulus-feedback combinations indicate that trial-to-trial fluctuations in power at both locations became more aligned under distinctive circumstances. Indeed, feedback can only lead to successful behavioral adaptations if it is properly interpreted in the context of the preceding stimulus and performed behavior. So far, investigations of functional interactions that support feedback-based learning have focused on interactions between midfrontal, (pre)motor, and prefrontal cortex (for an overview, see Luft, 2014). Most studies examined theta-band interactions (e.g., Cavanagh et al., 2010; Cavanagh et al., 2012a,b; Luft et al., 2013; van de Vijver et al., 2014, 2011), although feedback-related connectivity in the beta and gamma range has also been reported (De Pascalis et al., 2012; Luft et al., 2014). The current results add to these findings by showing that post-feedback connectivity extends to frontoposterior interactions, that this connectivity depends on stimulus-specific information, but also that feedback-related connectivity can involve multiple frequency bands. More specific neural mechanisms underlying such cross-frequency coupling, which have already been suggested in other cognitive domains (see, e.g., Bastos et al., 2015; Fries, 2015; Lisman and Jensen, 2013; Roux and Uhlhaas, 2014), may also play a role in post-feedback updating.

Instead of the hypothesized double dissociation, we found generally stronger feedback-related connectivity for color compared to face stimuli. This difference cannot be attributed to condition differences in learning accuracy or local power. Although RTs differed between stimulus dimensions, this likely reflected a difference in perceptual complexity of the stimuli rather than a learning difference. In line with

the localizer results, increased attention is thought to coincide with decreased alpha and beta power (Bauer et al., 2012; Donner and Siegel, 2011; Worden et al., 2000). Thus, a negative correlation between frontal low-frequency power and posterior beta power likely implies a simultaneous processing increase in both areas, whereas a positive correlation indicates a processing increase in one area and a simultaneous decrease in the other area. The current results therefore suggest that after negative feedback (and for color specifically), connectivity between ACC and precuneus decreased, whereas connectivity between aPFC and occipital cortex increased. After positive feedback, connectivity decreased between aPFC and temporal cortex.

The last two results seem to support the role of aPFC in tracking behavioral alternatives and switching to an explorative strategy (Cavanagh et al., 2012a; Daw et al., 2006): after negative feedback the alternative option must be considered, whereas after positive feedback the current association can be strengthened. Indeed, connectivity between aPFC and posterior task-relevant brain areas changes when switching to a different response alternative (Boorman et al., 2009) or when being instructed to attend a different stimulus dimension (Sakai and Passingham, 2006). Thus, we speculate that aPFC communicates to posterior regions which stimulus-response association should be adjusted or attended for future behavior.

4.4. Limitations of the current results

The interaction effects in frontal-posterior connectivity were fairly modest, which is not uncommon in measures of long-range oscillatory connectivity at the scalp level (see, for example, Nigbur et al., 2012; van de Vijver et al., 2011). Several factors may have contributed to this. First, more general trial-to-trial fluctuations in synchronization may have benefitted feedback processing but could have been removed during baseline subtraction. Second, task-relevant brain areas that did not differentiate between color and face stimuli were not included. Third, the interactions may have been too brief and local to be measurable at the level of MEG. Finally, our hypothesis-driven approach involved only a small number of ROIs and frequencies.

Relatedly, we did not find power-power correlations that reflected both stimulus processing and feedback valence and were more pronounced for the face than the color dimension of the stimuli. We speculate that this absence is due to the involvement of different cognitive processes in the color and face conditions during post-feedback adjustments. The color stimuli consisted of primary and secondary colors only. Adjustments in this condition may therefore mainly have involved changes in attention or discrimination at the visual level, represented along the visual stream. Adjustments in the face condition, however, may have involved changes in the memory representations of the different faces and their relation with the two response options. Such memory representations likely involve activity in hippocampal or surrounding brain regions, regions from which signals are more difficult to measure with MEG.

