Accepted Manuscript

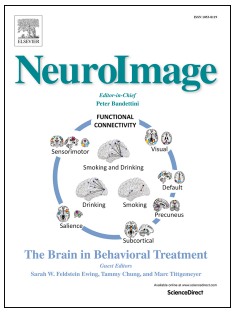

Decoding attentional states for neurofeedback: Mindfulness vs. wandering thoughts

A. Zhigalov, E. Heinilä, T. Parviainen, L. Parkkonen, A. Hyvärinen

PII: S1053-8119(18)31973-6

DOI: 10.1016/j.neuroimage.2018.10.014

Reference: YNIMG 15333

To appear in: *NeuroImage*

Received Date: 23 March 2018

Revised Date: 2 October 2018

Accepted Date: 4 October 2018

Please cite this article as: Zhigalov, A., Heinilä, E., Parviainen, T., Parkkonen, L., Hyvärinen, A., Decoding attentional states for neurofeedback: Mindfulness vs. wandering thoughts, *NeuroImage* (2018), doi: https://doi.org/10.1016/j.neuroimage.2018.10.014.

Decoding attentional states for neurofeedback: mindfulness vs. wandering thoughts

Zhigalov A[1,3*], Heinilä E[2], Parviainen T[2], Parkkonen L[3], Hyvärinen A[1,4]

[1]Department of Computer Science, University of Helsinki, Finland
[2]Department of Psychology, University of Jyväskylä, Finland
[3]Department of Neuroscience and Biomedical Engineering, Aalto University, Finland
[4]Gatsby Computational Neuroscience Unit, University College London, UK

* Correspondence should be addressed to,
Alexander Zhigalov
Department of Neuroscience and Biomedical Engineering
Biomedical Engineering
Aalto University School of Science
P.O. Box 12200
FI-00076 AALTO
alexander.zhigalov@aalto.fi

Abstract

Neurofeedback requires a direct translation of neuronal brain activity to sensory input given to the user or subject. However, decoding certain states, e.g., mindfulness or wandering thoughts, from ongoing brain activity remains an unresolved problem.

In this study, we used magnetoencephalography (MEG) to acquire brain activity during mindfulness meditation and thought-inducing tasks mimicking wandering thoughts. We used a novel real-time feature extraction to decode the mindfulness, i.e., to discriminate it from the thought-inducing tasks. The key methodological novelty of our approach is usage of MEG power spectra and functional connectivity of independent components as features underlying mindfulness states. Performance was measured as the classification accuracy on a separate session but within the same subject.

We found that the spectral- and connectivity-based classification approaches allowed discriminating mindfulness and thought-inducing tasks with an accuracy around 60% compared to the 50% chance-level. Both classification approaches showed similar accuracy, although the connectivity approach slightly outperformed the spectral one in a few cases. Detailed analysis showed that the classification coefficients and the associated independent components were highly individual among subjects and a straightforward transfer of the coefficients over subjects provided near chance-level classification accuracy.

Thus, discriminating between mindfulness and wandering thoughts seems to be possible, although with limited accuracy, by machine learning, especially on the subject-level. Our hope is that the developed spectral- and connectivity-based decoding methods can be utilized in real-

time neurofeedback to decode mindfulness states from ongoing neuronal activity, and hence, provide a basis for improved, individualized mindfulness training.

Keywords

Neurofeedback, magnetoencephalography, machine learning, mindfulness

Introduction

A brain computer interface (BCI), an essential component for neurofeedback, allows translating patterns of neuronal activity in the brain to inputs or commands for external devices (Wolpaw J. et al., 2002). Electroencephalography- (EEG) or magnetoencephalography-based (MEG) non-invasive BCIs provide opportunities for numerous clinical, assistive and entertainment applications. However, they require robust decoding of neuronal patterns: The person who controls the BCI should learn to produce robust neuronal patterns and/or a device that implements BCI should robustly identify these patterns.

Machine learning approaches have recently been successfully applied to detect some cognitive operations or mental states from underlying neuronal activity (Lemm et al., 2011; Lotte et al., 2007). Often, the classification is performed on the neuronal activity evoked by time-locked presentation of a target stimulus or task (Blankertz et al., 2011), which maximizes the signal-to-noise ratio of neuronal response. However, there are plenty of applications that require detection of cognitive states from stimulus-free ongoing neuronal activity.

One application where machine-learning BCI could be very useful is mindfulness meditation, which has been shown to have several positive behavioural outcomes (Tang et al., 2014). During mindfulness meditation, attention is supposed to be focussed on the breath or a similar target, but in reality, it varies over the time. Thus, it would be crucial to alert the user in real-time about wandering thoughts, i.e., temporary weakening of mindfulness. Detection of wandering thoughts can be considered as an example of neurofeedback system where a desired mental state can be defined by the user but the corresponding brain signal is not known. Such a sustained-attention neurofeedback system would presumably have many other applications, for instance, in driving assistance (Schmidt et al., 2009).

