Facilitating inhibition might inhibit facilitation;

a mental fatigue study

Martijn van Ackooy
BCN research master, C-track
Rijksuniversiteit Groningen

Supervised by
prof. dr. M. M. Lorist
dr. R. J. Renken

November, 2017

Abstract
The current study tries to identify changes in centrality of cortical components related to mental fatigue. Participants were subjected to a scanning session lasting two hours, consisting of either four blocks of modified Eriksen flanker task (fatigue) or two blocks of modified Eriksen flanker and two blocks of relaxing nature documentary watching (control). During scanning, both anatomical and functional scanning was performed. Cortical components were obtained through independent component analysis of a temporally concatenated dataset containing the preprocessed functional data of all subjects and sessions. Adjacency matrices of components within subject and session were constructed based on the time series uniquely defined for each component per subject and session. Subsequently, both positive and signed eigenvector centrality were computed and differences in changes in centrality over time between the control and fatigue groups were identified. Five components show significant different in positive centrality changes over time between groups, however, none of these effects survived multiple comparison correction. Signed centrality revealed additional information, but was deemed too unreliable for further analysis in the current study.

Key words: Mental fatigue; fMRI; ICA; Eigenvector centrality

1. Introduction

In a rapidly developing world where technology and machinery grows more complex by the year, human error can also have consequences that grow ever more complex. One of the leading causes of accidents in the workplace is mental fatigue. The effect is so large, and so predictable that large companies employ risk-prevention strategies based on fatigue effects amongst employees (for an example see: Sherry and Philbrick, 2016). The psychological sensation of fatigue, accompanied by a bodily feeling of weariness, is known to every single individual on the planet, which is why it deserves the proper amount of attention. Decades of work have already lead to a fair understanding of the underlying mechanisms of fatigue, and a clear description of several separate parts of the process of fatigue induction already exist, yet a single overarching description of the development of fatigue remains elusive.

Fatigue has been described as a psychobiologic al state associated with decreased efficiency and unwillingness to work (GrandJean, 1979), and as an acting stop signal intended to ensure effective resource management and prevention of full depletion (Hopstaken, van der Linden, Bakker & Kompier, 2015). It is a process starting with subjective sensations of effort and fatigue, leading to objective performance decline (Sun, Lim, Meng, Kwok, Thakor & Bezerianos, 2014). The same authors note that it is likely a defensive strategy for scenarios when excessive load is put on the frontal executive system. It is more than simple deterioration of task performance, as compromised performance monitoring is also observed. This decreased efficiency of performance monitoring is thought to be a result of changes in network connectivity, specifically changes in synchronization levels between cortical regions. When following GrandJean's perspective of Mental Fatigue as a psycho*biological* state change, and combining it with the specific proposed role of mental fatigue acting as a stop signal to modulate resource spending, we quickly arrive at the theoretical notion of mental fatigue being a grand scale inhibitory signaling strategy. To further elucidate the underlying

neural mechanisms involved in the process of fatigue, first it is necessary to understand what is already known about mental fatigue. There are several distinct fields from which we can gather a large amount of evidence towards the specific principles that govern fatigue development.

The first line of evidence is provided by behavioral studies. From these we can gather that fatigued individuals are very likely to show an increase in reaction times (Scheffers, Humphrey, Stanny, Kramer and Coles, 1999; Boksem, Meijman, and Lorist, 2005; Lorist, Boksem, Ridderinkhof, 2005; Lorist, Bezdan, ten Caat, Span, Roerdink, and Maurits, 2009; Zhao, Zhao, Liu, and Zheng, 2012; Guo, Chen, Zhang, Pan, and Wu, 2016), are more prone to making errors compared to well-rested individuals (Scheffers, et al., 1999; Boksem et al., 2005; Lorist et al., 2009; Faber, Maurits, Lorist, 2012; Hopstaken, van der Linden, Bakker, and Kompier, 2015; Guo et al., 2016), and are less attentive of the task at hand (Guo et al., 2016; Li, Song, and Miao, 2016). Van der Linden and Eling (2006) report that fatigued individuals have a reduced local processing as opposed to global processing of stimuli. Another study indirectly showed this reduced attention in fatigued individuals by providing evidence on direction of gaze, where fatigued individuals spent more time looking away from the task at hand (Hopstaken, van der Linden, Bakker, Kompier, and Leung, 2016). Combined with an increase in error rates this suggests that conflict monitoring and attentional control systems might perform suboptimally under conditions of mental fatigue, where the distinction between relevant and irrelevant stimuli blurs and participants allow themselves, whether consciously or subconsciously, to look around and fixate on task irrelevant features of their surroundings. Subsequent multilevel correlational analysis showed the direction of gaze was not only correlated with subjective attention, but a direct correlation between the gaze direction and other measures was observed as well. This provides evidence towards a possible disassociation between changes in several measures as time on task progresses, and the direct relationship with mental fatigue. For instance, when individuals had their gaze directed towards the task, it was correlated with an increase in pupil diameter, even though a different study found an overall decrease of pupil diameter as a function of time on task (Hopstaken et al., 2015). The same effect shows for the measure of task performance. Task performance generally declines with time on task, but Hopstaken et al. (2016) found a negative correlation between task performance and time spent gazing away from the screen. This provides us with additional evidence that the decline in task performance is not necessarily related to mental fatigue, but rather an emergent property of fatigue due to individual's increase in task-unrelated gazing. As such, we might state that a difference in gaze behavior underlies the decrease in task performance, the decrease in subjective attention and the decrease in pupil diameter.

The second line of evidence comes from investigations on the temporal dynamics of cortical activation as measured in electrophysiological signals through EEG, subsequently analyzed in the event-related potential style. Most important here is the relationship between mental fatigue and the amplitude of the P300. First of all, it follows from the work of Datta, Cusack, Hawkins, Heutink, Rorden, Robertson and Manly (2007) that measuring P300-related changes is accompanied by an increase in the measure of error rates. Following that, the amplitude of the P300 has been observed to decrease as a function of time on task in general (Zhao, Zheng, Zhao, and Liu, 2010; Zhao, Zhao, Liu, and Zheng, 2012; Möckel, Beste, and

Wascher, 2015; Hopstaken, van der Linden, Bakker, and Kompier, 2015; Hopstaken, van der Linden, Bakker, Kompier, and Leung, 2016; Guo, Chen, Zhang, Pan, and Wu, 2016), and P300 amplitudes are positively correlated with the times when individuals showed task-directed gazing behavior (Hopstaken et al., 2016), However, the distinction between gaze behavior and fatigue related effects is less sharp, as correlation between gaze towards screen and P300 amplitude can still exist, even when P300 amplitudes tend to decrease over time. Therefore, P300 amplitudes might still be a strong indication towards a state of mental fatigue, and it would serve us well to find evidence of a neural generator for the P300 to see if de-activation in that specific cortical region co-exists with a state of mental fatigue. Li, Wang and Hu (2009) have provided exactly such evidence for an anatomical basis for P300 generation. Anatomical generator sites of the P300 include the Anterior Cingulate Cortex, Parietal (visuospatial task), Middle Frontal, and Superior Temporal regions. In these specific areas, after induction of mental fatigue, we assume to find a deactivation, or reduced activation, to explain the reduced P300 amplitudes. As the P300 is commonly thought to reflect context updating (see Donchin and Coles, 1988), we might associate a reduced amplitude with poorer performance maintenance, or attentional control.

The cingulate cortex shown to be involved in P300 generation has also been linked to generation of the error related negativity (anterior cingulate cortex: Holroyd, Nieuwenhuis, Mars, and Coles, 2004; for a thorough and systematic review on the ERN, see: Wessel, 2012), the N200 (medial cingulate cortex), and is theorized to be part of the performance monitoring system (For example, see: van Noordt and Segalowitz, 2012). Additionally, the ERN has also been shown to decrease in magnitude as a function of time on task (Scheffers, Humphrey, Stanny, Kramer and Coles, 1999; Lorist, Boksem, Ridderinkhof, 2005), and Scheffers et al. also showed the relationship between a reduced ERN amplitude and poorer task performance. Increase of N200 amplitudes as a result of fatigue induction have also been reported (Möckel, Beste, and Wascher, 2015). Lastly, a decrease of amplitudes of the mismatch negativity and the N100 component have been reported as an effect of time on task (MMN: Li, Song, and Miao, 2016: N100: Boksem, Meijman, and Lorist, 2005; Faber, Maurits, Lorist, 2012). Coincidentally, there is evidence of two separate neural generators of the MMN, the superior temporal gyrus, and the inferior frontal gyrus (Alho, 1995). The neural generators of the N100 component are thought to lie in the superior temporal gyrus and the planum temporale (Zouradakis, Simos, and Papanicolaou, 1998; Godey, Schwartz, de Graaf, Chauvel, and Liégeois-Chauvel, 2001). When we combine these pieces of evidence, taking special notice of the elevated error rates fatigued individuals seem to experience along with the decrease of both P300 and ERN amplitudes, it all strongly suggest the involvement of performance monitoring, conflict monitoring and attentional controlling systems in the development of mental fatigue. Based on the aforementioned evidence, we expect to see reduced activation of the anterior cingulate cortex, frontal areas, and temporal areas, and increased activation in the medial cingulate cortex as a result of mental fatigue.

A third line of evidence comes from literature pertaining neural oscillations. Several independent teams of researchers have investigated the link between frequency band specific changes over time and mental fatigue. Instead of looking at the fatigue induced changes in brain activity specifically related to stimulus processing or error detection, time frequency analysis has the advantage of showing changes in cortical activation related to ongoing

processes. Inducing fatigue has been shown to coincide with changes in broad-spectrum EEG data. Both power in the theta frequency band (4Hz – 8Hz) and the alpha frequency band (8Hz – 13Hz) have been shown to increase as a function of time on task (For results pertaining theta frequency band increase see: Boksem, Meijman, and Lorist, 2005; Zhao, Zheng, Zhao, and Liu, 2010; Zhao, Zhao, Liu, and Zheng, 2012; Tanaka, Shigihara, Ishii, Funakura, Kanai, Watanabe, 2012; Wascher, Rasch, Sänger, Hoffmann, Schneider, Rinkenauer, Heuer, Gutberlet, 2014; for results concerning alpha, see: Boksem et al., 2005; Zhao et al., 2010; Zhao et al., 2012; Shigihara, Tanaka, Ishii, Kanai, Funakura, and Watanabe, 2013; Wascher et al., 2014; Fan, Zhou, Liu, and Xie, 2015; Xie, Xu, Wang, Li, Han, and Jia, 2016). Power in the beta frequency band (13Hz-25Hz) has been shown to decrease as a function of time on task (Zhao et al., 2010; Zhao et al., 2012; Tanaka et al., 2012; Shigihara et al., 2013; Fan et al., 2015). Lastly, power in the delta frequency band has also been investigated in one paradigm, where it was shown to decrease as a function of time on task, specifically in the dorsolateral prefrontal cortex and the posterior cingulate gyrus (Ishii, Tanaka, Watanabe, 2014). A positive correlation between activation in the dorsolateral prefrontal cortex and posterior cingulate cortex was also reported.