We assessed connectivity with power-power correlations, because frontal and posterior ROIs differed in frequency band. However, this connectivity measure does not provide information on whether the connection is direct, in which direction the information flows, or what the underlying physiological mechanism could be (although see Siegel et al., 2012). Future research should further investigate the nature of these relations, examining, for example, cross-frequency phase-amplitude coupling (Canolty and Knight, 2010) or cross-frequency directionality (Jiang et al., 2015).

The current results did not show a direct relation between power-power correlations and learning accuracy. The absence of a direct relation may have been caused by a decrease in power, because separating the first and second halves of the learning task reduced the number of trials per analysis compared with investigating the whole task. Relatedly, participants may have varied in their learning rates. Whereas the average accuracy suggests that learning mainly took place in the first half of each

block, the current split in two halves may in fact not appropriately represent separate stages of learning in some or even most participants. Moreover, it is possible that the current power-power correlations are related to feedback processing, but not at all to learning, which may reside in other parts of the network. However, as a direct relation between brain processes and learning may show in only one of a multitude of possible measures (for example, subject-level brain-behavior correlations or within-subject trial-to-trial fluctuations), the current results do not allow us to draw strong conclusions about the behavioral relevance of frontoposterior connectivity. More sensitive measures of learning behavior, the fitting of an RL model to examine individual behavior at a more detailed level including differences in learning rates, and a larger subject sample may help to investigate this relation more thoroughly.

Finally, like any measurement system of brain activity, MEG does not allow the investigation of all relevant brain regions. Although the striatum is very important for RL, the non-parallel orientation of the neurons in its structures prevents much of their signal from reaching the scalp and the MEG sensors (but see, e.g., Boon et al., 2017; Huang et al., 2017; Kaido et al., 2010). Thus, although we demonstrated variability in oscillation-based connectivity between PFC and visual areas, how they interact with the striatum remains to be determined.

5. Conclusions

Learning and feedback-based adjustment of stimulus-action associations require large-scale integration of information across multiple sensory, motor, and frontal areas. We provided evidence that neural oscillations and their coupling are signatures of this mechanism. Adjustments in connectivity may alter subsequent processing of the same stimulus, possibly by shifting attention to different response alternatives. Future studies should examine whether such adjustments are also reflected in the connectivity between stimulus- and response-processing areas.

Acknowledgements

Appendix A. Supplementary data

Supplementary data related to this article can be found at https://doi.org/10.1016/j.neuroimage.2018.07.014.

References

Adjmian, P., Worthen, S.F., Hillebrand, A., Furlong, P.L., Chizh, B.A., Hobson, A.R., Aziz, Q., Barnes, G.R., 2009. Effective electromagnetic noise cancellation with beamformers and synthetic gradiometry in shielded and partly shielded environments. J. Neurosci. Meth. 178, 120–127. https://doi.org/10.1016/j.jneumeth.2008.12.006.

Ahlfors, S.P., Han, J., Belliveau, J.W., Hämäläinen, M.S., 2010. Sensitivity of MEG and EEG to source orientation. Brain Topogr. 23, 227–232. https://doi.org/10.1007/s10548-010-0154-x.

Akam, T., Kullmann, D.M., 2010. Oscillations and filtering networks support flexible routing of information. Neuron 67, 308–320. https://doi.org/10.1016/j.neuron.2010.06.019.

Bastos, A.M., Vezoli, J., Fries, P., 2015. Communication through coherence with inter-areal delays. Curr. Opin. Neurobiol. 31, 173–180. https://doi.org/10.1016/j.conb.2014.11.001.

Bauer, M., Kluge, C., Bach, D., Bradbury, D., Heinze, H.J., Dolan, R.J., Driver, J., 2012. Cholinergic enhancement of visual attention and neural oscillations in the human brain. Curr. Biol. 22, 397–402. https://doi.org/10.1016/j.cub.2012.01.022.

Boon, L.I., Hillebrand, A., Olde Dubbelink, K.T.E., Stam, C.J., Berendse, H.W., 2017. Changes in resting-state directed connectivity in cortico-subcortical networks correlate with cognitive function in Parkinson's disease. Clin. Neurophysiol. 128, 1319–1326. https://doi.org/10.1016/j.clinph.2017.04.024.