Several functional magnetic resonance imaging (fMRI) studies suggest that activity of default-mode network may underlie mind-wandering thoughts (Andrews-Hanna et al., 2014; Christoff et al., 2009). On the other hand, a recent study suggests that mind wandering can be represented through dynamic connectivity of the brain networks (Kucyi, 2017). Nevertheless, the indirect coupling of BOLD signal with neuronal activity and the relatively low temporal resolution of fMRI do not allow assessment of the contribution of short-lasting cognitive processes that largely constitute wandering thoughts. Several attempts have been made to find the brain correlates of wandering thoughts during sustained attention tasks with millisecond resolution using electroencephalography (EEG; (Baldwin et al., 2017; Braboszcz and Delorme,

2011). These studies have reported changes in the power at alpha frequency during mind wandering, but the neuroanatomical basis remains poorly understood because of low spatial resolution of EEG. A technically more advanced, simultaneous EEG-fMRI neurofeedback study (Ros et al., 2013) showed an increase of connectivity in default-mode network, which was positively correlated with changes in mind-wandering as well as resting state alpha rhythm. Although this study provided a link between the brain rhythms and anatomy for mind-wandering, the relationship between EEG and BOLD signal is rather indirect (Scheeringa et al., 2011). Surprisingly, there is a lack of magnetoencephalography (MEG) studies. MEG provides better temporal resolution compared to fMRI and better spatial resolution compared to EEG, which allow assessing fast neuronal processes and brain networks that are closely related to fMRI networks (Brookes et al., 2011).

In this study, we designed a behavioural paradigm where the subject performed mindfulness meditation, and two different tasks mimicking wandering thoughts, in consequent blocks. In contrast to mindfulness meditation, the latter tasks supposedly induced numerous thoughts, e.g., related to positive future plans or anxious emotional scenes. We used active tasks instead of resting state as a contrast for mindfulness meditation, because variability of resting state activity seems to be too unspecific, possibly related to various preceding mental states, including mindfulness itself. We developed and applied spectral- and connectivity-based classification approaches to discriminate the behavioural states based on their underlying neuronal activity, in particular focussing on the choice of feature extraction methods.

Our results show that it is possible, to some extent, to detect wandering thoughts in on-going MEG measurements, and we provide a sketch of a pipeline for optimizing such detection. Furthermore, in addition to the conventional view that mindfulness meditation is characterized by changes in the power of alpha oscillations, the changes in (dynamic) functional connectivity may provide an alternative, possibly even better, description of mindfulness states.

Materials and Methods

Experimental design

Twenty-four subjects (9 females, 27±5.5 years (mean±SD)) with moderate or no previous experience in mindfulness meditation participated in the study. Prior to the study, we performed a screening to include subjects with no history of neurological disorders, head trauma or substance abuse. All participants had normal or corrected to normal vision. Ten subjects had no previous experience, while other subjects had experience in either focused attention or open monitoring meditation practices ranged from 0.5 to 10 years.

After a 2-minute resting state, participants were instructed to perform one of the tasks while undergoing MEG (Elekta Neuromag, TRIUX). The tasks were organized into 2-minute blocks with counterbalanced order and the participants performed each task four times in a single session. The session ended with a resting state 2-minutes block. We conducted two sessions per participant with a 5-minutes break between the sessions (Fig. 1).

Figure 1 about here

Fig. 1. Diagram of the experiment. The tasks are MF (mindfulness meditation), FP (reflection of future planning) and EP (reflection on anxiousness-inducing emotional pictures).

The tasks were mindfulness meditation (MF), reflection on future planning (FP) and reflection on anxiousness-inducing emotional pictures (EP). In all tasks, subjects were instructed to sit still, fix the gaze on the crosshair, and perform a task after a short (seven seconds) visual instruction. The visual instruction was shown at the beginning and at the middle of each task (Fig. 1) to keep subjects' attention. After each task, the subject was asked to evaluate his/her involvement in the task by answering two questions "How focussed were you on the task?" and "How did you feel during the task? (pleasantness)" using a touch pad that provided a gradual response within range from 0 to 1.

For the *mindfulness meditation task*, the subject was instructed to focus his/her attention on the sensations of breathing and move his/her focus of attention back to the task if mind-wandering occurs. The task started with a visual instruction "Please focus on your breathing" accompanied by a picture of clouds. For the future planning and anxiety-inducing tasks, the subject individually selected 16 (out of 40) relevant pictures prior to the experiment. In the *future planning task*, subject was asked to perform a planning related to the picture, presumably following the ensuing chains of thought and keeping his/her mind busy. The task started with an instruction "Please make plans related to the picture" accompanied by a relevant picture. The *anxiety-inducing task* was similar to the future planning, but instead of neutral pictures, disturbing, scary, disgusting or other unpleasant pictures were shown to the subject. The task started with a visual instruction "Place yourself or someone close to you in this situation" accompanied by a relevant picture. For the FP and EP tasks two different pictures were presented for each 2-minute block (at the beginning and in the middle), and for the MF task, the same picture of clouds was presented twice.