After disentangling the different frequency band specific effects, one might consider each of these effects in turn to better elucidate the mechanisms underlying the development of fatigue. Klimesch (1999) notes that, in order for proper task performance to occur, the expected behavior of the alpha and theta frequency bands are counterbalanced. He states that good performance is associated with an increase in alpha activation and a decrease of power in the theta-frequency band. The evidence from mental fatigue paradigms is in line with this statement. Both alpha and theta show an increase in power after fatigue induction in several different studies, leading to a different alpha/theta ratio, and there is ample evidence of poor performance in fatigued individuals. Klimesch (1999) subsequently writes that the reported findings are associated with information encoding and memory systems in general. The alpha frequency band (specifically, the upper alpha frequency band) is thought to reflect search and retrieval processes in semantic long-term memory, and the theta frequency band is thought to reflect novel information encoding. However, more recent work on alpha-frequency band oscillations show a more complex role for the alpha-frequency band, with speculated roles including cortical inhibition of task irrelevant regions, cortical idling, regulating the locus of attention, and binding within the global neuronal workspace (Palva, and Palva, 2012). It is beyond the scope of the current article to delineate the exact role of the alpha-frequency band oscillations in the brain, yet it must be noted that the role of the alpha frequency band is obviously widespread and important for cortical information processing, attention, and consciousness in general. Therefore, it is unsurprising to see effects of time on task on the alpha frequency band specific power in mental fatigue paradigms. We expect to see changes in event-related synchronization over cortical areas which are not critical to the maintenance of homeostatic equilibrium or task maintenance. It is also expected that task relevant areas, and conflict monitoring areas show event related desynchronization within the alpha frequency band. Lastly, we expect that, in the current paradigm, this translates into increase in activation in task-relevant areas and areas which are critical for maintaining internal homeostasis, and a decrease in activation in task irrelevant areas. Stated otherwise, this means that when brain states are written in terms of linear summations of localized states, the weight

coinciding with task-relevant areas will increase, whereas the weight related to task-irrelevant areas will decrease.

The beta frequency waves shows an inverse relationship with alpha frequency band activity, appearing mostly in wakeful and conscious states, rather than the relaxed state with eyes closed in which alpha is generally seen. Oscillations in the beta frequency band have been associated with gamma-aminobutyric acid based signaling. For example, see the induction of beta-waves after benzodiazepine ingestion (Feshchenko, Veselis, and Reinsel, 1997), and the relationship between benzodiazepines and GABA (Olsen and DeLorey, 1999; Olsen and Betz, 2006). From this we can postulate that beta-wave activity is associated with inhibitory processing, as GABA is the main inhibitory neurotransmitter of the central nervous system. It has already been shown that an increase in GABA, for example, by oral administration, decreases beta waves and increases alpha waves (Abdou, Higashiguchi, Horie, Kim, Hatta, and Yokogoshi, 2006). Furthermore, beta frequency band oscillations have been shown to relate to response inhibition (Huster, Enriquez-Geppert, Lavallee, Falkenstein, Hermann, 2013). However, it has also been shown that beta band activity reflects arousal of the visual system during increased visual attention (Wróbel , 2000), and has been specifically linked to emotional context. An increase in beta band activity was seen in dogs when they expected a picture of a rewarding piece of meat (Da Silva, van Rotterdam, van Leeuwen, and Tielen, 1970), and increased beta power was also observed as a response to emotionally negative stimuli (Güntekin and Basar, 2007, 2010; Woodruff, Daut, Brower and Bragg, 2011). Lastly, a strong link exists between beta power, somatosensory and motor functions (Pfurtscheller, Stancak, and Neuper, 1996). More specifically, beta band activity has been theorized to be a signature of an active promotion of the existing motor set, whilst compromising neuronal processing of new movements, following the discovery that beta band power is greatest during holding periods following voluntary movement, and its suppression or even replacement by gamma oscillations during planning and execution of movement. Combining these disparate sources of information provides us with evidence of the widespread utility of beta frequency band oscillations, and it remains unclear what the specific role of the change in beta-frequency band power in the context of fatigue development is. One might postulate the decrease of beta frequency band activity to reflect the decrease of emotional intensity experienced by individuals in states of mental fatigue. Alternatively, it can be postulated that the observed decrease in beta power is related to a reduced inhibition strength of task irrelevant motor actions, reflected in the increase in error rates coming from erroneous, and involuntary button presses.

Two specific points of interest to notice are the following. Although in general alpha power increases and beta power decreases, there are studies providing more localized results which show the opposite effect. A decrease in alpha power in the right middle and superior frontal gyri, and an increase in beta power was observed in the middle temporal gyrus and the inferior parietal lobule by Shigihara, Tanaka, Ishii, Kanai, Funakura, and Watanabe (2013). An increase in beta power in the right inferior and middle frontal gyri was found by Tanaka, Ishii, Watanabe (2014), These results provide evidence towards a theory where the right frontal gyrus might play a special role in the development of fatigue, which coincides with the reports. Secondly, it has been shown that adding an extra incentive can reverse behavioral, event related potential, and frequency band specific effects and return them to baseline values

(Hopstaken, van der Linden, Bakker, and Kompier, 2015; Hopstaken, van der Linden, Bakker, Kompier, and Leung, 2016).

The evidence provided thus far suggests few specific anatomical sites related to fatigue development, but does not yet contain investigation on the behavior of the full brain in fatigue paradigms. Neuroimaging studies can offer more clarity concerning possible neural correlates of fatigue. Increased activation as a function of time on task is reported in the superior temporal gyrus, cingulate regions, and inferior frontal regions (Cook, O'Connor, Lange, and Steffener, 2007). Cook et al. also reported increased activation in the hippocampus, thalamus, cerebellar clusters and the cerebellar vermis, and decreased activation was observed in the posterior parietal cortex. The decrease in posterior parietal is explained in terms of the posterior attentional network. Changes in general monitoring systems and attention related systems were coupled to the observed increase in activation of the frontal regions, cingulate regions and cerebellar regions. A second functional MRI study tried to expand the knowledge on fatigue through compensation strategy research, and reported midbrain deactivation specifically under conditions of compensatory effort (Nakagawa, Sugiura et al., 2013). These authors subsequently provided a large set of areas which show a decreased activation in fatigued individuals, split over two conditions. For sake of simplicity, only the areas showing changes in the hard condition will be described below, assuming that the hard condition was more fatiguing than the easy condition. Areas showing highly significant decreases in activation consist of right superior and bilateral middle frontal gyri, left superior temporal gyrus, left middle occipital gyrus, right fusiform gyrus, right precuneus and the right posterior lobe of the cerebellum. In the fatigue condition, the midbrain deactivation was also observed, suggesting compensatory effort to play a role in the development of fatigue. Note the discrepancy in direction of activation reported for the superior temporal gyrus. As we have seen that the P300 and MMN amplitudes decrease, and the superior temporal gyrus has been put forward as a generator for both of these phenomena, we expect to see a deactivation in this region, in line with what was observed by Nakagawa et al. (2013).

Having looked at results using different techniques, Ullsperger and von Cramon (2001) furthered the knowledge on mental fatigue by making use of the fact that there is an intricate relationship between the temporal and spatial domain of cortical activation. They specifically used both EEG and fMRI to investigate the internal monitoring of performance. Due to the increased error rates and increased reaction times reported several times, it is quite clear that performance suffers when fatigue grows. Ullsperger and von Cramon postulated their hypotheses concerning performance monitoring in a framework containing both neural oscillations and anatomical regions, and tried to distinguish between neural correlates of response competition and error processing. They subsequently tested these hypotheses in a paradigm where both dimensions were measured through a combination of EEG and fMRI data. During a Flanker task, response competition was thought to be highest in the compatible trials compared to the incompatible trials. To disassociate between error processing and response competition, all trials which were answered incorrectly were left out of the contrast aimed at revealing response competition specific processes. Furthermore, the contrast created to investigate error processing specific effects consisted of incompatible trials only, where the contrast revealed differences between trials answered correctly, and those answered incorrectly. The response competition specific contrast revealed higher activation in the

mesial superior frontal gyrus when there was a stronger contrast between correct answer and incongruent flankers. This effect extended into regions surrounding the anterior cingulate sulcus, and was observed in the posterior cingulate cortex. Lateral effects of response competition include the lateral premotor cortex, right inferior frontal gyrus, and areas around the right intraparietal sulcus, and bilaterally in the anterior superior insula and on the middle frontal gyrus. Coincidentally, the error processing contrast revealed higher activation during erroneous trials in pre-supplementary motor area, regions surrounding the anterior cingulate sulcus, the anterior insula, and the regions surrounding the intraparietal sulcus. The subsequent ERP-analysis showed a strong distinction between activation patterns of response competition and error processing. Systems governing response competition and error processing show high overlap, but can be separated into two distinct networks. Both of these networks contain cortical areas reported in several experiments designed to investigate the effects of mental fatigue, giving rise to the suspicion that error processing performs suboptimally in fatigued individuals, and response competition is naturally stronger in fatigued individuals. The latter point is carefully stated in light of the gaze-direction effects mentioned earlier, and the increase in error rates seen in fatigued individuals. We expect to find changes in either if not both of these networks during the development of fatigue.

Furthering this line of thinking, lastly, we consider a complete model specific to mental fatigue, as postulated by Ishii, Tanaka, and Watanabe (2014). Instead of trying to distinguish specific cortical areas related to fatigue development, their aim was to redefine the effects reported in the literature in terms of balancing acts between inhibitory and facilitating systems. Their work tries to first make the distinction between these two systems, and subsequently tries to unify their interactions in one model of mental fatigue. The facilitation system consists of the limbic system, basal ganglia, thalamus, orbitofrontal cortex, dorsolateral prefrontal cortex, anterior cingulate cortex, premotor area, supplementary motor area, and M1. Sensory information from the peripheral nervous system is used as a biofeedback mechanism to indicate periods of fatigue, and the facilitation system then inhibits motor output. Conversely, adding an incentive will reverse the effect and stimulate motor output. This reversal is strongly reminiscent of the reversal of mental fatigue related effects as reported by Hopstaken, van der Linden, Bakker, and Kompier (2015), and Hopstaken, van der Linden, Bakker, Kompier, and Leung (2016). The inhibitory system consists of thalamus, secondary somatosensory cortex, insular cortex, posterior cingulate cortex, anterior cingulate cortex, premotor area, supplementary motor area, and M1, and its role is to either block novel or to suppress ongoing activity of motor units. Notice the similarities between the constituents of both networks, and with the areas reported through neuroimaging studies on fatigue. The ideas presented above provide a picture of anatomical sites in which we expect to see changes in activation patterns during fatigue, but more importantly, it provides us with the idea of a balancing act between cortical networks making use of a shared, and limited, resource pool.