Boorman, E.D., Behrens, T.E., Woolrich, M.W., Rushworth, M.F., 2009. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. Neuron 62, 733–743. https://doi.org/10.1016/j.neuron.2009.05.014.

Brookes, M.J., Vrba, J., Robinson, S.E., Stevenson, C.M., Peters, A.M., Barnes, G.R., Hillebrand, A., Morris, P.G., 2008. Optimising experimental design for MEG beamformer imaging. Neuroimage 39, 1788–1802. https://doi.org/10.1016/j.neuroimage.2007.09.050.

Bruns, A., Eckhorn, R., Jokeit, H., Ebner, A., 2000. Amplitude envelope correlation detects coupling among incoherent brain signals. Neuroreport 11, 1509–1514.

Canolty, R.T., Knight, R.T., 2010. The functional role of cross-frequency coupling. Trends Cognit. Sci. 14, 506–515. https://doi.org/10.1016/j.tics.2010.09.001.

Cashdollar, N., Malecki, U., Rugg-Gunn, F.J., Duncan, J.S., Lavie, N., Duzel, E., 2009. Hippocampus-dependent and -independent theta-networks of active maintenance. Proc. Natl. Acad. Sci. U. S. A 106, 20493–20498. https://doi.org/10.1073/pnas.0904823106.

Cavanagh, J.F., Figueroa, C.M., Cohen, M.X., Frank, M.J., 2012a. Frontal theta reflects uncertainty and unexpectedness during exploration and exploitation. Cerebr. Cortex 22, 2575–2586. https://doi.org/10.1093/cercor/bhr332.

Cavanagh, J.F., Frank, M.J., Klein, T.J., Allen, J.J.B., 2010. Frontal theta links prediction errors to behavioral adaptation in reinforcement learning. Neuroimage 49, 3198–3209. https://doi.org/10.1016/j.neuroimage.2009.11.080. S1053-8119(09)01266-X [pii].

Cavanagh, J.F., Zambrano-Vazquez, L., Allen, J.J.B., 2012b. Theta lingua franca: a common mid-frontal substrate for action monitoring processes. Psychophysiology 49, 220–238. https://doi.org/10.1111/j.1469-8986.2011.01293.x.

Cohen, M.X., 2014. Analyzing Neural Time Series Data: Theory and Practice. MIT Press, Cambridge, MA.

Cohen, M.X., Elger, C.E., Ranganath, C., 2007. Reward expectation modulates feedback-related negativity and EEG spectra. Neuroimage 35, 968–978. https://doi.org/10.1016/j.neuroimage.2006.11.056. S1053-8119(06)01180-3 [pii].

Cohen, M.X., Ridderinkhof, K.R., 2013. EEG source reconstruction reveals frontal-parietal dynamics of spatial conflict processing. PLoS One 8, e57293. https://doi.org/10.1371/journal.pone.0057293.

Cohen, M.X., van Gaal, S., 2013. Dynamic interactions between large-scale brain networks predict behavioral adaptation after perceptual errors. Cerebr. Cortex 23, 1061–1072. https://doi.org/10.1093/cercor/bhs069.

Cohen, M.X., van Gaal, S., Ridderinkhof, K.R., Lamme, V.A.F., 2009. Unconscious errors enhance prefrontal-occipital oscillatory synchrony. Front. Hum. Neurosci. 3, 54. https://doi.org/10.3389/neuro.09.054.2009.

Cohen, M.X., Wilmes, K., van de Vijver, I., 2011. Cortical electrophysiological network dynamics of feedback learning. Trends Cognit. Sci. 15, 558–566. https://doi.org/10.1016/j.tics.2011.10.004.

Dalal, S.S., Guggisberg, A.G., Edwards, E., Sekihara, K., Findlay, A.M., Canolty, R.T., Berger, M.S., Knight, R.T., Barbaro, N.M., Kirsch, H.E., Nagarajan, S.S., 2008. Five-dimensional neuroimaging: localization of the time–frequency dynamics of cortical activity. Neuroimage 40, 1686–1700. https://doi.org/10.1016/J.NEUROIMAGE.2008.01.023.