Analysis of behavioural data

We analysed the subject's responses for the question "How focussed were you on the task?" (see, Fig. 1) ranged from 0 to 1. The average values were the following 0.65±0.012 (mindfulness), 0.70±0.010 (future planning task) and 0.66±0.013 (anxiety-inducing task) respectively. This result showed that the focus during different tasks was at a reasonable level. Additionally, we compared the responses between different tasks using the Wilcoxon rank sum test. The results showed no significant difference ($p > 0.05$) between the tasks, suggesting that the participants performed these tasks equally well.

Data pre-processing

In the analysis, we used the 204 planar gradiometers of the MEG scanner. The Signal Space Separation (SSS) method (Taulu et al., 2004) was applied to supress the external interference and sensor artefacts.

Independent component analysis

To obtain neurophysiologically realistic sources of neuronal activity, we applied complex-valued independent component analysis in the frequency-domain (Fourier-ICA; (Hyvärinen et al., 2010)) to MEG sensor's time series. The time series were divided into four-second epochs with 75 percent overlap and the epochs were Fourier transformed within the range of 4–24 Hz. A complex-valued ICA using 64 PCA components was applied to the epochs (from first sessions only), concatenated across subjects, which provided a group-level ICA un-mixing matrix. Note that the effective dimension of the data was reduced to approximately 64 by SSS.

Component selection criteria

Often, several independent components in MEG reflect physiological activity that is unrelated to neuronal activity of the brain (Jas et al., 2017). To exclude possible confounders, we applied three criteria to select independent components. *First*, we included only components that had a spectral peak within the range 8–16 Hz and the power of this peak was at least 50% larger than the power in theta (4–7 Hz) or high-beta (17–24 Hz) bands. Before this comparison, we equalized the spectral power over frequencies by subtracting the best fitting power-law function from the power spectra. *Second*, we analysed the component's spatial maps and excluded those components that had more than three blobs (i.e., continuous areas containing 5% of largest values) in the spatial map. *Third*, we excluded the components if their maximum in spatial map was located in the frontal or tempo-frontal areas. As a result, 38 components (out of 64) were selected for further analysis. Generally speaking, brain sources are spatially localized and band-pass, which leads to our formulation of the first two criteria, and the third criterion was included to exclude eye artifacts.

Spectral features extraction

The individual subjects' spectral features were extracted in the following manner. The sensor time series were divided into four-second epochs with 75 percent overlap and the epochs were Fourier transformed within the range of 4–24 Hz. The group-level ICA un-mixing matrix was applied to these epochs, and the amplitude spectra of frequency-domain independent components were computed (Suppl. Fig. 1A). In order to increase robustness of the spectral approach, we averaged the spectral amplitudes within four frequency bands: theta (4–7 Hz), alpha (8–12 Hz), low-beta (13–16 Hz) and high-beta (17–24 Hz). The amplitude spectra averaged inside these four frequency bands were then used as features for classification.

Connectivity features extraction

As an alternative to the spectral features, individual connectivity features were extracted as follows. Again, the sensor time series were divided into four seconds epochs with 75 percent overlap and the epochs were Fourier transformed within the range of 4–24 Hz. The group-level ICA un-mixing matrix was applied to these epochs and inverse Fourier transform was computed

to reconstruct components' time series. The independent component time series were filtered in four frequency bands: theta (4–7 Hz), alpha (8–12 Hz), low-beta (13–16 Hz) and high-beta (17–24 Hz). In contrast to the spectral feature case, the Pearson correlation coefficients were computed for each pair of the components (Suppl. Fig. 1B), separately for each frequency band. The correlation coefficients between epochs of independent component time series were well above zero (Suppl. Fig. 2). The connectivity matrices were then vectorised and used as features for classification.

Feature dimensionality reduction

To further improve the robustness of classification, we reduced the dimensionality of the spectral- and connectivity-based features by applying an algorithm described by Kauppi and colleagues (Kauppi et al., 2013). The idea is to compute, for each epoch and component, the most discriminating spectral or connectivity feature, and only use that in the classification. The dimensionality of features were thus reduced as follows,

$$V_{i,j} = \sum_{t \in T1} F_{i,t,j} - \sum_{t \in T2} F_{i,t,j}$$

$$P_{i,t} = \sum_{j} F_{i,t,j} \cdot \frac{V_{i,j}}{\|V_{i,}\|}$$

where **F** denotes a tensor containing the spectral or connectivity features (whose dimensions are either: components=38 *by* epochs *by* frequency_bands=4; or: component_pairs=703 *by* epochs *by* frequency_bands=4, respectively). Here, *i* denotes component or connectivity pair index; T1 and T2 are indices of task 1 and task 2 in the training dataset, respectively; ‖·‖ denotes vector norm operator; $V_{i.}$ means the vector consisting of all the $V_{i,j}$ for different *j*. The final result is given in matrix **P** which contains the resulting feature vectors of the epochs, with dimensions: components=38 *by* epochs; or component_pairs=703 *by* epochs, for spectral- or connectivity-based classification, respectively.