This resource management is not specific to mental fatigue. Interference resolution, i.e., the solution to an internal conflict between two independent inputs vying for the same resources, has been studied in an fMRI paradigm by Zhu, Zacks, and Slade (2010). Their specific goal was to study differences between younger and older adults, however, their paradigm still provides us with valuable information concerning interference systems. Using a flanker task to introduce interference, the authors found a distinctive difference in inferior and

middle frontal gyrus between congruent and incongruent conditions. More interestingly was the clear right-sided lateralization of the correct trials as opposed to bilateral activation on error trials. The current paradigm is not so lenient towards disassociating the networks in question. For this, an alternative solution must be found. Following the theory on the two-sidedness of mental fatigue, we expect to see these two balancing forces in action in a simple principle. Due to the large cortical scale, and high importance of these networks, we expect to find parts of these networks in a decomposition of the entire cortical activation pattern into its principal components. Subsequently, after directing ourselves to the language of graph theory, we expect to see changes in this balancing act. We specifically expect to see an increased influence of the ACC, PCC, frontal areas, insular cortex, primary-, pre-, and supplementary motor areas, secondary sensory area and subcortical regions, reflecting both inhibition and facilitation systems becoming more active when fatigue develops.

2. Methods

2.1. Subjects

Thirty-seven subjects took part in the study. Scanning was performed using a Philips Intera 3.0T scanner at the NeuroImaging Centre, Groningen. All participants filled out a questionnaire to verify their MR-compatibility. All subjects had normal or corrected to normal vision, and were right-handed. Participation was voluntary, and participants received compensation in the form of money (15 euros), or course-credits. Prior to the study, informed consent from each participant was obtained.

2.2. Experimental Procedure

Participants were randomly assigned to the control or fatigue condition. Those in the fatigue condition were asked to perform a modified version of the Eriksen Flanker task (Eriksen and Eriksen, 1974) for a period of four blocks, lasting approximately half an hour each (For a clear description of the task, see Faber, Maurits, and Lorist, 2012). Participants in the control condition performed the modified Eriksen Flanker task for the first and last block of the scanning session and watched a minimally-stimulating nature documentary in between.

2.3. fMRI data acquisition

For each subject, four functional scans and one anatomical scan were acquired. Each anatomical scan consisted of 160 slices and voxel sizes of $1mm^3$. Echo planar imaging was used for functional scanning, with a TR of 2 seconds. Each functional scan consisted of 39 slices, scanned using an interleaved slice acquisition order, with a voxel size of $3,5mm^3$, and each functional set consisted of 780 scans.

2.4. Preprocessing

Data of the first and last block of all participants was subjected to a standard batch of preprocessing steps. Functional data was subjected to a high-pass temporal filter with a 100 second cut-off. Slice-timing correction was applied, and FMRIB's linear registration tool was used for motion correction (Jenkinson and Smith, 2001; Jenkinson, Bannister, Brady and Smith, 2002). Subsequently, data for each individual was coregistered to their respective

anatomical data. Anatomical data was segmented into brain and non-brain tissues, and the brain extracted images were normalized to standard space after which the normalization parameters were applied to the functional data to register all functional datasets to the standard space as defined by the MNI-152 template. Spatial smoothing was performed using a 3mm FWHM kernel. The data was subsequently concatenated across the temporal domain and group independent component analysis was performed to identify spatial maps of components which all participants, regardless of group, shared, whilst retaining information on optimal timecourses per component on the level of the individual. i.e. which set of unique, independent variables could be used to best explain the data of all participants. For this purpose, FSL's MELODIC was used with implemented automatic dimensionality estimation prior to independent component analysis. During the ICA, voxel-wise demeaning was performed, and voxel-wise variance normalization was applied. MELODIC identified 207 independent components (IC's) with explained variance percentages ranging from 0.60% to 0.41%.

2.5. Component classification

We subsequently split the independent components into noise and signal components based on several objective criteria described below, and a subjective classification of what defines a component of interest based on earlier agreements within the field (for an up to date in depth review of what constitutes noise and signal within fMRI data, see Griffanti et al., 2016) .

2.5.1. Binary tissue masks

To ensure correct identification of components which represent cerebrospinal fluid, white matter tracts and grey matter, overlap scores were computed between each IC's spatial map and a mask corresponding to one of the three tissue types. To minimize misclassification of components as noise due to large overlaps with small structures such as the low percentages of white matter still found in the grey-matter dominant regions of the brain, binarization of the tissue maps was performed prior to computing overlap values. Binary masks were created by using SPM's segmented MNI-152 template tissue priors, and subsequently performing binarization of the tissue specific masks based on thresholding scores, thereby representing each voxel by one or zero. All scores in the standard space masks ranged from zero to one, and each mask showed its own characteristic distribution of intensity values over voxels. Therefore, we decided to base each threshold on visual inspection of sorted voxel values seen in figure 1.

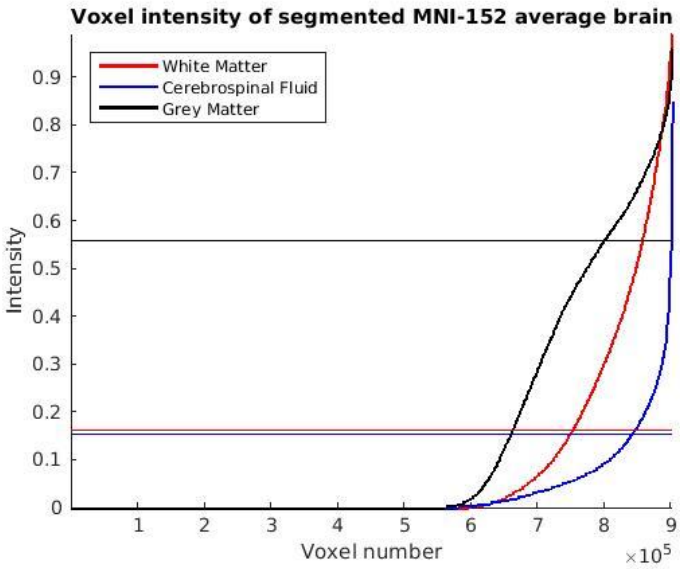

Figure 1. Single voxel intensity values per mask for all voxels. Thresholds were chosen to be approximate to the 'elbow' of each graph, and are shown per tissue type as the horizontal line of matching color.

For CSF and white matter, an appropriate threshold was easily found, and was set to be 0.1529 for CSF and 0.1608 for white matter. All voxels showing intensity higher than these values were set to one and everything else to zero. For grey matter, it was harder to establish an optimal threshold, due to the non-logarithmic nature of the curve, and was eventually set to 0.5569. For ease of further handling of overlap scores, the binary mask for grey matter was flipped in polarity, giving all voxels with intensity *lower* than the threshold a value of one, and all other voxels zero. This can also be thought of creating a map for areas which are high in non-grey matter material. (*Note:* during subsequent sections of the article, references to grey matter overlap will be given from the point of view of true overlap). The result of this last choice is that noise components will have high overlap scores for at least one tissue type, making the distinction between unwanted and wanted components fairly straightforward after overlap scores are computed.

2.5.2. Component-Mask overlap scores

To compute overlap scores, binary tissue masks were first reconstructed into vectors, as well as each component, where information concerning spatial location was represented in the vector form. Subsequently, the inner product between two such vectors was computed by:

$$\vec{v}^T\vec{w} = x = \sum_{i=1}^{n} v_i\, w_i \vee$$

With $\vec{v}^T$ being the transposed binary mask data, $\vec{w}$ the absolute form of the component under investigation, and x representing the projection of $\vec{w}$ on $\vec{v}^T$. Here, x can be seen to represent the amount of overlap between a component and a mask.

The result of this computation was a set of three vectors, one for each tissue type, composed of overlap scores for each independent component. See figure 2 for the sorted overlap values per tissue type.

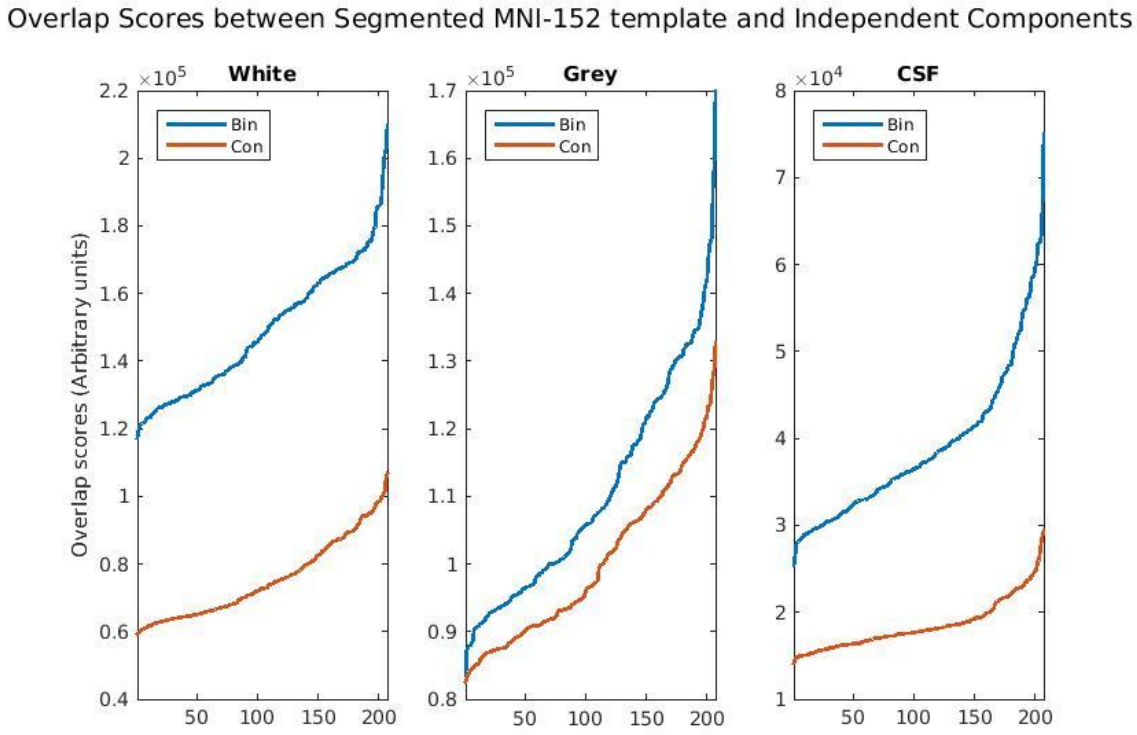

Figure 2. Overlap scores between the segmented tissue priors of the MNI-152 template with the maps of the independent components. For sake of comparison both tissue priors consisting of continuous values, shown in red, as the thresholded and subsequently binarized tissue priors, shown in blue, have been used to compute overlap scores.