Dalal, S.S., Zumer, J.M., Agrawal, V., Hild, K.E., Sekihara, K., Nagarajan, S.S., 2004. NUTMEG: a neuromagnetic source reconstruction toolbox. Neurol. Clin. Neurophysiol. Annu. Vol. 2004, 52.

Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., Dolan, R.J., 2006. Cortical substrates for exploratory decisions in humans. Nature 441, 876–879. https://doi.org/10.1038/nature04766.

De Pascalis, V., Varriale, V., Rotonda, M., 2012. EEG oscillatory activity associated to monetary gain and loss signals in a learning task: effects of attentional impulsivity and learning ability. Int. J. Psychophysiol. 85, 68–78. https://doi.org/10.1016/j.ijpsycho.2011.06.005.

Delorme, A., Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Meth. 134, 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009.

den Ouden, H.E.M., Daunizeau, J., Roiser, J., Friston, K.J., Stephan, K.E., 2010. Striatal prediction error modulates cortical coupling. J. Neurosci. 30, 3210–3219. https://doi.org/10.1523/JNEUROSCI.4458-09.2010.

Doñamayor, N., Schoenfeld, M.A., Münte, T.F., 2012. Magneto- and electroencephalographic manifestations of reward anticipation and delivery. Neuroimage 62, 17–29. https://doi.org/10.1016/j.neuroimage.2012.04.038.

Donner, T.H., Siegel, M., 2011. A framework for local cortical oscillation patterns. Trends Cognit. Sci. 15, 191–199. https://doi.org/10.1016/j.tics.2011.03.007.

Ebner, N.C., 2008. Age of face matters: age-group differences in ratings of young and old faces. Behav. Res. Meth. 40, 130–136.

Engel, A.K., Fries, P., Singer, W., 2001. Dynamic predictions: oscillations and synchrony in top-down processing. Nat. Rev. Neurosci. 2, 704–716. https://doi.org/10.1038/35094565.

Fries, P., 2015. Rhythms for cognition: communication through coherence. Neuron 88, 220–235. https://doi.org/10.1016/j.neuron.2015.09.034.

Fries, P., 2005. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cognit. Sci. 9, 474–480. https://doi.org/10.1016/j.tics.2005.08.011.

Garrison, J., Erdeniz, B., Done, J., 2013. Prediction error in reinforcement learning: a meta-analysis of neuroimaging studies. Neurosci. Biobehav. Rev. 37, 1297–1310. https://doi.org/10.1016/j.neubiorev.2013.03.023.

Gregoriou, G.G., Gotts, S.J., Zhou, H., Desimone, R., 2009. High-frequency, long-range coupling between prefrontal and visual cortex during attention. Science 324, 1207–1210. https://doi.org/10.1126/science.1171402.

Gross, J., Baillet, S., Barnes, G.R., Henson, R.N., Hillebrand, A., Jensen, O., Jerbi, K., Litvak, V., Maess, B., Oostenveld, R., Parkkonen, L., Taylor, J.R., van Wassenhove, V., Wibral, M., Schoffelen, J.-M., 2013. Good practice for conducting and reporting MEG research. Neuroimage 65, 349–363. https://doi.org/10.1016/j.neuroimage.2012.10.001.

Haber, S.N., Knutson, B., 2010. The reward circuit: linking primate anatomy and human imaging. Neuropsychopharmacology 35, 4–26. https://doi.org/10.1038/npp.2009.129.

Hillebrand, A., Barnes, G.R., 2002. A quantitative assessment of the sensitivity of whole-head MEG to activity in the adult human cortex. Neuroimage 16, 638–650.

Hillebrand, A., Fazio, P., de Munck, J.C., van Dijk, B.W., 2013. Feasibility of clinical magnetoencephalography (MEG) functional mapping in the presence of dental artefacts. Clin. Neurophysiol. 124, 107–113. https://doi.org/10.1016/j.clinph.2012.06.013.

