

Tracking Motivational Biases and Their Suppression in Time and Space

Johannes Algermissen (j.algermissen@donders.ru.nl),
Jennifer C. Swart, Emma J. van Dijk, René Scheeringa,
Roshan Cools, Hanneke E. M. den Ouden
Donders Centre for Brain, Cognition, and Behaviour,
Radboud University Nijmegen, the Netherlands

Abstract

Action selection is not only based on acquired knowledge about action-outcome contingencies, but also by evolutionary “priors” such as *motivational biases*: Organisms show a tendency to invigorate responding when hoping for rewards, and to hold back when attempting to avoid punishments. While these biases are likely adaptive in many situations, they need to be inhibited when maladaptive. We probed the neural basis of overcoming these biases by measuring simultaneous EEG and fMRI. Successful detection and suppression of biases was associated with an increased synchronization in the alpha band 175–325 ms post-stimulus, which on a trial-by-trial basis was negatively correlated with BOLD signal in left MFG and right SMG. At a later time window around responses, there was a much stronger synchronization for executed vs. withheld actions in lower frequencies (peak in theta band), which was positively correlated on a trial-by-trial basis with BOLD signal in ACC/ SMA as well as bilateral motor cortex and operculum. Our work spatially locates oscillatory signatures of action selection and motivational conflict resolution.

Keywords: decision-making; reinforcement learning; motivational biases; ACC; midfrontal theta

Introduction

When choosing what action to perform to maximize rewards and minimize punishments, organisms cannot always rely on slow, incremental learning from experience. Instead, they take into account priors that have evolved over evolution or over their lifetime. One candidate for such priors are *motivational biases* (Dayan, Niv, Seymour, & Daw, 2006; Guitart-Masip, Duzel, Dolan, & Dayan, 2014), for example agents’ tendency to exhibit active “Go” actions in face of rewards, but passive “NoGo” actions in face of punishments (Guitart-Masip et al., 2012; Swart et al., 2017, 2018).

Although these motivational biases might facilitate action selection in a majority of situations, they need to be inhibited if their suggested action is suboptimal, i.e. when agents have to actively avoid punishment or passively wait for rewards. Previous findings reported increased midfrontal theta band synchronization when people successfully overcome these biases (Cavanagh, Eisenberg, Guitart-Masip, Huys, & Frank, 2013; Swart et al., 2018). Using simultaneously recorded EEG and fMRI, we aimed to localize the source of this midfrontal theta band synchronization. We hypothesized that

theta synchronization reflects neural mechanisms in the anterior cingulate cortex (ACC) and (pre-) supplementary motor area (SMA) which detect conflict between bias-induced actions and action requirements, and in response increase the decision-threshold (i.e. bounds in a drift-diffusion framework) (Cavanagh et al., 2011; Cohen & Ridderinkhof, 2013; Frank et al., 2015).

Methods

Task

We simultaneously recorded EEG and fMRI while students ($N = 36$, $M_{age} = 23.58$, 25 female, all right-handed) performed the Motivational Go/NoGo Learning Task (Swart et al., 2017, 2018). In each trial, participants saw one of eight stimuli. They learned by trial-and-error whether a stimulus was a “Win” stimulus (yielding rewards or neutral outcomes) or “Avoid” stimulus (yielding neutral outcomes or punishments), and which action (Left Go button press, Right Go button press, or no button press) to perform to achieve their preferred outcome (i.e., rewards for Win stimuli, neutral outcomes for Avoid stimuli). Outcomes were probabilistic, with the preferred outcome delivered with a 80% chance for correct responses and a 20% chance for incorrect responses.

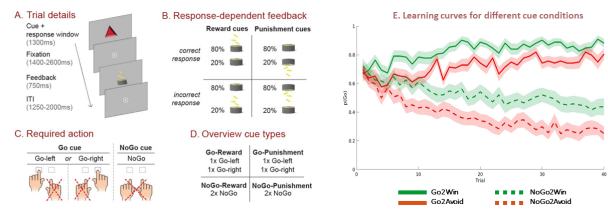


Figure 1: Task details and learning curves for the different cue conditions

EEG

We used a MRI-compatible EEG cap (BrainCap-MR-3-0 64Ch-Standard) plus extra channels for electrocardiogram, heartrate, and respiration, with a 1,000-Hz sampling rate (BrainCap; Brain Products, EasyCap; extended international 1020 layout). A Polhemus FASTRAK device was used to record the exact location of each EEG electrode on the participant’s head relative to three fiducial points.