Classification methods: individual vs. group classification

The spectral and connectivity features with reduced dimensionality were classified using the linear Support Vector Machine (SVM) algorithm as implemented in scikit-learn (Pedregosa et al., 2011). We used two scenarios to train and test the classifiers. In the first scenario "individual classifier", we trained the classifier using individual data from the first session and tested the classifier using data from the second session. In the second scenario "group classifier", we trained the classifier using data from both sessions and all subjects except one "testing" subject, and tested the classifier using the testing subject's data from the second session. The second scenario is more challenging, essentially providing information on the generalizability of the classifier across subjects (Jayaram et al., 2016; Kia et al., 2017).

Real-time computation

We tested the computational time for both feature extraction and classification algorithms to ensure that our approaches suit real-time applications. The algorithms implemented in Python

were launched on a Linux based laptop (Intel Core i5-3570 @ 3.40 GHz, 8.00 GB RAM). The average feature extraction time (±SD) for single epoch was 0.0118±0.0085 and 0.0385±0.0192 second for the spectral and connectivity approaches, respectively. The average classification time using linear SVM was 0.0212±0.0042 and 0.1310±0.0086 seconds for the spectral and connectivity approaches respectively. Thus, the computational time was negligibly small compared to the inter-epoch interval of 0.5 second.

Statistical analysis of classification accuracies

To assess the statistical differences between the classification accuracies for different tasks or for different feature sets, we applied the Wilcoxon rank sum test.

Statistical analysis of classification coefficients

To evaluate the contribution of different spectral and connectivity features to the resulting classification accuracy, we performed a statistical analysis of the significance of the classification coefficients. The coefficients were divided by their absolute sum and averaged across subjects and then compared against zero mean using a two-sided z-test. For the spectral-based approach, we reported both uncorrected ($p < 0.05$) and Bonferroni corrected classification coefficients, while for the connectivity-based approach, we reported only Bonferroni corrected classification coefficients, since there the problem of multiple testing was more serious.

Results

We applied the spectral and connectivity approaches to investigate whether and how it is possible to discriminate (decode) between mindfulness meditation (MF), future planning (FP) and reflection on anxious-inducing emotional pictures (EP) tasks. We then analysed the classification coefficients to identify neuronal correlates (spatial maps and spectral profiles) associated with the mental states during task performance.

Classification accuracy

We first computed the classification accuracies using spectral and connectivity approaches (Fig. 2).

Figure 2 about here

Fig. 2. Spectral- and connectivity-based classification accuracies averaged across subjects for (A) individual and (B) group classifiers. MF_FP denotes mindfulness meditation vs. future planning task, MF_EP denotes mindfulness meditation vs. reflection of anxious-inducing emotional pictures task, and FP_EP denotes future planning task vs. reflection of anxious-inducing emotional pictures task. Error bars represent SEM.

The spectral approach with individual classifier provided accuracies well above chance-level, 0.59±0.008 (MF vs. FP), 0.61±0.007 (MF vs. EP) and 0.55±0.005 (FP vs. EP). The accuracies MF vs. FP and MF vs. EP were significantly larger ($p < 0.004$, individual classifier; $p < 0.01$, group classifier) than the accuracy FP vs. EP, suggesting that FP and EP have similar neuronal correlates. The accuracies for the group spectral-based classifier were relatively low, 0.55±0.003 (MF vs. FP), 0.54±0.003 (MF vs. EP) and 0.52±0.002 (FP vs. EP), showing that the classifier had poor generalization over subjects.

The connectivity approach provided slightly higher accuracies compared to the spectral approach, 0.62±0.009 (MF vs. FP), 0.62±0.009 (MF vs. EP) and 0.55±0.004 (FP vs. EP). However, the accuracies were not significantly different ($p > 0.05$) between the spectral and connectivity classifiers. The accuracies at the group level classifier were similar to those of the spectral approach, 0.55±0.004 (MF vs. FP), 0.54±0.002 (MF vs. EP) and 0.53±0.002 (FP vs. EP).

We repeated the analysis by swapping training and testing sessions. There was no difference between accuracies for the original and swapped sessions for the individual classifier, but the difference became non-significant for the group classifier (Suppl. Fig. 3).

Relationship between spectral- and connectivity-based accuracies

To compare of the approaches further, we assessed the relationship (correlation) between accuracies of spectral- and connectivity-based classifiers at individual subject's level (Fig. 3). The results showed a significant correlation between spectral- and connectivity-based accuracies only for two tasks, MF vs. FP ($r = 0.53$, $p < 0.008$) and MF vs. EP ($r = 0.60$, $p < 0.002$) for individual classifier.

Figure 3 about here

Fig. 3. Scatter plots of individual subject's accuracies computed using spectral and connectivity approaches for the following tasks, (A) mindfulness mediation vs. reflection on future planning (MF_FP), (B) mindfulness mediation vs. reflection on anxiety-inducing emotional pictures (MF_EP) and (C) reflection on future planning vs. reflection on anxiety-inducing emotional pictures (FP_EP).