For each individual tissue type, there is a clear plateau of 'normal' overlap, and a set of components showing markedly higher overlap scores for each tissue type. It is this region of high overlap that was used to define components which white matter dominant, CSF dominant. Inversely, for grey matter, those components showing the *lowest* overlap scores were selected for pruning. A set of components showing intermediate overlap with both CSF and WM masks would normally be ignored, but when there was little to no grey matter present, it was assumed that these components instead resemble edge regions between tissue types and thus, did not carry information relevant for further analysis.

In a more general setting, it is preferable to also include analysis of temporal information, to identify unwanted components based on saw-tooth patterns in IC temporal signals. Due to our choice of temporally concatenating datasets before ICA, timecourse data was only available on the individual level. However, the concatenated timecourses could still be converted into frequency spectra, which were subsequently used as a source of additional information, where noisy components were expected to show diffuse spectra with higher energy in higher frequency bands.

After these initial pruning steps, the remainder of components was subjected to visual inspection for verification of proper cleansing. During visual inspection, care was taken to identify components related to edge/motion activity, and components related to physiological activity which were possibly left unidentified using the previous pruning steps. Of the 207 initial components, 91 were classified as signal of interest and selected for further analysis.

After successful data reduction, further analysis was performed on remaining components based on their temporal information, which is unique per session and per participant. Spatial patterns emerging from ICA are not represented at each point in time for each individual and variations in component activity over time can be used to create a more global map of cortical activation patterns. Timecourses represent temporal activity of one spatial component, and can therefore be used to distinguish these global patterns of activity and variations over time.

2.6. Adjacency scores/Linear Correlations

To describe these possible changes in global connectivity over time, we assume that there is some truth to Hebbian theory, and that what fires together, wires together. After combining this idea with that of a graph theoretical view of networks, the next step towards our goal of describing dynamical systems is that of connectedness, or adjacency within networks (Newman, 2008). Building on an already solid mathematical framework, which has been tested and validated in various applied settings countless times (e.g. Organic Chemistry: Estrada, 1996; Steganography: Pevny, Bas & Fridrich, 2010; Social network analysis: Freeman, Borgatti, & White, 1991), the validity of constructing adjacency matrices in the current context depends solely on the assumption that when two separate regions, or networks of the brain are highly active at the same time they are highly likely to be connected. The temporal information captured in the timecourse data for each component provides this information readily.

Consequently, the correlation between all timecourse pairs within subject and session was computed to provide an index of connectedness. This resulted in an *n* x *n* correlation matrix per participant, per session, where *n* is the number of components that survived cleansing steps. This *adjacency matrix* is a real-valued, symmetric matrix, representing a network where weighted connections are independent of direction. i.e. $r_{i,j} = r_{j,i}$ where r is the correlation coefficient, and *i* and *j* are adjacency matrix indices referring to component numbers, and r is computed as follows:

$$r = \frac{1}{(N-1)\sigma_x\sigma_y}\sum_{i=1}^{N}(x_i - \overline{x})(y_i - \overline{y})$$

Where $i = 1,2,\dots,N$ is the length of the time series of components x and y, σ denotes standard deviation of a time series and an overbar denotes the mean of a time series.

The symmetry imposed is a natural consequence of using correlation as a measure of connectivity, as correlation does not imply causation. Activity in component A is related to activity in component B. However, a measure of correlation does not provide the direction of information flow between components A and B. See figure 3 for an example of an adjacency matrix.

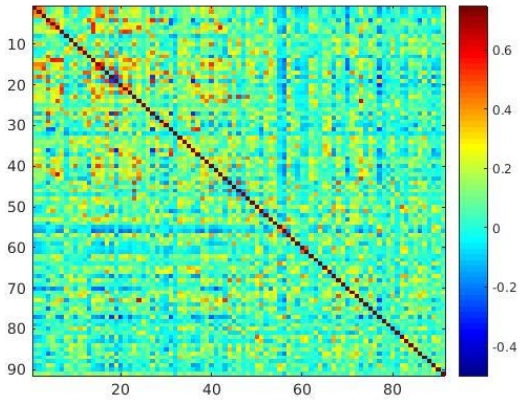

Figure 3. Adjacency matrix over all components, mean over all participants and sessions.

By revealing connections between components, and carrying information on natural clusters of components, these adjacency matrices provide a useful tool for further connectivity studies. Describing which components have a lot of strong connections, and which are weakly connected can provide a plethora of information on the natural flow of information through network structures. Several methods have been developed to reduce the complexity of these systems by viewing only a subset of possible variables. With the aim in mind of describing the large scale changes within these networks, the current study makes use of the measure of Eigenvector centrality to investigate possible changes in the influence a single node exerts in the entire network structure. Subsequently, expansion to a more general case of graph eigenvector analysis is employed as an additional form of dimensionality reduction.

2.7. Eigenvector Centrality

Quantitative statements about the internal workings of networks can be phrased in the terminology of graph theory. A graph is a set of nodes and edges, often explained using the analogy of a social network of friends. Each group of friends can be imagined using a graph, where each individual is a node, and edges represent whether two people are friendly towards each other. Some friends only really talk to one or two others in the group, whereas others are friendly towards everybody involved. Yet, it is clear that the whole structure consisting of individuals is one coherent group. One individual often does not like all of his or her friends equally well, leading to a ranking described in terms of liking ratings, or a similar subjective measure. When encoding these subjective liking ratings, we get a graph with weighted edges. The question then is which friend is the most important player in this specific group. Specifically, what part of the group's activities as a whole can be explained by the activity or presence of a single individual? How heavy does his vote count compared to the vote of others? How strong does her opinion resonate throughout the group? This way of computing social influence of an individual strongly resembles independent component analysis, where statements are made about percentages of variance explained by single components. In essence, computing eigenvector centrality is a close cousin to ICA and PCA.

For more detailed information on the measure of eigenvector centrality, see: Newman (2006) for a theoretical description; Lohmann, Margulies, Horstmann and Pleger (2010) for a

practical example in fMRI data analysis. For a possible implementation in Matlab, see: Brain Connectivity Toolbox (Rubinov and Sporns, 2010). For any real-valued non-negative symmetric matrix, eigenvector centrality is defined as the eigenvector associated to the largest absolute eigenvalue given by the matrix equation:

$$Ax = \lambda x$$

Which states that there exists a vector x, for which holds that it remains unchanged by multiplication with matrix A, with exception of a scaling term, lambda. If the matrix A is non-negative, i.e. all of its entries are positive or zero, the Perron-Frobenius theorem states that the vector x associated to the largest absolute λ consists only of positive entries. In the case of weighted correlation matrices, this means a transformation is required to ensure all matrix entries are non-negative. This introduces a potential risk of skewing information, introducing non-existent information, or discarding information altogether. It would serve the analyst well to carefully consider the different transformations available to ensure non-negativity of all the matrix entries.

Of all the possible rescaling techniques, some have grown to be favorites within the network analysis community. Within neuroscience, adding one to every entry, thereby rescaling all matrix entries to lie between zero and two (For examples see: Preti, Bolton, and Van De Ville, 2016; Wink, de Munck, van der Werf, van den Heuvel, & Barkhof, 2012; Lohmann et al., 2010). This method provides a good approximation of network structures, and provides information on individual components from the perspective of the entire network. A second option is ensuring only positive entries by creating a binary matrix where one denotes there is a connection, and zero denotes that there is no connection between components (for examples, see: Bagler, 2008; Borgatti, 1995; Rosenkopf & Schilling, 2007). This method serves well for those that are only interested in possible connectivity studies and cluster analysis, without giving credit to weights of individual connections. A binary adjacency matrix also has the advantage of being computationally straightforward. Alternatively, one can choose to look at only the positive side, or the negative side of the scale, for an example of such a split where both sides were subsequently analyzed, see Hosseinmardi, Han, Lv and Mishra (2014). However, in the current research there is an interest in preserving naturally occurring clique structures. One such division between cliques in neuroscience is the natural division between types of cortical networks, with a clear distinction between networks powered by different neurotransmitters, networks which are predominantly excitatory in nature or rather inhibitory in nature. The current study uses the first of these aforementioned methods, and also provides a comparison with another method described below, aimed at directly describing the overall status of a component in terms of excitation or inhibition.

2.8. Eigenvector centrality with negative relations

The current investigation not only aims to define which networks or cortical areas are the most influential players, but additionally tries to redefine the centrality measure in order to make a clear distinction between those components whose influence is mainly excitatory and those whose influence is mainly inhibitory. Components negatively correlated with a plethora of other cortical networks are thought to be either inhibitory themselves, or inhibited by everything else. Please note the two-sidedness of this statement, as correlation clearly does not imply causation. To preserve this structure and thus the distinction between positive and

negative correlations, the current research follows the precedent of Bonacich and Lloyd (2004) and uses matrices with negative entries to compute centrality scores. These centrality scores are subsequently compared to those resulting from more traditional metrics on centrality. See Seidel (1968) for a concise description on the properties of graphs with positive and negative entries, and their eigenvector decomposition. The resulting eigenvector is also signed, providing a natural distinction between two cliques. Given the direction of the original correlations, the clique distinction is thought to be related to inhibitory and excitatory signaling. The signed metric provides centrality scores relative to the clique to which a component belongs. This allows for retention of all information contained in the original time courses, without having to apply one of the earlier described transformations. The subtle difference between unsigned centrality, and signed centrality scores lies in the balancing act between positive and negative influence, where a component, or node, which has an equal amount of influence in both groups, will end up as a non-central player, i.e., it is unclear to which group or clique this component really belongs. It is neither liked by everyone, nor disliked by everyone. Its opposing influences balancing results in a player having neutral influence in the network. The resulting metric resembles the structure of centrality measures computed from a split dataset, such as in Hosseinmardi et al. (2014), and can be inspected below.