EEG data were cleaned for scanner and cardioballistic artifacts using BrainVisionAnalyzer, and subsequently

pre-processed using Fieldtrip. After rejection of noise-contaminated channels, data were epoched, re-referenced to the grand average, and bandpass filtered between 1–15 Hz. We performed ICAs to remove components associated with blinks, saccades, MR artifacts, and head motion. We removed global noise by computing Laplacian filters, which we also used to interpolate rejected channels, and subsequently performed time-frequency decomposition using Hanning tapers (1-15 Hz, 400 ms time windows) and decibel conversion. Baseline-correction was achieved by fitting a linear trend through the signal at stimulus onset, and removing the predicted baseline. For statistical tests, we used Fieldtrip's cluster-based nonparametric permutation test.

fMRI

We collected EPI data on a 3T Siemens Prisma using a whole-brain T2*-weighted multiband-4 sequence (TR/ TE = 1400/32 ms, 68 slices, voxel size 2.0 mm isotropic, 75° flip angle, A/P phase encoding direction). Anatomical images were acquired with a T1-weighted MP-RAGE (GRAPPA acceleration factor = 2, TR/TE = 2300/3.03 ms, voxel size 1 mm isotropic, 8° flip angle).

Using FSL 6.0.0, data were preprocessed by applying brain extraction (BET), realignment (MCFLIRT), and smoothing (FWHM 2 mm). We used ICA-AROMA to detect and remove independent components associated with head motion, and afterwards high-pass filtered at 100 s. At the second-level GLM, statistical maps were registered to structural space and normalized to MNI152 space (FNIRT). We used additionally collected fieldmaps to correct for field distortions.

We fit a GLM with a fully crossed design including stimulus valence (Win vs. Avoid), required action (Go vs. NoGo), and actually executed action (Go vs. NoGo). We additionally included a) regressors of no interest for left vs. right button presses, errors, outcome onsets and valences and invalid button presses, b) the six realignment parameters, mean CSF and out-of-brain signal as nuisance regressors, and c) single spike regressors for volumes with a relative displacement > 2 mm. First-level subject-specific contrasts were combined on a second level using FSL's mixed-effects models tool *FLAME*, with a cluster-forming threshold of $z > 3.1$ and cluster-correction at a level of $\alpha < .05$.

EEG-inspired fMRI

For preliminary analyses, we extracted the t -values for each time-frequency-channel combination (only channels Fz, FCz, Cz) within a significant cluster as identified with the permutation test. We then used these t -values as weights in a linear filter to compute a trial-by-trial EEG estimate of conflict-related alpha and action-related broadband-signal (peaked in the theta range). Future analyses will involve source-modeling of these effects using beamforming to suppress noise. We then added these regressors as parametric modulators to the GLM design specified above. We were interested in BOLD variability that was explained by EEG regressors above and beyond the task regressors.

Results

Behavior

Go vs. NoGo actions were analyzed with mixed-models logistic regression as a function of required action and stimulus valence. Participants successfully learned the task (main effect of required action, $p < .001$), but showed strong motivational biases, i.e. more Go actions for Win stimuli than for Avoid stimuli (main effect of valence, $p < .001$; Figure 1e). Using computational reinforcement learning models, we replicated previous findings that stimulus valence biased both action selection and learning (Swart et al., 2017, 2018).

Overcoming motivational conflict

Using a permutation test, we rejected the null hypothesis of no difference between correct congruent and incongruent trials in the theta-range (4–8 Hz) over midfrontal electrodes ($p = .023$). Closer inspection yielded a stronger synchronization for incongruent compared to congruent actions 175–325 ms post-stimulus, which was however rather located in the alpha range (8–13 Hz), leaking into the theta band (Figure 2a-c). Both time and frequency range deviated from findings by Swart et al. (2018), while timing was more similar to results by Cavanagh et al. (2013). While previous research found conflict-related midfrontal theta increases to correlate positively with reaction times (Cohen & Cavanagh, 2011), we found the conflict-related alpha signal to be correlated *negatively* with reaction times ($\beta = -0.03$, $t(27.62) = 2.52$, $p = .022$).