For the further analysis, we considered only results of the individual classifier and two tasks (MF_FP and MF_EP).

Spectral and connectivity projections

We analysed the projection weights provided by the dimensionality reduction algorithm (see, Materials and Methods), to evaluate the frequency-specific differences between the tasks.

The results showed that that largest difference between tasks was associated with alpha frequency (Fig. 4). For the spectral classifier, the projection weights at alpha frequency were a little larger compared to those of connectivity classifier. Surprisingly, there was a strong

variability of the projection weights for group spectral classifier for different tasks, in contrast to individual spectral and connectivity classifiers.

Figure 4 about here

Fig. 4. Group level spectral and connectivity projections. Error bars represent SEM.

This result suggested that indeed, alpha frequency play an important role in discriminating mindfulness and thought-inducing tasks.

Classification coefficients of the spectral approach

We analysed the SVM classification coefficients of spectral classifier to evaluate the impact of different independent components and frequency bands on the accuracy. The coefficients obeyed Gaussian distribution and we picked those coefficients whose absolute values were above the chance-level corresponding to $p < 0.05$ (Fig. 5).

Figure 5 about here

Fig. 5. (A) Classification coefficients (absolute values) of spectral approach for the individual classifier. The significant coefficients indicated by black ($p < 0.05$, uncorrected) and red ($p < 0.05$, Bonferroni corrected) arrows, respectively. (B) Spatial maps and spectral profiles of independent components associated with significant coefficients. The grey curves indicate power spectra with subtracted power-law fit.

There were a few independent components (all independent components are shown in Suppl. Fig. 4A and 4B) associated with significantly larger classification coefficients. The components were associated with both thought-inducing tasks, showing a slight increase in power spectra at alpha frequency and located in the occipital (components 16, 20 and 30) or central areas (component 11). After correcting for multiple testing, only one coefficient (component 30) remained significant. Importantly, none of the significant weights was associated with components characterised by strong (very clearly peaked) alpha oscillations, although several such components were available for classification (see Suppl. Fig. 4A and 4B).

These results showed that rhythmic activity in occipital areas, but apparently not exclusively in alpha frequency, makes the strongest contribution to the discrimination of mindfulness and thought-inducing tasks.

To further elucidate the classification weights for individual subjects, we selected four largest coefficients and associated components. The results showed that the components associated with largest coefficients were highly individual (Suppl. Fig. 5), which make generalization of the classifier coefficients over subjects impractical.

Classification coefficients of the connectivity approach

The connectivity approach provided similar classification accuracies to the spectral approach, although these approaches utilized different principles for feature extraction. Similarly to the spectral approach, the classification coefficients followed a Gaussian distribution. We applied a t-test and selected significant ($p < 0.05$, Bonferroni corrected) classification coefficients (Fig. 6).

Figure 6 about here

Fig. 6. (A) Significant classification coefficients representing connections (i.e., connectivity pairs), for individual classifier. (B) Node degree when considering only the significant connections shown in panel A. Degree of zero means that there are no significant connections. The significant components with node degree above one indicated by arrows. (C) Spatial maps and spectral profiles of independent components associated with component's node degree above one.

The results showed that multiple connections contribute to the classification accuracy (Fig. 6A). To identify the independent components that represented a hub (i.e., node with a high degree) and hence, more strongly influenced the classification accuracy, we analysed the node degrees of the components using only the significant connections (Fig. 6B). We selected the components with node degree above one, and further analysed them (Fig. 6C).

Similarly to the spectral approach, strongest connections were associated with the occipital and temporal components, some of which demonstrated a peak at alpha frequency. Only one component (component 17) was characterized by particularly prominent alpha oscillations (see Suppl. Fig. 4A and 4B for comparison), again suggesting that alpha oscillations may not be strongly related to mindfulness states.

These results demonstrated a considerable variety of spatial maps and spectral patterns, which can underlie the difference in mindfulness and thought-provoking tasks.

To evaluate the contribution of high node degree components in individual subjects, we selected four components with highest node degree for each subject. The results showed that the high node degree components were highly individual (Suppl. Fig. 6), and similarly to the spectral-based classification, generalization over the subjects seems impractical.

All components classification

As noted in the Materials and Methods section, we excluded independent components that seem to be strongly contaminated by non-brain physiological artifacts. However, it is interesting to assess the impact such components would have on the classification. To accomplish this, we selected all the 64 independent components and applied the spectral- and connectivity-based classification approaches (Fig. 7).

Figure 7 about here

Fig. 7. Classification accuracies for (A) spectral and (B) connectivity approaches using 38 (solid bars) and 64 (dashed bars) components for individual and group classifiers.