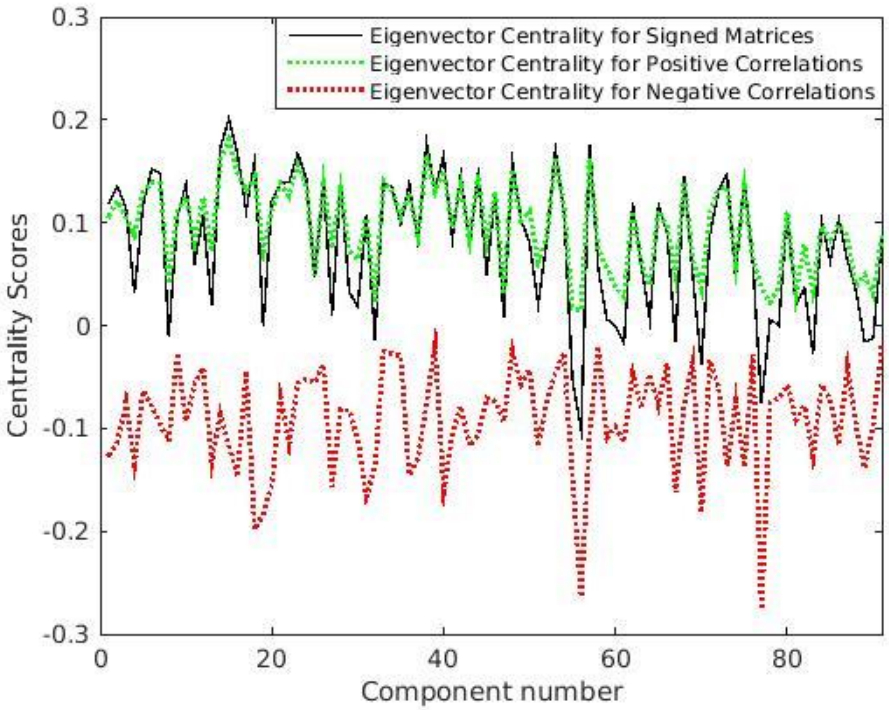

Figure 4. Three eigenvector centrality measures computed from the mean adjacency matrix shown in figure 1. Comparison between defining centrality based on only positive entries, only negative entries, or the entire adjacency matrix.

As can be seen from the signed metric in figure 4, in the current matrix, the network defined by negative entries has a relatively small influence in the global structure. Nonetheless, the influence of negative correlations can clearly be seen in those few components which have an overwhelming negative presence compared to their small positive influences, i.e., those

components which show very little to no positive outgroup communication. Figure 5 shows the natural clustering for the mean of all participants, with two clusters counterbalancing each other, and one cluster whose influence is mainly neutral.

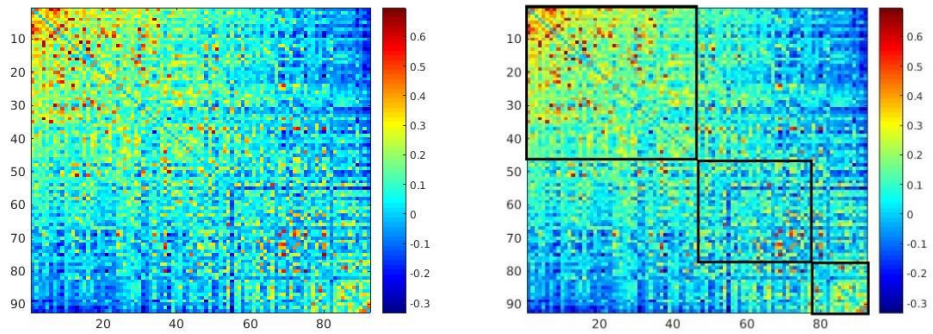

Figure 5. Mean adjacency matrix with components sorted by signed eigenvector centrality scores reveals natural community structures.

As a technical note, one of the usual arguments against using negative matrix entries for computation of eigenvectors is that signed matrices need not have a uniquely defined set of eigenvectors and eigenvalues, rather having a possibility of repetition in the eigenvalue set. However, eigenvector decompositions will only return eigenvalues which are identical or multiples in the case of degenerate matrices, i.e., when one of the vectors of the original adjacency matrix can be written in terms of one of the other vectors. It is expected that the unfortunate scenario of non-uniqueness of eigenvectors does not arise when handling weighted adjacency matrices where each weighted edge is defined up to 32 bit precision. Secondary, it can be shown that the signed metric, in critical ways, behaves the same as the original positive metric. For instance, the following equation holds true for both positive and signed eigenvectors:

$$|x|^2 = 1$$

i.e., the norm-squared (also referred to as the L-2 norm, or the 2-norm) of a vector x is equal to one. All eigenvectors, whether signed or positive, resulting from the decomposition as described above, will have a length of one.

2.9. Principal Component Analysis

To maintain proper nomenclature, it is important to realize that eigenvector centrality, and especially the centrality measure relying on both positive and negative relations is simply a form of principal component analysis. However, the version of eigenvalue decomposition implemented in the Brain connectivity toolbox, and the fast eigenvector centrality mapping toolbox (Wink, 2012; see equations 8 and 9) always ensures a non-negative matrix. To deal with signed matrices, MatLabs own eigenvector decomposition method was used (The MathWorks, 1993), following the QZ factorization algorithm (For more information see Kressner, 2005) and subsequent eigenvector normalization. This method was chosen over MatLab's PCA function for ease of comparison, as the Brain Connectivity Toolbox makes use of the same method, after ensuring non-negativity of the input adjacency matrix.

2.10. Permutation Testing

After all eigenvector sets are computed per participant and per block, a permutation test per component over the eigenvectors is run. Without evidence supporting the assumption of normality of eigenvector components, the median was used for statistical analysis of differences in component distributions. Within condition, the permutation testing made a comparison between centrality during the fourth block and the first block of scanning. Secondary to that, between condition permutation testing was performed for both the first and last block. As a last test, the interaction effect was investigated by taking difference scores between fourth and first block per condition, to identify changes in centrality, and subsequently comparing the distribution of differences between fatigue and control conditions. From 10.000 permutations, p-values per component were computed through the set of equations:

$$C_i = \bar{\bar{j}} \, \forall \{|m_i| > |n_j|\}$$

$$p_i = \frac{(N - C_i)}{N}$$

Where $\bar{\bar{j}}$ is the size (or cardinality) of the set of j for which the following inequality holds true. $|m_i|$ is the true absolute difference between medians of component i, $|n_j|$ is the absolute difference between medians of permutation j. C_i is then the number of occurrences where the difference between permuted medians is smaller than the true difference for component i, N is the total number of permutations, and p_i is the estimated p-value computed for component i.

2.11. False Discovery Rate

Having performed permutation statistics on several hundreds of values, a multiple comparison correction is in order. A multiple comparison correction in the form of a false discovery rate according to Benjamini and Hochberg's (1995) algorithm has been applied.

3. Results

3.1. Interaction

Based on permutation testing, five components show a significant difference between the first and last block of scanning when compared between the fatigue and control condition. Areas showing a significant difference between changes seen in the control and fatigue condition over time can be seen in figure 6 and include: Lateral central sulcus and surrounding gyri (p=0.039; red); Frontomedian areas, including anterior cingulate cortex (p=0.023; blue); Posterior cingulate gyrus (p=0.035; green); Right anteriofrontal areas (p=0.035; pink); Right insular cortex (p=0.033; yellow).

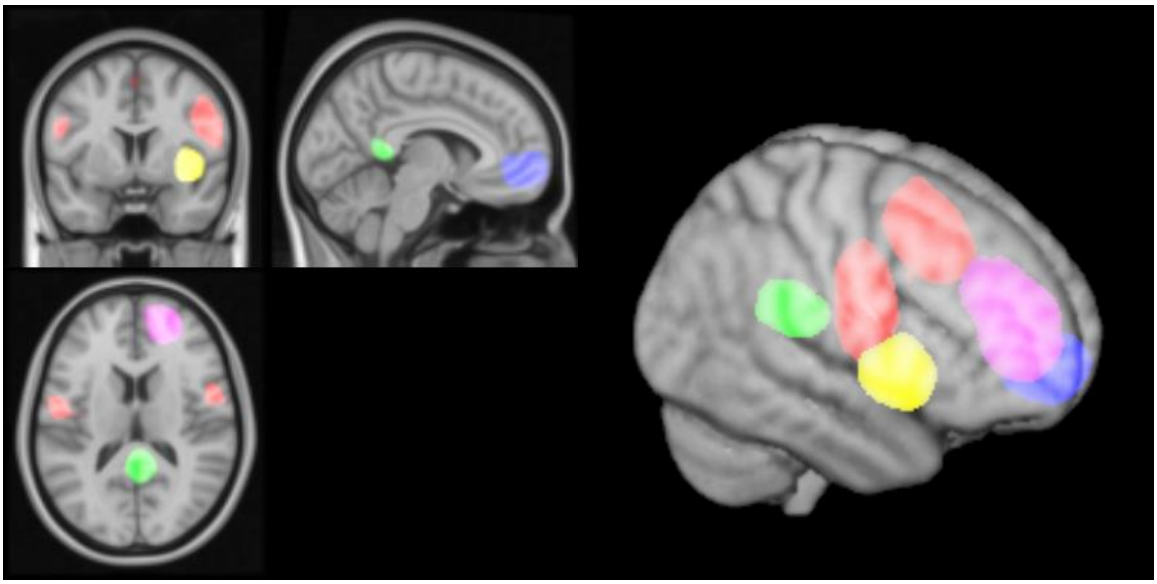

Figure 6. Locations of the five components showing significant differences in changes over time between control and fatigue conditions.

Further investigation on these significant components shows the direction of these changes. See figure 7 below.