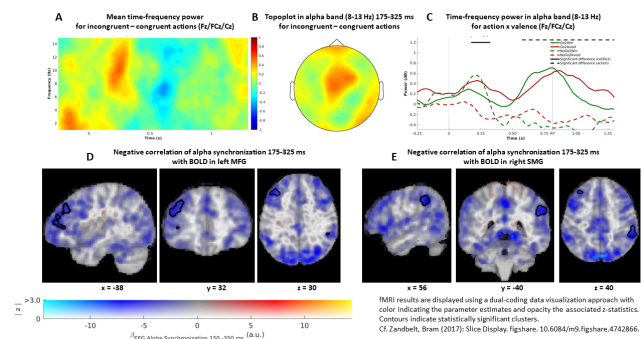


Figure 2: Conflict-related alpha (8-13 Hz) power synchronization over midfrontal channels (stimulus-locked) and its BOLD correlates

In our BOLD analyses, first, we contrasted bias-incongruent actions with bias-congruent actions, which yielded clusters of increased BOLD signal in bilateral superior frontal gyrus and precuneus (Figure 3a). Second, we selected only bias-incongruent actions and contrasted trials in which these actions were correct (i.e. the bias had to be overcome to perform correctly) with those trials in which they were incorrect (i.e. bias was overcome unnecessarily, potentially because action and/ or valence were not learned yet). This contrast

yielded significant clusters in dorsal ACC as well as in bilateral inferior lateral occipital cortex (Figure 3b).

Finally, when entering the trial-by-trial alpha power 175–325 ms post-stimulus as an additional regressor on top of the task-based regressors, we found clusters in left middle frontal gyrus (MFG) and right supramarginal gyrus (SMG) that correlated significantly negatively with alpha-power (Figure 2d-e). For $z > 3$ uncorrected, further areas in precuneus were found that correlated negatively with alpha power.

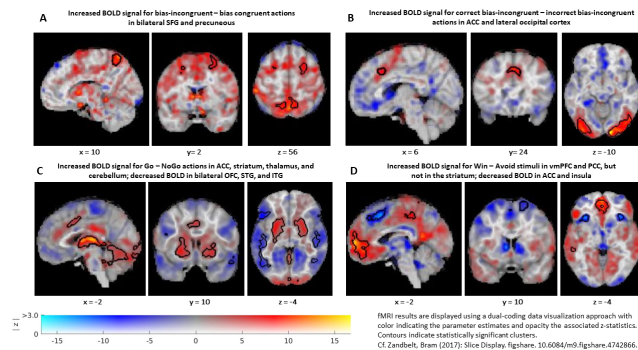


Figure 3: BOLD differences for incongruent vs. congruent action (a), correct incongruent vs. incorrect incongruent actions (b), Go vs. NoGo actions (c), and Win vs. Avoid stimuli (d)

Action dominates valence

Surprisingly, broadband power was dominated by a strong increase for Go compared to NoGo actions 500–1300 ms post-stimulus ($p = .006$) / -175–425 ms response-locked ($p = .004$), which was most accentuated in the theta band over frontopolar (Fpz) and central (Cz) electrodes, but extended into alpha and beta band (Figure 4a-c). The signal started ramping upwards prior to the response and peaked at the time of the response. This ramping signal might have occluded differences in midfrontal theta synchronization between incongruent and congruent actions reported previously (Cavanagh et al., 2011; Swart et al., 2018). Alternatively, it might reflect conflict in selecting between left and right Go responses in our right-handed participants. Indeed, theta power was significantly increased for left compared to right hand responses 550–700 ms post-stimulus ($p = .008$) / -225–50 ms response-locked ($p = .016$).

In our BOLD analyses, we found clusters of increased BOLD for Go vs. NoGo actions in dorsal ACC, striatum, thalamus, and cerebellum. BOLD was higher for NoGo vs. Go actions in clusters in bilateral inferior frontal gyri, superior and inferior temporal gyri, right SMG, and lateral occipital cortex (Figure 3c). Notably, the striatum appeared to encode action, but not stimulus valence (Win vs. Avoid stimuli); if anything, BOLD in medial caudate was higher for Avoid than for Win stimuli (Figure 3d).