Hipp, J.F., Siegel, M., 2015. Accounting for linear transformations of EEG and MEG data in source analysis. PLoS One 10, e0121048. https://doi.org/10.1371/journal.pone.0121048.

Horga, G., Maia, T.V., Marsh, R., Hao, X., Xu, D., Duan, Y., Tau, G.Z., Graniello, B., Wang, Z., Kangarlu, A., Martinez, D., Packard, M.G., Peterson, B.S., 2015. Changes in corticostriatal connectivity during reinforcement learning in humans. Hum. Brain Mapp. 36, 793–803. https://doi.org/10.1002/hbm.22665.

Huang, M.-X., Harrington, D.L., Robb Swan, A., Angeles Quinto, A., Nichols, S., Drake, A., Song, T., Diwakar, M., Huang, C.W., Risbrough, V.B., Dale, A., Bartsch, H., Matthews, S., Huang, J.W., Lee, R.R., Baker, D.G., 2017. Resting-State magnetoencephalography reveals different patterns of aberrant functional connectivity in combat-related mild traumatic brain injury. J. Neurotrauma 34, 1412–1426. https://doi.org/10.1089/neu.2016.4581.

Jiang, H., Bahramisharif, A., van Gerven, M.A.J., Jensen, O., 2015. Measuring directionality between neuronal oscillations of different frequencies. Neuroimage 118, 359–367. https://doi.org/10.1016/j.neuroimage.2015.05.044.

Kaido, T., Otsuki, T., Kaneko, Y., Takahashi, A., Kakita, A., Takahashi, H., Saito, Y., Nakagawa, E., Sugai, K., Sasaki, M., 2010. Anterior striatum with dysmorphic neurons associated with the epileptogenesis of focal cortical dysplasia. Seizure 19, 256–259. https://doi.org/10.1016/j.seizure.2010.02.003.

Knyazev, G.G., 2012. EEG delta oscillations as a correlate of basic homeostatic and motivational processes. Neurosci. Biobehav. Rev. 36, 677–695. https://doi.org/10.1016/j.neubiorev.2011.10.002.

Kolev, V., Beste, C., Falkenstein, M., Yordanova, J., 2009. Error-related oscillations. Effects of aging on neural systems for behavioral monitoring. J. Psychophysiol. 23, 216–223.

Kolev, V., Falkenstein, M., Yordanova, J., 2005. Aging and error processing: time-frequency analysis of error-related potentials. J. Psychophysiol. 19, 289–297.

Lee, D., Seo, H., Jung, M.W., 2012. Neural basis of reinforcement learning and decision making. Annu. Rev. Neurosci. 35, 287–308. https://doi.org/10.1146/annurev-neuro-062111-150512.

Lisman, J.E., Jensen, O., 2013. The theta-gamma neural code. Neuron 77, 1002–1016. https://doi.org/10.1016/j.neuron.2013.03.007.

Luft, C.D.B., 2014. Learning from feedback: the neural mechanisms of feedback processing facilitating better performance. Behav. Brain Res. 261, 356–368. https://doi.org/10.1016/j.bbr.2013.12.043.

Luft, C.D.B., Meeson, A., Welchman, A.E., Kourtzi, Z., 2015. Decoding the future from past experience: learning shapes predictions in early visual cortex. J. Neurophysiol. 113, 3159–3171. https://doi.org/10.1152/jn.00753.2014.

Luft, C.D.B., Nolte, G., Bhattacharya, J., 2013. High-learners present larger mid-frontal theta power and connectivity in response to incorrect performance feedback. J. Neurosci. 33, 2029–2038. https://doi.org/10.1523/JNEUROSCI.2565-12.2013.

Luft, C.D.B., Takase, E., Bhattacharya, J., 2014. Processing graded feedback: electrophysiological correlates of learning from small and large errors. J. Cognit. Neurosci. 26, 1180–1193. https://doi.org/10.1162/jocn_a_00543.