We observed no differences in the accuracy for 38 and 64 independent components for the individual spectral and connectivity classifiers, as well as for the group connectivity classifier (Fig. 7). However, there were large differences in the accuracy for the group spectral classifier in MF vs. FP ($p < 0.001$) and MF vs. EP ($p < 0.03$).

To clarify the difference, we analysed the classification coefficients of the group spectral classifier (Fig. 8).

Figure 8 about here

Fig. 8. (A) The coefficients of the group spectral classifier for 64 independent components. The significant coefficients indicated by black ($p < 0.05$, uncorrected) and red ($p < 0.05$, Bonferroni corrected) arrows, respectively. (B) The spatial maps and spectral profiles associated with the significant coefficients.

The results clearly showed that when artifacts were not removed, the classifier picked the components that were mainly associated with physiological artifacts such as eye-blinks (component 6 and 18), cardiac activity (component 52), instrumental noise or other physiological artifacts (20, 50, 56). On the other hand, the possibly sensory-motor component (15) may have been erroneously excluded from the main analysis. After correcting for multiple testing, three coefficients (components 18, 50 and 52) remained significant.

In this case, the higher accuracy of the group spectral classifier likely to be related to the fact that the artifacts were highly consistent across subjects (Smith and Nichols, 2018), and thus provided a better basis for generalizing to new subjects.

Individual and group ICA

To evaluate the impact of ICA on classification accuracy, we recomputed the accuracies using spectral- and connectivity-based classifiers for individual and group ICA weights. Because the spectral profiles and spatial maps of individual independent components strongly varied across subjects, we did not apply the component selection criteria (see, *Materials and Methods*) and computed the classification accuracies for all 64 components (Fig. 9). There were no significant differences in the accuracies for individual and group ICA. This means that ICA spatial filters can be precomputed in advance and the classification approach can be operated in real time.

Figure 9 about here

Fig. 9. Classification accuracies for the individual and group ICA weights using (A) spectral and (B) connectivity approaches.

Discussion

In this study, we developed spectral- and connectivity-based classification approaches and showed that the mental states underlying mindfulness and thought-provoking tasks can be discriminated using MEG recordings and machine learning approaches.

We observed a variety of spatial and spectral patterns that contribute to discriminating mindfulness meditation, suggesting that several neuronal mechanisms may underlie mindfulness state. While the results in Fig. 4 indicate that the alpha frequency band is the most important, they also show the major contribution of other bands for classification. Moreover, our detailed analysis showed that none of the strongest classification weights was associated with components characterised by particularly strong alpha oscillations. Indeed, among many components with a noticeable peak at the alpha frequency, the components with the most pronounced peaks in alpha frequency (clearly oscillatory components) did not show any significant contribution to the discrimination. This seems to refute the conventional view on mindfulness meditation where the mindfulness state is tightly related to alpha frequency oscillations (Kerr et al., 2013). However, one caveat is that we used rather inexperienced meditators, and the situation might be different in the case of more experienced meditators.

In our results, spectral and connectivity features gave similar classification accuracies, which raises the question of whether they contain the same information. This does not seem to be so based on comparisons of the most important sources in Figures 5 and 6, and so it would seem that the similar classification accuracies may be only a coincidence. However, since our connectivity measures used zero lag, some overlapping information is likely to be present.

We visualized the spatial weights enabling classification between mindfulness and the tasks simulating wandering thoughts. It should be noted that the connection between these classifier weights and the neural correlates is not straight-forward. Interpretation of the weights can lead to wrong conclusions regarding the origin of neural signals of interest, since significant nonzero weights may also be associated with task-irrelevant signals (Haufe et al., 2014). However, from a neurofeedback viewpoint, the classifier weights reported here show which brain areas should be measured, e.g. in a case of building a portable EEG cap with a small number of sensors.

Considering that the task-relevant frequencies and brain regions may not be simply linked to the classification weights, there have been a few attempts to clarify the neuronal basis of mindfulness meditation. Gil and colleagues (Navarro Gil et al., 2018) showed that an EEG neurofeedback that aims at increasing power at the alpha frequency, improves mindfulness

outcome, and thus, may be effective for increasing mindfulness in healthy individuals. Unfortunately, the neurofeedback signal in this study was derived by averaging a set of occipital-parietal electrodes, which makes difficult to assess the location of underlying sources. Another EEG study (van Lutterveld et al., 2017) overcame such limitation by deriving the neurofeedback signal in a source space. The neurofeedback was provided based on gamma-band activity (40–57 Hz) from the posterior parietal cortex. The subjects were able to volitionally control the neurofeedback signal in the direction associated with effortless awareness by practicing effortless awareness meditation. Hence, these two studies suggested that not only alpha but also other frequencies are associated with mindfulness meditation, and parietal cortex may have a key role in mindfulness.

We observed relatively low classification accuracy in discriminating between future planning and reflection on anxiousness-inducing emotional pictures tasks, which suggests similarity of the rhythmic neuronal activity as captured by MEG during these tasks. Although these tasks are behaviourally quite different, and likely to be different in terms of amplitudes of the evoked responses in an affective picture paradigm (Olofsson et al., 2008), they may be similar in terms of task-nonspecific cortical processes related to attentional states. Consequently, our analysis, focusing on ongoing brain activity, may not be sensitive to this difference.