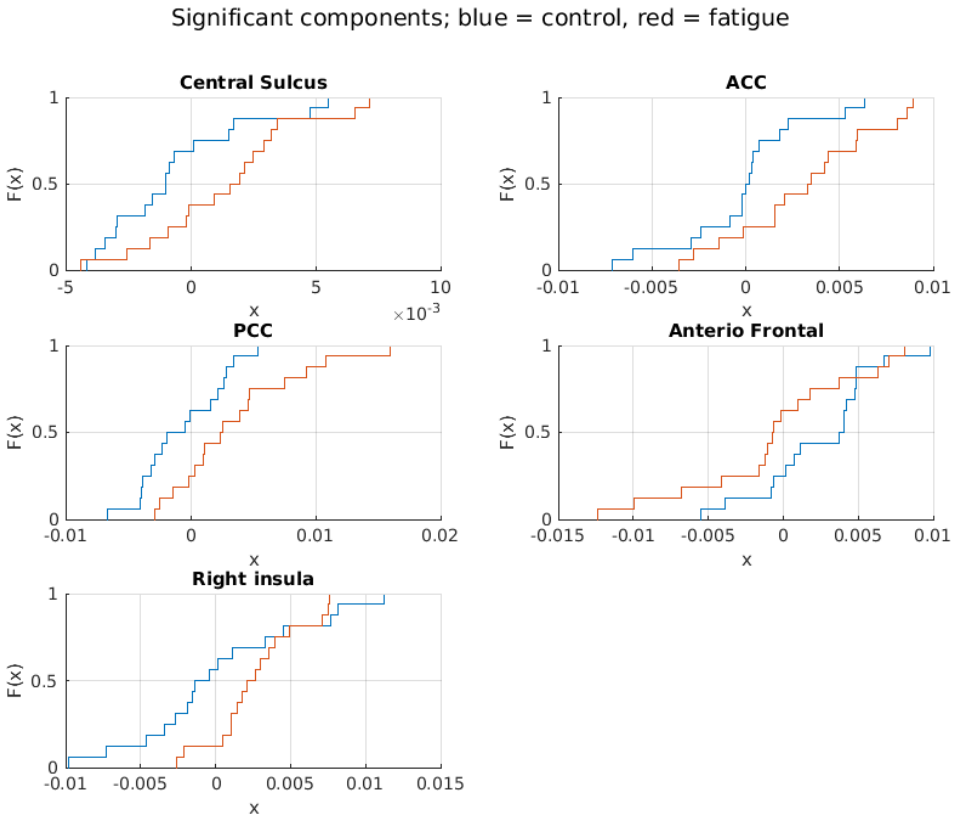

Figure 7. Probability density functions of differences (Difference = Centrality block four - Centrality block one) in the control condition (Blue) and the fatigue condition (Red) for all significant components.

From figure 7 we see that the areas surrounding the central sulcus, the anterior cingulate cortex, the posterior cingulate cortex and the right insula become more central over time in the fatigue condition. The anteriofrontal areas seem to be distributed around no change over time in the fatigue group, rather showing an increase in centrality over time in the control group.

3.2. Control

Within the control condition, no significant differences between block one and four were found.

3.3. Fatigue

Within the fatigue condition, four components were found to differ significantly in their activation between block four and one. Figure 8 shows the regions changing significantly including: Orbitofrontal regions (p = 0.039; blue); Posterior cingulate cortex (p =0.008; red); Lateral occipital cortex and superior parietal lobule (p = 0.033; green); bilateral hippocampal regions (p = 0.012; pink).

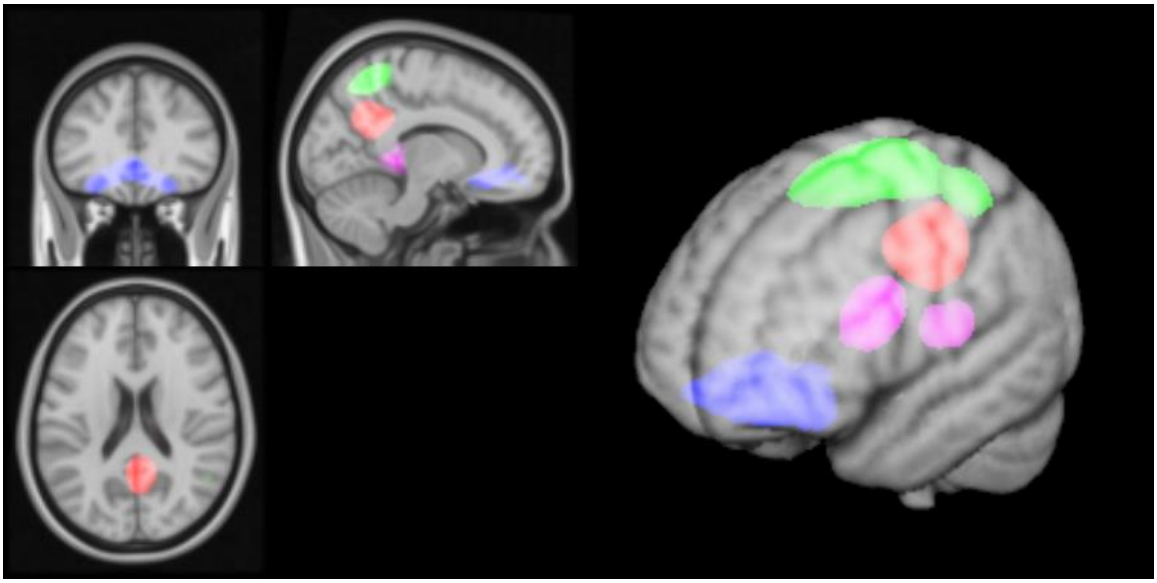

Figure 8. Anatomical sites showing significant changes over time within fatigue condition.

3.4. Block one

Four component centralities were already significantly different between the control and fatigue condition during the first block of scanning, including: Right parietal regions (p = 0.034); Left parietal regions (p = 0.031); Cerebellar regions (p = 0.013); Right insular region, specifically right amygdala (p = 0.049).

3.5. Block four

After fatigue was successfully induced in the experimental group, five components show a significant difference in centrality scores between control and experimental group during the last block of scanning. These regions include: Lateral central sulcus and surrounding gyri (p = 0.028); Lateral occipital cortex and superior parietal lobule (p = 0.010); Left insular cortex (p = 0.011); Inferior and middle frontal gyri (p = 0.009); Bilateral postcentral gyrus (p = 0.029).

3.6. Correction for multiple comparisons

Analysis on the false discovery rate revealed that all of the aforementioned effects are likely found by chance only, and no significant effects survived correcting for multiple comparisons.

3.7. Comparison between positive and signed centrality metrics

Comparing eigenvectors obtained from classical methods with the signed metric results in some interesting properties. Although it provides a similar split between networks of components, the signed metric has the additional property of switching signs over time. For reference, see figures 9a and 9b below. Figure 9a shows the eigenvectors computed based on non-negative matrices, whereas figure 9b shows the output of eigenvector decomposition of the signed correlation matrices.

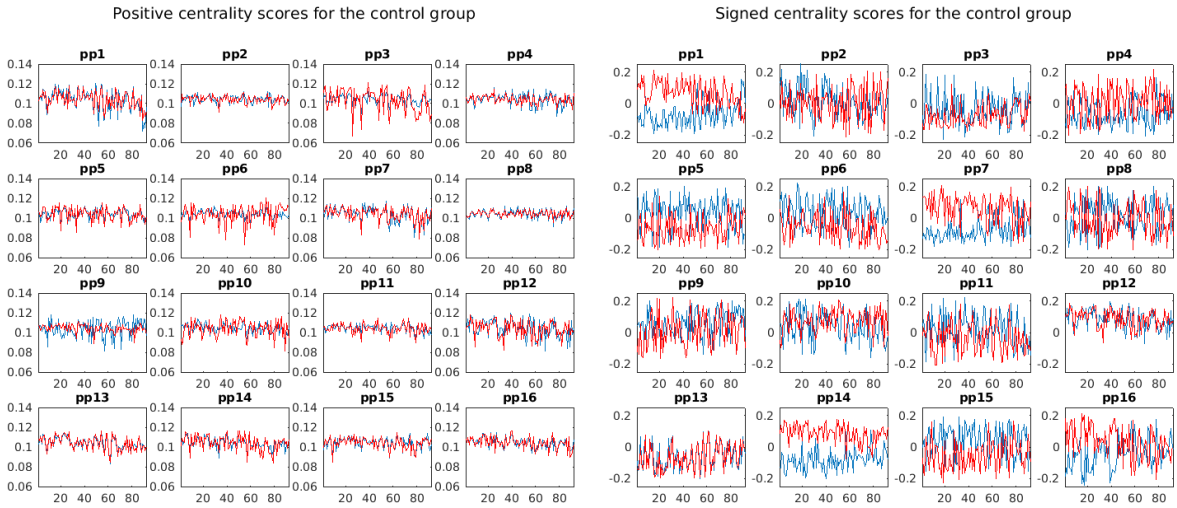

Figure 9a (left) and 9b (right). Centrality scores for block one (blue) and block four (red) for all participants in the control group. 9a shows centrality scores based on the transformed non-negative matrix, whereas 9b shows centrality scores based on the original signed matrices.

Although it appears the signed metric is not reliable for measuring the temporal dynamics of network interactions, we would like to mention one interesting result. For one random participant (Selected through the "lazy selection algorithm", i.e., the first participant), the ten components showing the highest positive centrality and the ten showing the highest negative centrality have been investigated and mapped. See table 1. and 2. for the major constituents of each of these networks, and the overlap between the components and the earlier described inhibition and facilitation systems.

Positive components and their constituents of pp1	**Overlap with proposed inhibition system**	**Overlap with proposed facilitation system**
DLPFC, broader frontal region	-	DLPFC
ACC, orbitofrontal cortex, caudate, nucleus accumbens	ACC	ACC, Orbitofrontal cortex, Caudate, Nucleus accumbens
Hippocampus, Amygdala	-	Hippocampus, Amygdala

Amygdala, Hippocampus	-	Amygdala, Hippocampus
Temporal fusiform cortex, Middle temporal gyrus, Hippocampus	-	Hippocampus
Temporal fusiform cortex, Amygdala, Hippocampus	-	Amygdala, Hippocampus
Temporal Pole, Superior temporal gyrus, Middle temporal gyrus, Frontal pole	-	
Hippocampus, Lingual gyrus, Temporal occipital fusiform cortex, Precuneus cortex	-	Hippocampus
Amygdala, Hippocampus, DLPFC	-	Amygdala, Hippocampus, DLPFC
Middle frontal gyrus, ACC, inferior frontal gyrus, Postcentral gyrus	ACC	DLPFC, ACC
Total component overlap:	**2**	**18**

Table 1. Ten components showing the highest positive centrality scores, subdivided into the regions that showed the highest intensity in the spatial maps. Column two and three denote the overlap between the observed networks and the theoretical systems proposed by Ishii, Tanaka, and Watanabe (2014).

Negative components and their constituents of pp1	**Overlap with proposed inhibition system**	**Overlap with proposed facilitation system**
Frontal pole	-	-
PCC, secondary somatosensory cortex, Angular gyrus	PCC, SSC	-
Thalamus, ACC	Thalamus, ACC	ACC
Insular cortex, Superior temporal gyrus, Inferior frontal gyrus	Insular cortex	-
ACC, DLPFC, Secondary somatosensory cortex, Central opercular cortex	ACC, SSC	ACC, DLPFC
ACC, PCC, superior frontal gyrus, frontomedian wall	ACC, PCC	ACC
MCC, SMA, Precuneus	SMA	SMA
Putamen, Heschl's gyrus, Insular cortex, Secondary somatosensory cortex	Insular cortex, SSC	Putamen
Superior temporal gyrus, Secondary somatosensory cortex, Insular cortex, Heschl's gyrus	SSC, Insular cortex	-
Visual areas surrounding parieto-occipital sulcus	-	-
Total component overlap:	**13**	**6**

Table 2. Ten components showing the highest negative centrality scores, subdivided into the regions that showed the highest intensity in the spatial maps. Column two and three denote the overlap between the observed networks and the theoretical systems proposed by Ishii, Tanaka, and Watanabe (2014).