Maia, T.V., 2009. Reinforcement learning, conditioning, and the brain: successes and challenges. Cognit. Affect Behav. Neurosci. 9, 343–364. https://doi.org/10.3758/CABN.9.4.343.

Marco-Pallarés, J., Münte, T.F., Rodríguez-Fornells, A., 2015. The role of high-frequency oscillatory activity in reward processing and learning. Neurosci. Biobehav. Rev. 49, 1–7. https://doi.org/10.1016/j.neubiorev.2014.11.014.

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG- and MEG-data. J. Neurosci. Meth. 164, 177–190. https://doi.org/10.1016/j.jneumeth.2007.03.024.

Mas-Herrero, E., Marco-Pallarés, J., 2014. Frontal theta oscillatory activity is a common mechanism for the computation of unexpected outcomes and learning rate. J. Cognit. Neurosci. 26, 447–458. https://doi.org/10.1162/jocn_a_00516.

Mazaheri, A., Nieuwenhuis, I., van Dijk, H., Jensen, O., 2009. Prestimulus alpha and mu activity predicts failure to inhibit motor responses. Hum. Brain Mapp. 30, 1791–1800.

Nigbur, R., Cohen, M.X., Ridderinkhof, K.R., Stürmer, B., 2012. Theta dynamics reveal domain-specific control over stimulus and response conflict. J. Cognit. Neurosci. 24, 1264–1274. https://doi.org/10.1162/jocn_a_00128.

Nigbur, R., Ivanova, G., Stürmer, B., 2011. Theta power as a marker for cognitive interference. Clin. Neurophysiol. 122, 2185–2194. https://doi.org/10.1016/J.CLINPH.2011.03.030.

Nolte, G., 2003. The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors. Phys. Med. Biol. 48, 3637–3652. https://doi.org/10.1088/0031-9155/48/22/002.

Nyberg, L., Habib, R., McIntosh, A.R., Tulving, E., 2000. Reactivation of encoding-related brain activity during memory retrieval. Proc. Natl. Acad. Sci. Unit. States Am. 97, 11120–11124. https://doi.org/10.1073/pnas.97.20.11120.

Nyberg, L., Petersson, K.M., Nilsson, L.G., Sandblom, J., Aberg, C., Ingvar, M., 2001. Reactivation of motor brain areas during explicit memory for actions. Neuroimage 14, 521–528. https://doi.org/10.1006/nimg.2001.0801.

O'Neill, G.C., Barratt, E.L., Hunt, B.A.E., Tewarie, P.K., Brookes, M.J., 2015. Measuring electrophysiological connectivity by power envelope correlation: a technical review on MEG methods. Phys. Med. Biol. 60, R271–R295. https://doi.org/10.1088/0031-9155/60/21/R271.

Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M., 2011. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput. Intell. Neurosci. 2011, 156869. https://doi.org/10.1155/2011/156869.

Palva, J.M., Monto, S., Kulashekhar, S., Palva, S., 2010. Neuronal synchrony reveals working memory networks and predicts individual memory capacity. Proc. Natl. Acad. Sci. U. S. A 107, 7580–7585. https://doi.org/10.1073/pnas.0913113107.

Rösler, F., Heil, M., Hennighausen, E., 2007. Distinct cortical activation patterns during long-term memory retrieval of verbal, spatial, and color information. J. Cognit. Neurosci. 7, 51–65.

Roux, F., Uhlhaas, P.J., 2014. Working memory and neural oscillations: alpha–gamma versus theta–gamma codes for distinct WM information? Trends Cognit. Sci. 18, 16–25. https://doi.org/10.1016/j.tics.2013.10.010.

Rushworth, M.F., Noonan, M.P., Boorman, E.D., Walton, M.E., Behrens, T.E., 2011. Frontal cortex and reward-guided learning and decision-making. Neuron 70, 1054–1069. https://doi.org/10.1016/j.neuron.2011.05.014.

Sakai, K., Passingham, R.E., 2006. Prefrontal set activity predicts rule-specific neural processing during subsequent cognitive performance. J. Neurosci. 26, 1211–1218. https://doi.org/10.1523/JNEUROSCI.3887-05.2006.