The overall classification accuracy in this study was nearly sixty percent, which is relatively low for a neurofeedback system. However, accuracies for a few participants were around seventy percent or more, which may allow a significant improvement in mindfulness meditation. Moreover, this relatively low accuracy may be explained by the fact that the subjects did not have previous experience in mindfulness meditation. Possibly, the neurofeedback might work much better after subjects gain more experience. Generalization over subjects was even more difficult, presumably due to the large individual differences and the methods are more likely to work when a large amount of data can be collected from each single subject. However, it should be noted that advanced multi-task classification methods might be able to generalize better by either finding some structure in data that is invariant across subjects or finding some structure in the decision rules between different subjects (Jayaram et al., 2016; Kia et al., 2017). Furthermore, combining brain measurements with psychophysiological measurements (for instance, heart-rate variability (Nesvold et al., 2012)) and other modalities might be useful to obtain practically useful classification performance in future research.

Acknowledgements

This work was supported by the Academy of Finland [grant number 295075] to AH, TP, LP; the Finnish Cultural Foundation [grant number 00161132] to AZ.

References

Andrews-Hanna, J.R., Smallwood, J., Spreng, R.N., 2014. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. Ann. N. Y. Acad. Sci. 1316,

29–52. https://doi.org/10.1111/nyas.12360
Baldwin, C.L., Roberts, D.M., Barragan, D., Lee, J.D., Lerner, N., Higgins, J.S., 2017. Detecting and Quantifying Mind Wandering during Simulated Driving. Front. Hum. Neurosci. 11, 406. https://doi.org/10.3389/fnhum.2017.00406
Blankertz, B., Lemm, S., Treder, M., Haufe, S., Müller, K.-R., 2011. Single-trial analysis and classification of ERP components — A tutorial. Neuroimage 56, 814–825. https://doi.org/10.1016/j.neuroimage.2010.06.048
Braboszcz, C., Delorme, A., 2011. Lost in thoughts: neural markers of low alertness during mind wandering. Neuroimage 54, 3040–7. https://doi.org/10.1016/j.neuroimage.2010.10.008
Brookes, M.J., Woolrich, M., Luckhoo, H., Price, D., Hale, J.R., Stephenson, M.C., Barnes, G.R., Smith, S.M., Morris, P.G., 2011. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proc. Natl. Acad. Sci. U. S. A. 108, 16783–8. https://doi.org/10.1073/pnas.1112685108
Christoff, K., Gordon, A.M., Smallwood, J., Smith, R., Schooler, J.W., 2009. Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. Proc. Natl. Acad. Sci. 106, 8719–8724. https://doi.org/10.1073/pnas.0900234106
Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., Bießmann, F., 2014. On the interpretation of weight vectors of linear models in multivariate neuroimaging. Neuroimage 87, 96–110. https://doi.org/10.1016/j.neuroimage.2013.10.067
Hyvärinen, A., Ramkumar, P., Parkkonen, L., Hari, R., 2010. Independent component analysis of short-time Fourier transforms for spontaneous EEG/MEG analysis. Neuroimage 49, 257–271. https://doi.org/10.1016/j.neuroimage.2009.08.028
Jas, M., Engemann, D.A., Bekhti, Y., Raimondo, F., Gramfort, A., 2017. Autoreject: Automated artifact rejection for MEG and EEG data. Neuroimage 159, 417–429. https://doi.org/10.1016/j.neuroimage.2017.06.030
Jayaram, V., Alamgir, M., Altun, Y., Scholkopf, B., Grosse-Wentrup, M., 2016. Transfer Learning in Brain-Computer Interfaces Abstract\uFFFDThe performance of brain-computer interfaces (BCIs) improves with the amount of avail. IEEE Comput. Intell. Mag. 11, 20–31. https://doi.org/10.1109/MCI.2015.2501545
Kauppi, J.-P., Parkkonen, L., Hari, R., Hyvärinen, A., 2013. Decoding magnetoencephalographic rhythmic activity using spectrospatial information. Neuroimage 83, 921–936. https://doi.org/10.1016/j.neuroimage.2013.07.026
Kerr, C.E., Sacchet, M.D., Lazar, S.W., Moore, C.I., Jones, S.R., 2013. Mindfulness starts with the body: somatosensory attention and top-down modulation of cortical alpha rhythms in mindfulness meditation. Front. Hum. Neurosci. 7, 12. https://doi.org/10.3389/fnhum.2013.00012
Kia, S.M., Pedregosa, F., Blumenthal, A., Passerini, A., 2017. Group-level spatio-temporal pattern recovery in MEG decoding using multi-task joint feature learning. J. Neurosci. Methods 285, 97–108. https://doi.org/10.1016/j.jneumeth.2017.05.004
Kucyi, A., 2017. Just a thought: How mind-wandering is represented in dynamic brain connectivity. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.07.001
Lemm, S., Blankertz, B., Dickhaus, T., Müller, K.-R., 2011. Introduction to machine learning for brain imaging. Neuroimage 56, 387–399. https://doi.org/10.1016/j.neuroimage.2010.11.004
Lotte, F., Congedo, M., Lécuyer, A., Lamarche, F., Arnaldi, B., 2007. A review of classification algorithms for EEG-based brain–computer interfaces. J. Neural Eng. 4, R1–R13. https://doi.org/10.1088/1741-2560/4/2/R01
Navarro Gil, M., Escolano Marco, C., Montero-Marín, J., Minguez Zafra, J., Shonin, E., García Campayo, J., 2018. Efficacy of Neurofeedback on the Increase of Mindfulness-Related Capacities in Healthy Individuals: a Controlled Trial. Mindfulness (N. Y). 9, 303–311. https://doi.org/10.1007/s12671-017-0775-1
Nesvold, A., Fagerland, M.W., Davanger, S., Ellingsen, Ø., Solberg, E.E., Holen, A., Sevre, K., Atar, D., 2012. Increased heart rate variability during nondirective meditation. Eur. J. Prev. Cardiol. 19, 773–