When entering the trial-by-trial theta power -175–425 ms around responses as an additional regressor on top of the

task-based regressors, we found clusters in dorsal ACC/ SMA, bilateral motor cortex, and bilateral operculum/ insula that correlated significantly positively with theta-power (Figure 4d-e).

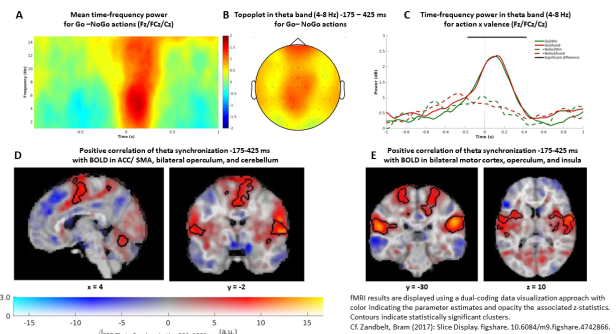


Figure 4: Action-related theta (4-8 Hz) power synchronization over midfrontal channels (response-locked) and its BOLD correlates

Discussion

We found successful inhibition of motivational biases to be associated with increased synchronization in the alpha band 175–325 ms post-stimulus, which was negatively correlated with BOLD in left MFG and right SMG. This signal was different in timing and frequency band from a previous finding with the same task (Swart et al., 2018) which showed increased synchronization in the theta band 450-650 ms post-stimulus. However, timing was more in line with findings by Cavanagh et al. (2013).

The brain areas negatively correlated with this alpha signal were part of the fronto-parietal network and spatially close to regions found more active for NoGo compared to Go actions. However, these regions were *less* activated on trials with stronger conflict-related alpha synchronization, suggesting *less* response inhibition when alpha was strong. We thus speculate that alpha synchronization might reflect mechanisms complementary to response inhibition. Given that we also found areas in the dorsal attention network to be correlated negatively with alpha, we speculate that directing attention away from stimulus valence might be an alternative strategy to suppress motivational biases (Kanske, Heissler, Schönfelder, Bongers, & Wessa, 2011). This might explain our finding of increased activity in visual cortex for correct bias-incongruent actions. These findings are in line with the suggested role of alpha in top-down control of attention (Klimesch, Sauseng, & Hanslmayr, 2007; Jensen & Mazaheri, 2010). Further analyses might involve EEG source modeling to better denoise our EEG regressors.

In our initialized hypothesized time-frequency window (i. e. 450–650 ms in the theta band), we observed a strongly ramping signal dissociating Go vs. NoGo actions. One possibility is that our right-handed subjects experienced conflict in selecting left vs. right Go actions. Another possibility is that this

signal reflects a role of theta in action-selection and decision-making, more generally (Womelsdorf, Vinck, Leung, & Everling, 2010), which is in line with our observation of BOLD in ACC/ SMA, motor cortex, and operculum correlating positively with theta synchronization. The signal's time course and spatial topography bear striking resemblance to the drift signal in a drift-diffusion process, reflecting increasing motor excitability until a threshold is passed and an action executed (Polanía, Krajbich, Grueschow, & Ruff, 2014). This theta signal was paralleled by BOLD in striatum and ACC encoding executed action rather than stimulus valence, replicating earlier work (Guitart-Masip et al., 2012). These results are in line with recent suggestions that the striatum does not encode value per se, but only the *value of work*, i.e. an active, effortful action (Collins & Frank, 2015; Hamid et al., 2015). Future analyses might attempt to link this signal to a drift process fitted to behavioral data.

In sum, our work contributes to the understanding of how humans overcome motivational biases by putatively directing attention away from stimulus valence, as reflected in phasic alpha synchronization. Also, we conclude that theta synchronization might reflect processes of action selection and initiation in ACC/ SMA, motor cortex, and striatum.