Salazar, R.F., Dotson, N.M., Bressler, S.L., Gray, C.M., 2012. Content-specific fronto-parietal synchronization during visual working memory. Science 80, 338.

Sauseng, P., Klimesch, W., Schabus, M., Doppelmayr, M., 2005. Fronto-parietal EEG coherence in theta and upper alpha reflect central executive functions of working memory. Int. J. Psychophysiol. 57, 97–103. https://doi.org/10.1016/j.ijpsycho.2005.03.018.

Schiffer, A.-M., Muller, T., Yeung, N., Waszak, F., 2014. Reward activates stimulus-specific and task-dependent representations in visual association cortices. J. Neurosci. 34, 15610–15620. https://doi.org/10.1523/JNEUROSCI.1640-14.2014.

Siebenhühner, F., Wang, S.H., Palva, J.M., Palva, S., 2016. Cross-frequency synchronization connects networks of fast and slow oscillations during visual working memory maintenance. Elife 5. https://doi.org/10.7554/eLife.13451.

Siegel, M., Donner, T.H., Engel, A.K., 2012. Spectral fingerprints of large-scale neuronal interactions. Nat. Rev. Neurosci. 13, 121. https://doi.org/10.1038/nrn3137.

Singer, W., Gray, C.M., 1995. Visual feature integration and the temporal correlation hypothesis. Annu. Rev. Neurosci. 18, 555–586. https://doi.org/10.1146/annurev.ne.18.030195.003011.

Slotnick, S.D., 2009. Memory for color reactivates color processing region. Neuroreport 20, 1568–1571. https://doi.org/10.1097/WNR.0b013e328332d35e.

Taulu, S., Hari, R., 2009. Removal of magnetoencephalographic artifacts with temporal signal-space separation: demonstration with single-trial auditory-evoked responses. Hum. Brain Mapp. 30, 1524–1534. https://doi.org/10.1002/hbm.20627.

Taulu, S., Simola, J., 2006. Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. Phys. Med. Biol. 51, 1759–1768. https://doi.org/10.1088/0031-9155/51/7/008.

van de Vijver, I., Cohen, M.X., Ridderinkhof, K.R., 2014. Aging affects medial but not anterior frontal learning-related theta oscillations. Neurobiol. Aging 35, 692–704. https://doi.org/10.1016/j.neurobiolaging.2013.09.006.

van de Vijver, I., Ridderinkhof, K.R., Cohen, M.X., 2011. Frontal oscillatory dynamics predict feedback learning and action adjustment. J. Cognit. Neurosci. 23, 4106–4121. https://doi.org/10.1162/jocn_a_00110.

van Driel, J., Ridderinkhof, K.R., Cohen, M.X., 2012. Not all errors are alike: theta and alpha EEG dynamics relate to differences in error-processing dynamics. J. Neurosci. 32, 16795–16806. https://doi.org/10.1523/JNEUROSCI.0802-12.2012.

Wang, H.-P., Spencer, D., Fellous, J.-M., Sejnowski, T.J., 2010. Synchrony of thalamocortical inputs maximizes cortical reliability. Science 328, 106–109. https://doi.org/10.1126/science.1183108.

Wessel, J.R., Aron, A.R., 2013. Unexpected events induce motor slowing via a brain mechanism for action-stopping with global suppressive effects. J. Neurosci. 33.

Whalen, C., Maclin, E.L., Fabiani, M., Gratton, G., 2008. Validation of a method for coregistering scalp recording locations with 3D structural MR images. Hum. Brain Mapp. 29, 1288–1301. https://doi.org/10.1002/hbm.20465.

Womelsdorf, T., Schoffelen, J.-M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., Fries, P., 2007. Modulation of neuronal interactions through neuronal synchronization. Science 316, 1609–1612. https://doi.org/10.1126/science.1139597.

Worden, M.S., Foxe, J.J., Wang, N., Simpson, G.V., 2000. Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. J. Neurosci. 20. RC63.