4. Discussion

4.1. Fatigue effects

Although striking differences in changes over time between the two conditions have been found, none of these effects survive corrections for multiple comparisons. However, the areas that were found to change in a different manner over time when comparing control and fatigue conditions are interesting nonetheless. They include Lateral Central Sulcus and surrounding gyri; Frontomedian areas, including Anterior Cingulate Cortex; Posterior Cingulate Cortex; Right Anteriofrontal areas; and Right Insular Cortex. These areas do not contain the full set of areas where we expected to find changes over time, but those areas that belong to both observed results, and predicted anatomical sites, do show changes in the direction that was expected based on previous literature. The notable exception to this statement is the region containing the right sided anteriofrontal areas, whose difference in centrality over time was best explained by a drop of centrality in the control group, relative to a stable centrality for the two time points of the fatigue group. Due to the lack of statistical power for these effects, we cannot safely discuss the implications or interpretations of these findings, rather only remarking upon the peculiar fit there seems to be with the current findings and those proposed throughout the literature. The only exception is a tentative and speculative suggestion of the roles of the inhibition and facilitation systems. Assuming that these systems indeed exist as coherent identities, and assuming that the results reported above are indeed real effects instead of random fluctuations in the data, we might suggest the following additional notes on the theory proposed by Ishii et al. (2014).

Whenever an individual grows fatigued, it is of paramount importance that the remaining resources are distributed efficiently over the full system. This means inhibiting those regions or parts of a network that are not absolutely necessary to maintain proper task performance in the current context. I.e., the inhibition system and its auxiliaries become more active in order to reduce the activation of as many non-critical functions as possible. Coinciding with this activation of the inhibition system, the facilitation system is activated to ensure the proper workings of those regions or functions which are paramount to the current context, or task at hand. Restructuring the distribution of energy from small doses of energy to a lot of functions, to a larger dose of energy to only a few critical systems is, in itself, already quite a daunting task for any manager. Therefore, we would expect to see that the regions thought to belong to the facilitation system become highly influential in the domain, whilst it is working on this difficult task. Subsequently, when we see increased global inhibition, and increased local facilitation, we also expect to see an increase in faulty outcomes of these few systems still active. For example, as in any company structure, it is the safe option to have multiple individuals go over important documents, to ensure that no minor mistakes are made by any single individual. This means that all individuals perform less work per unit (the unit being one individual), but the total amount of energy or effort required to deliver the end product is likely higher than when the entire document is made, reviewed, corrected, and sent by only one individual. I.e., in terms of energy it might be more efficient to put only one individual on the project at hand, but the likelihood of making errors will be higher. Translating this back into the language of cortical structures, whilst retaining the essence of the message, and we do not find it a striking result that increasing inhibition and facilitation at the same time leads to

an increased efficiency in terms of resources or energy spent by the cortex as a whole, whilst allowing minor lapses of perfection to exit the cortical system as potential output is processed by fewer cortical subsystems.

Within the fatigue group, changes were seen in several sites which partially overlap with the sites contained in component set showing significant interaction effects. Centrality of the orbitofrontal cortex increased with time on task. This region is associated with emotion and reward in the cognitive processing of decision making. Secondly, it is known to receive information through projections from the magnocellular, medial nucleus of the mediodorsal thalamus. The magnocellular pathways process mainly broader features of a input field. This can potentially be linked to the evidence that global versus local processing is handled differently after mental fatigue has developed. Van der Linden and Eling (2006) report that local processing is compromised as a function of time on task, and we observe an increase in centrality of a system that receives mainly global information. The orbitofrontal cortex is also one of the candidate players of the proposed facilitation system. The hippocampus is another, and there too, we have observed an increase in centrality over time within the fatigue condition. But more interestingly, the PCC, thought to be part of the inhibition system, also shows an increase in centrality over time. This means three out of four observed effects are part of the systems where we expected to see the largest changes. However, since none of these effects survive corrections for multiple comparisons, it must be left as just an interesting note for the reader, who is also referred to the earlier paragraph on potential interpretation of an increase in both facilitation and inhibition systems.

We suggest that the changes in correlations between components, induced by states of mental fatigue, might not be large enough and do not carry a sufficient weight in their global influence. Stated otherwise, although a difference in flow of information is apparent between the first and last block of testing, this 'rewiring' leads to very subtle changes on the global scale of cortical networks. This leads to difficulties in interpretation, as we cannot state with certainty whether the effects of mental fatigue are nonexistent, or whether the centrality measure used to identify changes in global scale influencing is not sensitive enough to pick up these differences.

4.2. Signed centrality

Due to the changing signs in the signed eigenvector centrality computation, problems arise with the analysis of possible changes over time. Some participants show changes in all components, leading to difficulties in deciding whether the sign changed or not. Therefore, it was deemed an insufficient addition towards the current investigation. Until it is clear why the sign changes over time, and corrections based on a rigorous method surrounding the signed metric have been worked out, the measure of signed centrality does not suffice for data with a temporal dimension. Regardless of the actual statistics, and remaining problems with the signed eigenvector metric, those components showing a strong negative centrality for one of the participants show a remarkable overlap with the theoretical inhibition system as proposed by Ishii et al. (2014), just as the most central components on the positive side of the scale show remarkable overlap with the proposed facilitation system. This is reminiscent of the work performed by Bonacich and Lloyd (2004), where negative centrality was only observed for the monks that were disliked by many of their peers, and liked by very few of their peers.

Thus, we suggest that the signed centrality metric (and by translation, the other centrality metrics), does indeed have the capacity to not only identify the amount of influence of a single participant or component of a network, but additionally, provides information on the *direction* of the influence. This might provide a valuable source of information in future studies, when the observations need to be classified into positively and negatively influential players.

4.3. Limitations and suggestions for improvement

A possible shortcoming of the current design is the set-up of the control group. Although the task itself has been shown to induce a state of mental fatigue, fatigue effects can already be measured after only half an hour (Shigihara et al., 2013; Tanaka et al., 2014;). The control group watched a thoroughly non-stimulating nature documentary for the second and third block, and resumed the flanker task for the last block of scanning. This led to the conclusion that the fatigued participants, when compared to the control group, did not show significant differences in global activation patterns between the first and last block of scanning. One possible explanation for these findings is that the control group, having been in the scanner for two hours, might have shown similar symptoms of being fatigued as the experimental group, albeit to a lesser degree. A reason for considering this possibility is that the scanner room is thoroughly darkened, during the second and third blocks of scanning, the control group participants were instructed to stay still, yet not fall asleep. It could well be that for some participants watching the nature documentary with the explicit instructions of not falling asleep might actually be thoroughly taxing. Lastly, it must be kept in mind that fatigue is not a pure binary state, but rather a continuum. We suspect that during these two hours of scanning, all participants showed some relative motion on that continuum towards the side of fatigue.

One potential interesting investigation for future studies is found in the process of handling the output of matrix decompositions. After the set of eigenvalues and eigenvectors has been obtained, the question remains whether the expected changes occurring in conditions of mental fatigue are large enough to be part of the vector which explains the highest degree of variance in the data, or whether the changes have a more subtle nature. Taking into account that the leading eigenvector is also the principal component explaining the largest percentage of variance in the data, it could well be possible that it does not include the potentially minute changes which could be classified as fatigue related changes in connectivity. As a natural extension of leading eigenvector inspection, all eigenvectors up to a certain percentage of variance explained might be of interest in future studies.

References

Abdou, A. M., Higashiguchi, S., Horie, K., Kim, M., Hatta, H., & Yokogoshi, H. (2006). Relaxation and immunity enhancement effects of γ-aminobutyric acid (GABA) administration in humans. *Biofactors*, *26*(3), 201-208.

Alho, K. (1995). Cerebral generators of mismatch negativity (MMN) and its magnetic counterpart (MMNm) elicited by sound changes. *Ear and hearing*, *16*(1), 38-51.

Bagler, G. (2008). Analysis of the airport network of India as a complex weighted network. *Physica A: Statistical Mechanics and its Applications*, *387*(12), 2972-2980.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, 289-300.

Boksem, M. A., Meijman, T. F., & Lorist, M. M. (2005). Effects of mental fatigue on attention: an ERP study. *Cognitive brain research*, *25*(1), 107-116.

Bonacich, P., & Lloyd, P. (2004). Calculating status with negative relations. *Social networks*, *26*(4), 331-338.

Borgatti, S. P. (1995). Centrality and AIDS. *connections*, *18*(1), 112-114.

Cook, D. B., O'Connor, P. J., Lange, G., & Steffener, J. (2007). Functional neuroimaging correlates of mental fatigue induced by cognition among chronic fatigue syndrome patients and controls. *Neuroimage*, *36*(1), 108-122.

Da Silva, F. L., Van Rotterdam, A., Van Leeuwen, W. S., & Tielen, A. M. (1970). Dynamic characteristics of visual evoked potentials in the dog. II. Beta frequency selectivity in evoked potentials and background activity. *Electroencephalography and clinical neurophysiology*, *29*(3), 260-268.

Datta, A., Cusack, R., Hawkins, K., Heutink, J., Rorden, C., Robertson, I. H., & Manly, T. (2007). The p300 as a marker of waning attention and error propensity. *Computational intelligence and neuroscience*, *2007*.

Donchin, E., & Coles, M. G. (1988). Is the P300 component a manifestation of context updating?. *Behavioral and brain sciences*, *11*(3), 357-374.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Attention, Perception, & Psychophysics*, *16*(1), 143-149.

Estrada, E. (1996). Spectral moments of the edge adjacency matrix in molecular graphs. 1. Definition and applications to the prediction of physical properties of alkanes. *Journal of chemical information and computer sciences*, *36*(4), 844-849.

Faber, L. G., Maurits, N. M., & Lorist, M. M. (2012). Mental fatigue affects visual selective attention. *PloS one*, *7*(10), e48073.

Fan, X., Zhou, Q., Liu, Z., & Xie, F. (2015). Electroencephalogram assessment of mental fatigue in visual search. *Bio-medical materials and engineering*, *26*(s1), S1455-S1463.

Feshchenko, V. A., Veselis, R. A., & Reinsel, R. A. (1997). Comparison of the EEG effects of midazolam, thiopental, and propofol: the role of underlying oscillatory systems. *Neuropsychobiology*, *35*(4), 211-220.