Acknowledgments

This work was supported by grants by the Netherlands Organization for Scientific Research (NWO; grant number 406-14-028) to JCS; Netherlands Organization for Scientific Research (NWO; grant number 453-14-005) to RC; James S. McDonnell Foundation James McDonnell Scholar Award to RC; Netherlands Organization for Scientific Research (NWO; grant number 451-11-004) to HEMdO; Netherlands Organization for Scientific Research (NWO; grant number 452-17-016) to HEMdO.

References

Cavanagh, J. F., Eisenberg, I., Guitart-Masip, M., Huys, Q. J. M., & Frank, M. J. (2013). Frontal theta overrides Pavlovian learning biases. *Journal of Neuroscience*, *33*(19), 8541–8548. doi: 10.1523/JNEUROSCI.5754-12.2013

Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, *14*(11), 1462–1467. doi: 10.1038/nn.2925

Cohen, M. X., & Cavanagh, J. F. (2011). Single-trial regression elucidates the role of prefrontal theta oscillations in response conflict. *Frontiers in Psychology*, *2*, 1–12. doi: 10.3389/fpsyg.2011.00030

Cohen, M. X., & Ridderinkhof, K. R. (2013). EEG source reconstruction reveals frontal-parietal dynamics of spatial conflict processing. *PLoS ONE*, *8*(2), e57293. doi: 10.1371/journal.pone.0057293

Collins, A. G. E., & Frank, M. J. (2015). Surprise! Dopamine signals mix action, value and error. *Nature Neuroscience*, *19*(1), 3–5. doi: 10.1038/nn.4207

Dayan, P., Niv, Y., Seymour, B., & Daw, N. (2006). The misbehavior of value and the discipline of the will. *Neural Networks*, *19*(8), 1153–1160. doi: 10.1016/j.neunet.2006.03.002

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, *35*(2), 485–494. doi: 10.1523/JNEUROSCI.2036-14.2015

Guitart-Masip, M., Duzel, E., Dolan, R., & Dayan, P. (2014). Action versus valence in decision making. *Trends in Cognitive Sciences*, *18*(4), 194–202. doi: 10.1016/j.tics.2014.01.003

Guitart-Masip, M., Huys, Q. J., Fuentemilla, L., Dayan, P., Duzel, E., & Dolan, R. J. (2012). Go and no-go learning in reward and punishment: Interactions between affect and effect. *NeuroImage*, *62*(1), 154–166. doi: 10.1016/j.neuroimage.2012.04.024

Hamid, A. A., Pettibone, J. R., Mabrouk, O. S., Hetrick, V. L., Schmidt, R., Vander Weele, C. M., ... Berke, J. D. (2015). Mesolimbic dopamine signals the value of work. *Nature Neuroscience*, *19*(1), 117–126. doi: 10.1038/nn.4173

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 1–8. doi: 10.3389/fnhum.2010.00186

Kanske, P., Heissler, J., Schönfelder, S., Bongers, A., & Wessa, M. (2011). How to regulate emotion? Neural networks for reappraisal and distraction. *Cerebral Cortex*, *21*(6), 1379–1388. doi: 10.1093/cercor/bhq216

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition-timing hypothesis. *Brain Research Reviews*, *53*(1), 63–88. doi: 10.1016/j.brainresrev.2006.06.003

Polanía, R., Krajbich, I., Grueschow, M., & Ruff, C. C. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron*, *82*(3), 709–720. doi: 10.1016/j.neuron.2014.03.014

Swart, J. C., Frank, M. J., Määttä, J. I., Jensen, O., Cools, R., & den Ouden, H. E. M. (2018). Frontal network dynamics reflect neurocomputational mechanisms for reducing maladaptive biases in motivated action. *PLoS Biology*, *16*(10), e2005979. doi: 10.1371/journal.pbio.2005979

Swart, J. C., Froböse, M. I., Cook, J. L., Geurts, D. E., Frank, M. J., Cools, R., & den Ouden, H. E. (2017). Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. *eLife*, *6*, 1–54. doi: 10.7554/eLife.22169

Womelsdorf, T., Vinck, M., Leung, L. S., & Everling, S. (2010). Selective theta-synchronization of choice-relevant information subserves goal-directed behavior. *Frontiers in Human Neuroscience*, *4*, 1–13. doi: 10.3389/fnhum.2010.00210