780. https://doi.org/10.1177/1741826711414625
Olofsson, J.K., Nordin, S., Sequeira, H., Polich, J., 2008. Affective picture processing: an integrative review of ERP findings. Biol. Psychol. 77, 247–65. https://doi.org/10.1016/j.biopsycho.2007.11.006
Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, C., Thirion, B., Grisel, O., Blondel, M., Prettenhoffer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., 2011. Scikit-learn: Machine learning in Python. J. Mach. Learn. Res. https://doi.org/10.1007/s13398-014-0173-7.2
Ros, T., Théberge, J., Frewen, P.A., Kluetsch, R., Densmore, M., Calhoun, V.D., Lanius, R.A., 2013. Mind over chatter: plastic up-regulation of the fMRI salience network directly after EEG neurofeedback. Neuroimage 65, 324–35. https://doi.org/10.1016/j.neuroimage.2012.09.046
Scheeringa, R., Fries, P., Petersson, K.-M., Oostenveld, R., Grothe, I., Norris, D.G., Hagoort, P., Bastiaansen, M.C.M., 2011. Neuronal Dynamics Underlying High- and Low-Frequency EEG Oscillations Contribute Independently to the Human BOLD Signal. Neuron 69, 572–583. https://doi.org/10.1016/J.NEURON.2010.11.044
Schmidt, E.A., Schrauf, M., Simon, M., Fritzsche, M., Buchner, A., Kincses, W.E., 2009. Drivers' misjudgement of vigilance state during prolonged monotonous daytime driving. Accid. Anal. Prev. 41, 1087–1093. https://doi.org/10.1016/j.aap.2009.06.007
Smith, S.M., Nichols, T.E., 2018. Statistical Challenges in "Big Data" Human Neuroimaging. Neuron 97, 263–268. https://doi.org/10.1016/j.neuron.2017.12.018
Tang, Y.-Y., Posner, M.I., Rothbart, M.K., 2014. Meditation improves self-regulation over the life span. Ann. N. Y. Acad. Sci. 1307, 104–111. https://doi.org/10.1111/nyas.12227
Taulu, S., Kajola, M., Simola, J., 2004. Suppression of interference and artifacts by the Signal Space Separation Method. Brain Topogr. 16, 269–75.
van Lutterveld, R., Houlihan, S.D., Pal, P., Sacchet, M.D., McFarlane-Blake, C., Patel, P.R., Sullivan, J.S., Ossadtchi, A., Druker, S., Bauer, C., Brewer, J.A., 2017. Source-space EEG neurofeedback links subjective experience with brain activity during effortless awareness meditation. Neuroimage 151, 117–127. https://doi.org/10.1016/j.neuroimage.2016.02.047
Wolpaw J., Birbaumer N., McFarland D., Pfurtscheller G., Vaughan T., 2002. Brain-computer interfaces for communication and control. Clin. Neurophysiol.

Supplementary figure legends

Suppl. Fig. 1. Schematic representation of the feature extraction in (A) spectral and (B) connectivity approaches.

Suppl. Fig. 2. Distribution of the correlation coefficients for different frequency bands in connectivity approach.

Suppl. Fig. 3. Classification accuracy for swapped training and testing sessions.

Suppl. Fig. 4A. The spatial maps and spectral profiles of group-level independent components (1 to 20).

Suppl. Fig. 4B. The spatial maps and spectral profiles of group-level independent components (21 to 38).

Suppl. Fig. 5. (A, B) Four largest individual classification coefficients associated with independent components for spectral classifier in (A) MF_FP and (B) MF_EP tasks.

579 Suppl. Fig. 6. (A, B) Four individual independent components with highest degree node for
580 connectivity classifier in (A) MF_FP and (B) MF_EP tasks.