Freeman, L. C., Borgatti, S. P., & White, D. R. (1991). Centrality in valued graphs: A measure of betweenness based on network flow. *Social networks*, *13*(2), 141-154.

Godey, B., Schwartz, D., De Graaf, J. B., Chauvel, P., & Liegeois-Chauvel, C. (2001). Neuromagnetic source localization of auditory evoked fields and intracerebral evoked potentials: a comparison of data in the same patients. *Clinical neurophysiology*, *112*(10), 1850-1859.

Grandjean, E. (1979). Fatigue in industry. *Occupational and Environmental Medicine*, *36*(3), 175-186.

Griffanti, L., Douaud, G., Bijsterbosch, J., Evangelisti, S., Alfaro-Almagro, F., Glasser, M. F., ... & Beckmann, C. F. (2016). Hand classification of fMRI ICA noise components. *NeuroImage*.

Güntekin, B., & Basar, E. (2007). Emotional face expressions are differentiated with brain oscillations. *International Journal of Psychophysiology*, *64*(1), 91-100.

Güntekin, B., & Başar, E. (2010). Event-related beta oscillations are affected by emotional eliciting stimuli. *Neuroscience letters*, *483*(3), 173-178.

Guo, Z., Chen, R., Zhang, K., Pan, Y., & Wu, J. (2016). The Impairing Effect of Mental Fatigue on Visual Sustained Attention under Monotonous Multi-Object Visual Attention Task in Long Durations: An Event-Related Potential Based Study. *PloS one*, *11*(9), e0163360.

Holroyd, C. B., Nieuwenhuis, S., Mars, R. B., & Coles, M. G. (2004). Anterior cingulate cortex, selection for action, and error processing. *Cognitive neuroscience of attention*, 219-231.

Hopstaken, J. F., van der Linden, D., Bakker, A. B., & Kompier, M. A. (2015). The window of my eyes: Task disengagement and mental fatigue covary with pupil dynamics. *Biological psychology*, *110*, 100-106.

Hopstaken, J. F., van der Linden, D., Bakker, A. B., Kompier, M. A., & Leung, Y. K. (2016). Shifts in attention during mental fatigue: Evidence from subjective, behavioral, physiological, and eye-tracking data. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(6), 878.

Hosseinmardi, H., Han, R., Lv, Q., Mishra, S., & Ghasemianlangroodi, A. (2014). Analyzing negative user behavior in a semi-anonymous social network. *CoRR abs*, *1404*.

Huster, R. J., Enriquez-Geppert, S., Lavallee, C. F., Falkenstein, M., & Herrmann, C. S. (2013). Electroencephalography of response inhibition tasks: functional networks and cognitive contributions. *International Journal of Psychophysiology*, *87*(3), 217-233.

Ishii, A., Tanaka, M., & Watanabe, Y. (2014). Neural mechanisms of mental fatigue. *Reviews in the Neurosciences*, *25*(4), 469-479.

Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, *17*(2), 825-841.

Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical image analysis*, *5*(2), 143-156.

Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain research reviews*, *29*(2), 169-195.

Kressner, D. (2005). The QZ Algorithm. *Numerical Methods for General and Structured Eigenvalue Problems*, 67-111.

Li, J., Song, G., & Miao, D. (2016). Effect of mental fatigue on nonattention: a visual mismatch negativity study. *NeuroReport*, *27*(18), 1323-1330.

Li, Y., Wang, L. Q., & Hu, Y. (2009). Localizing P300 generators in high-density event-related potential with fMRI. *Medical Science Monitor*, *15*(3), MT47-MT53.

Lohmann, G., Margulies, D. S., Horstmann, A., Pleger, B., Lepsien, J., Goldhahn, D., ... & Turner, R. (2010). Eigenvector centrality mapping for analyzing connectivity patterns in fMRI data of the human brain. *PloS one*, *5*(4), e10232.

Lorist, M. M., Bezdan, E., ten Caat, M., Span, M. M., Roerdink, J. B., & Maurits, N. M. (2009). The influence of mental fatigue and motivation on neural network dynamics; an EEG coherence study. *Brain Research*, *1270*, 95-106.

Lorist, M. M., Boksem, M. A., & Ridderinkhof, K. R. (2005). Impaired cognitive control and reduced cingulate activity during mental fatigue. *Cognitive Brain Research*, *24*(2), 199-205.

Möckel, T., Beste, C., & Wascher, E. (2015). The effects of time on task in response selection-an erp study of mental fatigue. *Scientific reports*, *5*.

Nakagawa, S., Sugiura, M., Akitsuki, Y., Hosseini, S. H., Kotozaki, Y., Miyauchi, C. M., ... & Kawashima, R. (2013). Compensatory effort parallels midbrain deactivation during mental fatigue: an fMRI study. *PLoS One*, *8*(2), e56606.

Newman, M. E. (2008). The mathematics of networks. *The new palgrave encyclopedia of economics*, *2*(2008), 1-12.

Olsen RW, Betz H (2006). "GABA and glycine". In Siegel GJ, Albers RW, Brady S, Price DD. *Basic Neurochemistry: Molecular, Cellular and Medical Aspects (7th ed.)*. Elsevier. pp. 291–302. ISBN 0-12-088397-X.

Olsen, R. W., & DeLorey, T. M. (1999). GABA and glycine. *Basic neurochemistry: molecular, cellular and medical aspects*, *6*.

Palva, S., & Palva, J. M. (2012). Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. *Trends in cognitive sciences*, *16*(4), 219-230.

Pevny, T., Bas, P., & Fridrich, J. (2010). Steganalysis by subtractive pixel adjacency matrix. *IEEE Transactions on information Forensics and Security*, *5*(2), 215-224.

Pfurtscheller, G., Stancak, A., & Neuper, C. (1996). Post-movement beta synchronization. A correlate of an idling motor area?. *Electroencephalography and clinical neurophysiology*, *98*(4), 281-293.

Preti, M. G., Bolton, T. A., & Van De Ville, D. (2016). The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage*.

Rosenkopf, L., & Schilling, M. A. (2007). Comparing alliance network structure across industries: observations and explanations. *Strategic Entrepreneurship Journal*, *1*(3-4), 191-209.

Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, *52*(3), 1059-1069.

Scheffers, M. K., Humphrey, D. G., Stanny, R. R., Kramer, A. F., & Coles, M. G. (1999). Error-related processing during a period of extended wakefulness. *Psychophysiology*, *36*(2), 149-157.

Seidel, J. J. (1968). Strongly regular graphs with (—1, 1, 0) adjacency matrix having eigenvalue 3. *Linear algebra and its Applications*, *1*(2), 281-298.

Sherry, P., & Philbrick, K. (2016). A44-Models of Accident Risk and Fatigue in Railroad Operations. *Journal of Transport & Health*, *3*(2), S30-S31.

Shigihara, Y., Tanaka, M., Ishii, A., Kanai, E., Funakura, M., & Watanabe, Y. (2013). Two types of mental fatigue affect spontaneous oscillatory brain activities in different ways. *Behavioral and Brain Functions*, *9*(1), 2.

Simulink, M., & Natick, M. A. (1993). The mathworks.

Sun, Y., Lim, J., Meng, J., Kwok, K., Thakor, N., & Bezerianos, A. (2014). Discriminative analysis of brain functional connectivity patterns for mental fatigue classification. *Annals of biomedical engineering*, *42*(10), 2084-2094.

Tanaka, M., Ishii, A., & Watanabe, Y. (2014). Neural effects of mental fatigue caused by continuous attention load: a magnetoencephalography study. *Brain research*, *1561*, 60-66.

Tanaka, M., Shigihara, Y., Ishii, A., Funakura, M., Kanai, E., & Watanabe, Y. (2012). Effect of mental fatigue on the central nervous system: an electroencephalography study. *Behavioral and brain functions*, *8*(1), 48.

Ullsperger, M., & Von Cramon, D. Y. (2001). Subprocesses of performance monitoring: a dissociation of error processing and response competition revealed by event-related fMRI and ERPs. *Neuroimage*, *14*(6), 1387-1401.

Van der Linden, D., & Eling, P. (2006). Mental fatigue disturbs local processing more than global processing. *Psychological research*, *70*(5), 395-402.

Van Noordt, S. J., & Segalowitz, S. J. (2012). Performance monitoring and the medial prefrontal cortex: a review of individual differences and context effects as a window on self-regulation. *Frontiers in human neuroscience*, *6*.

Wascher, E., Rasch, B., Sänger, J., Hoffmann, S., Schneider, D., Rinkenauer, G., ... & Gutberlet, I. (2014). Frontal theta activity reflects distinct aspects of mental fatigue. *Biological psychology*, *96*, 57-65.

Wessel, J. R. (2012). Error awareness and the error-related negativity: evaluating the first decade of evidence. *Frontiers in Human Neuroscience*, *6*.

Wink, A. M., de Munck, J. C., van der Werf, Y. D., van den Heuvel, O. A., & Barkhof, F. (2012). Fast eigenvector centrality mapping of voxel-wise connectivity in functional magnetic resonance imaging: implementation, validation, and interpretation. *Brain connectivity*, *2*(5), 265-274.

Woodruff, C. C., Daut, R., Brower, M., & Bragg, A. (2011). Electroencephalographic α-band and β-band correlates of perspective-taking and personal distress. *Neuroreport*, *22*(15), 744-748.

Wróbel, A. (2000). Beta activity: a carrier for visual attention. *Acta neurobiologiae experimentalis*, *60*(2), 247-260.

Xie, J., Xu, G., Wang, J., Li, M., Han, C., & Jia, Y. (2016). Effects of mental load and fatigue on steady-state evoked potential based brain computer interface tasks: a comparison of periodic flickering and motion-reversal based visual attention. *PloS one*, *11*(9), e0163426.

Zhao, C., Zhao, M., Liu, J., & Zheng, C. (2012). Electroencephalogram and electrocardiograph assessment of mental fatigue in a driving simulator. *Accident Analysis & Prevention*, *45*, 83-90.

Zhao, C., Zheng, C., Zhao, M., & Liu, J. (2010). Physiological assessment of driving mental fatigue using wavelet packet energy and random forests. *The American Journal of Biomedical Sciences*, *2*(3), 262-274.

Zhu, D. C., Zacks, R. T., & Slade, J. M. (2010). Brain activation during interference resolution in young and older adults: an fMRI study. *Neuroimage*, *50*(2), 810-817.

Zouridakis, G., Simos, P. G., & Papanicolaou, A. C. (1998). Multiple bilaterally asymmetric cortical sources account for the auditory N1m component. *Brain topography*, *10*(3), 183-189.